

University of Southampton Research Repository

Copyright © and Moral Rights for this thesis and, where applicable, any accompanying data are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis and the accompanying data cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content of the thesis and accompanying research data (where applicable) must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holder/s.

When referring to this thesis and any accompanying data, full bibliographic details must be given, e.g.

Thesis: Belal Asad (2025) "AI-Driven Resilience: Advancing DDoS Detection in Virtualised IoT Environments Through an FCNN-Based Framework", University of Southampton, Faculty of Engineering and Physical Sciences, School of Electronics and Computer Science, PhD Thesis, 216pp.

Data: Author (Year) Title. URI [dataset]

University of Southampton

Faculty of Engineering and Physical Sciences
School of Electronics and Computer Science

**AI-Driven Resilience: Advancing DDoS
Detection in Virtualised IoT
Environments Through an FCNN-Based
Framework**

by

Belal Asad

BSc, MSc

ORCID: [0000-0001-6310-0170](https://orcid.org/0000-0001-6310-0170)

*A thesis for the degree of
Doctor of Philosophy*

April 2026

University of Southampton

Abstract

Faculty of Engineering and Physical Sciences
School of Electronics and Computer Science

Doctor of Philosophy

**AI-Driven Resilience: Advancing DDoS Detection in Virtualised IoT
Environments Through an FCNN-Based Framework**

by Belal Asad

The rapid adoption of Internet of Things systems has accelerated the use of virtualisation technologies, including software-defined networking and cloud-based sensor platforms, to enhance scalability and manageability. However, these technologies introduce new security vulnerabilities, particularly distributed denial-of-service attacks, which pose significant risks to the stability and performance of virtualised IoT environments. This thesis addresses these challenges by proposing a novel AI-based DDoS detection framework specifically designed for virtualised IoT networks.

A lightweight five-layer fully connected neural network comprising 49,729 parameters is developed and evaluated using eight benchmark datasets, including a custom virtualised IoT testbed with 168 hours of baseline traffic characterisation. The framework integrates batch normalisation, adaptive dropout regularisation, Adam optimisation, and SMOTE-based class balancing to enhance training stability and generalisation. The experimental results demonstrate a mean detection accuracy of 99.85% across all datasets, achieving a peak accuracy of 99.92% on the virtualised IoT testbed. The proposed model significantly outperforms established approaches, including LSTM, Random Forest, and Support Vector Machine classifiers, with all improvements statistically significant ($p < 0.001$) and showing large effect sizes.

This research contributes to IoT security by quantifying virtualisation-specific vulnerabilities, demonstrating the effectiveness of lightweight AI-based detection in resource-constrained environments, and establishing experimentally validated scalability limits for deployment. The findings provide practical insights for designing efficient, adaptive DDoS detection systems for virtualised IoT infrastructures, while highlighting future directions in model adaptability, cross-domain evaluation, and privacy-preserving collaborative learning.

Contents

List of Figures	xvii
List of Tables	xxiii
Declaration of Authorship	xxvii
Acknowledgements	xxix
1 Introduction	1
1.1 Research Motivation	1
1.1.1 The Imperative of Smart City Security	2
1.1.2 The Research Gap	2
1.1.3 Principal Research Motivations	3
1.2 Research Justification	3
1.3 Research Questions and Objectives	4
1.3.1 Research Objectives	5
1.4 Research Contributions	6
1.4.1 Contribution 1: Quantitatively Validated AI Framework for vIoT DDoS Detection	6
1.4.2 Contribution 2: Novel Testbed Architecture with Reproducibility Protocols	6
1.4.3 Contribution 3: Cross-Model Validation and Contemporary Dataset Integration	7
1.5 Research Scope and Limitations	7
1.5.1 Application Context: Smart City IoT	8
1.5.2 Research Significance	8
1.6 Thesis Structure	8
2 Background and Literature Review	11
2.1 IoT and Virtualisation Fundamentals	12
2.1.1 IoT Architecture	12
2.1.2 Network Function Virtualisation	15
2.1.3 Software-Defined Networking	16
2.1.4 IoT Cloud Virtualisation	18
2.1.5 Security Challenges Overview	20
2.2 Security in Virtualised IoT	21
2.2.1 Security Fundamentals	21
2.2.2 SDN Vulnerabilities	22

2.2.3	NFV Vulnerabilities	24
2.2.4	DDoS Attacks in IoT	26
2.2.4.1	Characteristics and types of DDoS attack	26
2.2.4.2	IoT-Specific DDoS Considerations	26
2.2.5	Attack Detection Approaches	28
2.3	DDoS Detection Mechanisms	29
2.3.1	Systematic Review Methodology	29
2.3.2	Traditional vs Virtualised IoT Security	30
2.3.3	Defence Mechanisms Taxonomy	32
2.3.4	Machine Learning in Detection	34
2.3.4.1	Advantages of Machine Learning Approaches	34
2.3.4.2	Machine Learning Algorithms for Attack Detection	34
2.3.4.3	Comparative Performance of ML Algorithms	35
2.3.4.4	Deep Learning for Attack Detection	36
2.3.4.5	Challenges in ML Deployment for IoT	36
2.3.5	FCNN vs CNN Justification	37
2.4	Research Gap and Positioning	39
2.5	Summary	41
3	Experimental Methodology and Requirements	43
3.1	Research Design and Methodological Foundations	43
3.1.1	Research Philosophy and Epistemological Approach	43
3.1.2	Validation Strategy and Methodological Rigour	44
3.1.2.1	Internal Validation	44
3.1.2.2	External Validation	45
3.1.2.3	Reproducibility	45
3.2	System Requirements Analysis	46
3.2.1	Requirements Specification and Validation	46
3.2.1.1	Non-Functional Requirements	46
3.2.1.2	Functional Requirements	48
3.2.2	Requirements Selection Rationale	51
3.2.3	Requirements Prioritisation	52
3.2.4	Requirements Validation Summary	52
3.3	Experimental Design Principles	53
3.3.1	Testbed Design Philosophy	53
3.3.2	Baseline Characterisation Strategy	55
3.3.3	Multi-Dataset Evaluation Strategy	56
3.3.4	Experimental Scenarios Design Rationale	57
3.4	AI Framework Methodology	58
3.4.1	Model Architecture Selection Rationale	58
3.4.2	Layer Depth Determination: “Why 5 Layers?”	59
3.4.3	Neuron Configuration and Hierarchical Feature Learning	60
3.4.4	Training Strategy and Hyperparameter Selection	62
3.4.5	Regularisation Strategy	63
3.5	Evaluation Metrics and Statistical Validation	64
3.5.1	Performance Metrics Selection	64
3.5.2	Statistical Validation Methodology	65

3.6	Reproducibility Protocols	65
3.6.1	Documentation Requirements	65
3.6.2	Dataset Availability	66
3.6.3	Code Availability	66
3.7	Ethical Considerations	67
3.7.1	Data Privacy and Protection	67
3.7.2	Attack Isolation Protocols	67
4	Testbed Implementation and System Development	71
4.1	Virtualised IoT Testbed Deployment	73
4.1.1	Physical Infrastructure Setup	73
4.1.1.1	Host Computer Configuration	73
4.1.1.2	Raspberry Pi IoT Device Setup	74
4.1.1.3	Network Equipment and Topology	74
4.1.2	Virtualisation Layer Configuration	75
4.1.2.1	GNS3 Configuration	75
4.1.2.2	Virtual IoT Device Configuration	75
4.1.3	IoT Device Integration	76
4.1.3.1	Sensor Configuration	76
4.1.3.2	Data Transmission Protocol	77
4.1.4	Network Services and Orchestration	77
4.1.4.1	Data Logging Server	77
4.1.4.2	Network Monitoring Tools	78
4.1.4.3	Orchestration Scripts	78
4.1.5	Complete Testbed Architecture	79
4.1.5.1	Attack Traffic Entry Points and Flow Topology	80
4.2	Baseline Network Characterisation	83
4.2.1	Normal Traffic Patterns	83
4.2.1.1	Traffic Generation Methods	83
4.2.1.2	Monitoring Tools and Procedures	84
4.2.1.3	Data Collection Duration and Conditions	84
4.2.1.4	Baseline Traffic Characteristics	86
	Temporal Traffic Patterns	86
	Protocol Distribution	87
	Packet Size Characteristics	87
	Physical vs Virtualised Traffic Patterns	87
4.2.2	Resource Utilisation Under Benign Conditions	88
4.2.2.1	CPU and Memory Profiling	88
4.2.2.2	Network Bandwidth Utilisation	89
4.2.2.3	Virtual Machine Resource Overhead	90
4.2.3	Performance Thresholds	91
4.2.3.1	Latency and Throughput Measurements	91
4.2.3.2	Threshold Establishment and Validation	92
4.2.3.3	Physical IoT vs. Virtualised IoT Traffic Comparison	93
4.3	Attack Scenario Design and Dataset Generation	95
4.3.1	DDoS Attack Simulation Framework	95
4.3.1.1	Attack Tool Configuration	96

4.3.1.2	Attack Orchestration System	97
4.3.2	Attack Intensity Variations	98
4.3.3	Multiple Dataset Generation	102
4.3.3.1	CIC-DDoS2019	102
4.3.3.2	IoT-23	103
4.3.3.3	IoT-DH	104
4.3.3.4	SDN-DDoS	105
4.3.3.5	CICIoT2023	106
4.3.3.6	EdgeIIoTset	108
4.3.3.7	CICIoMT2024	111
4.3.3.8	Generated Virtualised IoT Dataset	115
4.3.3.9	Dataset Complementarity and Rationale	120
4.3.4	Data Labelling and Validation	124
4.3.4.1	Labelling Methodology	124
4.3.4.2	Cross-Validation with Multiple Sources	125
4.3.4.3	Inter-Rater Reliability	125
4.4	AI Model Training and Optimisation	127
4.4.1	FCNN Architecture Design	127
4.4.1.1	Architecture Design Rationale	127
	Modified Pyramidal Structure	127
	Strategic Dropout Regularisation	127
	Batch Normalisation Integration	129
	Activation Function Selection	129
4.4.1.2	Parameter Distribution and Computational Efficiency	129
4.4.2	Training Dataset Preparation	130
4.4.2.1	Train/Test Split Rationale	130
4.4.2.2	Data Balancing Strategies	132
4.4.3	Model Training Process	134
4.4.4	Hyperparameter Tuning	135
4.4.5	Model Convergence and Validation	136
4.4.5.1	Training Convergence Analysis	137
	Loss Convergence	138
	Accuracy Plateau	139
	Learning Rate Schedule	139
	Gradient Stability	139
	Reproducibility Validation	139
4.5	Competing Methods Implementation	140
4.5.1	Selection of State-of-the-Art Methods	140
4.5.1.1	Method 1: Random Forest-based Detection	141
4.5.1.2	Method 2: SVM-based Approach	141
4.5.1.3	Method 3: CNN-LSTM Hybrid Architecture	141
4.5.2	Implementation Details	142
4.5.3	Configuration for Fair Comparison	142
4.5.4	Validation of Implementations	142
4.6	Experimental Scenarios and Test Cases	142
4.6.1	Scenario Overview and Categories	143
4.6.2	Baseline Scenarios (S1–S4)	144

4.6.2.1	S1: Pure Benign Traffic	145
4.6.2.2	S2: Single SYN Flood Attack	146
4.6.2.3	S3: Sequential Multi-Type Attacks	147
4.6.2.4	S4: Multi-Vector Simultaneous Attack	148
4.6.3	Attack Intensity Variations (S5–S8)	149
4.6.3.1	S5: Very Low Intensity (Stealthy Attack)	150
4.6.3.2	S6: Gradual Ramp-Up Attack	151
4.6.3.3	S7: Pulsing Attack Pattern	152
4.6.3.4	S8: Sustained High-Intensity Attack	153
4.6.4	Network Congestion Scenarios (S9–S12)	154
4.6.4.1	S9: Legitimate Traffic Spike (Flash Crowd)	154
4.6.4.2	S10: High Network Load + Medium Attack	155
4.6.4.3	S11: Bandwidth Saturation (Near-Capacity)	156
4.6.4.4	S12: Partial Network Outage + Attack	157
4.6.5	Resource Constraint Scenarios (S13–S16)	158
4.6.5.1	S13: CPU Throttling (Limited Processing)	158
4.6.5.2	S14: Memory Limitation	159
4.6.5.3	S15: Concurrent Process Interference	160
4.6.5.4	S16: Combined Resource Stress (Extreme Constraint)	161
4.6.6	Edge Case Scenarios (S17–S20)	162
4.6.6.1	S17: Cold-Start Detection	162
4.6.6.2	S18: Intermittent Attack (On-Off Pattern)	163
4.6.6.3	S19: Zero-Day Attack Simulation	164
4.6.6.4	S20: Adversarial Evasion Attempt	165
4.6.7	Scenario Summary and Requirements Traceability	166
4.7	Summary	166
5	Results, Analysis and Discussion	171
5.1	Baseline Network Performance Results	172
5.1.1	Normal Traffic Behaviour	172
5.1.2	Resource Utilisation Patterns	174
5.1.3	Performance Benchmarks	175
5.1.4	Validation of Virtualised IoT Vulnerability Claims	176
5.2	AI Model Performance Evaluation	176
5.2.1	Detection Accuracy and Precision	177
	Statistical Reliability and Reproducibility.	177
	Exceptional Reproducibility and Statistical Reliability.	178
	Highest Performance on Generated vIoT Dataset.	179
	Performance Progression Across Dataset Chronology.	180
	Cross-Dataset Robustness and Generalisation.	181
	Comparative Performance Context.	182
	Computational Investment for Rigorous Validation.	183
5.2.1.1	CIC-DDoS2019 Dataset Results	183
5.2.1.2	IoT-23 Dataset Results	184
5.2.1.3	IoT Device Hack (IoT-DH) Dataset Results	186
5.2.1.4	SDN-DDoS Dataset Results	188
5.2.1.5	CICIoT2023 Dataset Results	189

5.2.1.6	EdgeIIoTset Dataset Results	191
5.2.1.7	CICIoMT2024 Dataset Results	193
5.2.1.8	Generated Virtualised IoT Dataset Results	196
5.2.2	Recall and F1-Score Analysis	197
5.2.3	ROC Curves and AUC Analysis	198
5.2.4	Confusion Matrices	199
5.2.5	Performance Under Varying Conditions	203
5.2.5.1	Different Attack Intensities	203
5.2.5.2	Network Congestion Scenarios	204
5.2.6	Edge Cases and Robustness Testing	205
5.2.7	False Positive and False Negative Analysis	206
	False Positives (Mean Rate: 0.19%, Range: 0.11–0.26%). . .	206
	False Negatives (Mean Rate: 0.10%, Range: 0.05–0.16%). .	206
5.3	Comparative Experimental Analysis	207
5.3.1	Direct Comparison with Implemented Competing Methods	208
5.3.1.1	Performance on Identical Datasets	208
5.3.1.2	Method-by-Method Comparison	210
	Deep Learning Approach (Multi-Layer LSTM with Atten- tion).	210
	Random Forest Ensemble (500 Trees, Optimised Hyperpa- rameters).	210
	Support Vector Machine (RBF Kernel, Grid-Searches Hy- perparameters).	211
5.3.1.3	Statistical Significance of Differences	211
5.3.2	Performance Benchmarking Against State-of-the-Art	212
5.3.2.1	Analysis of Relative Strengths	213
5.3.3	Computational Overhead Analysis	214
5.3.3.1	Training Time Comparison	214
5.3.3.2	Inference Time Comparison	215
5.3.3.3	Resource Consumption	216
5.3.4	Scalability Comparison	216
5.3.5	Interpretation of Comparative Results	217
5.4	Statistical Validation and Significance Testing	218
5.4.1	Confidence Intervals and Error Margins	218
5.4.2	Statistical Significance Testing	220
5.4.2.1	Paired T-Tests: Proposed Model vs Competing Methods	220
5.4.2.2	ANOVA: Performance Across Datasets	221
5.4.2.3	Non-Parametric Tests: Robustness to Distributional As- sumptions	222
5.4.3	Sensitivity Analysis	222
5.4.3.1	Hyperparameter Variation Impact	222
5.4.3.2	Robustness to Dataset Characteristics	223
5.4.4	Error Analysis and Failure Modes	225
5.4.4.1	False Positive Analysis (Mean Rate: 0.19%, Range: 0.11– 0.26%)	225
5.4.4.2	False Negative Analysis (Mean Rate: 0.10%, Range: 0.05– 0.16%)	225

5.4.4.3	Failure Pattern Identification	226
5.4.5	Repeatability and Consistency	227
5.4.5.1	Multiple Run Analysis	227
5.4.5.2	Mean Values and Variance Across Runs	228
5.5	Bottleneck Analysis and System Limits	230
5.5.1	Data Volume Impact on Detection Performance	230
5.5.1.1	Scalability Testing with Increasing Data Sizes	230
5.5.1.2	Performance Degradation Thresholds	232
5.5.2	User Scalability Analysis	232
5.5.2.1	Maximum Concurrent Users/Devices	233
5.5.2.2	Performance Under High Load	233
5.5.3	Bandwidth and Traffic Load Thresholds	235
5.5.3.1	Maximum Sustainable Traffic Rates	235
5.5.3.2	Network Saturation Points	236
5.5.4	Resource Exhaustion and System Degradation Points	237
5.5.4.1	CPU/Memory Bottlenecks	237
5.5.4.2	Breaking Points Identification	238
5.5.5	Distinguishing Attack from Legitimate Overload	240
5.5.5.1	Characteristics of Resource Exhaustion	240
	Traffic Pattern Differences (Temporal Characteristics).	240
	Statistical Feature Differences (Figure 5.11(c)).	242
	Detection Confidence Patterns (Figure 5.11(d)).	243
5.5.5.2	Detection Challenges at System Limits	244
5.6	Discussion: Addressing Research Questions	245
5.6.1	RQ1: Critical Vulnerabilities in Virtualised IoT – Findings and Insights	246
	Shared Resource Vulnerabilities (Validated through Sec- tion 5.5.4).	246
	SDN Controller Single Point of Failure (Validated through SDN-DDoS Dataset Results, Section 5.2.1.4).	247
	Virtualisation Layer Attack Surface Expansion (Quantified through Comparative Analysis, Section 5.1.4).	247
	Difficulty Distinguishing Attacks from Legitimate Behaviour (Addressed through Section 5.5.5).	248
5.6.2	RQ2: AI Mechanism Design and Effectiveness	249
	Architecture Optimisation for Virtualised Traffic (Validated through Multi-Dataset Evaluation, Section 5.2.1).	249
	Balance of Accuracy and Efficiency (Demonstrated through Computational Overhead Analysis, Section 5.3.3).	250
	Robustness Across Operating Conditions (Validated through Sensitivity Analysis, Section 5.4.3, and Varying Conditions Testing, Section 5.2.5).	251
	Incorporation of Multiple Security Principles (Architectural Integration of Defence-in-Depth, Validated through Conjecture C2).	252
5.6.3	RQ3: Performance Compared with Existing Approaches	253

	Detection Accuracy Improvements (Quantified through Direct Comparison, Section 5.3.1).	253
	Computational Efficiency Trade-offs (Balanced Performance Analysis, Section 5.3.3).	254
	Comparison with Contemporary Research (Positioned through State-of-the-Art Benchmarking, Section 5.3.2).	254
	Scalability Considerations (Operational Limits Established through Section 6.4.2 and Section 5.5).	255
5.6.4	Novel Contributions Validation	256
5.6.5	Implications for Virtualised IoT Security	257
5.6.6	Theoretical and Practical Insights	258
	Theoretical Insights.	258
	Practical Insights.	259
5.7	Limitations and Threats to Validity	260
5.7.1	Methodological Limitations	260
5.7.2	Generalisability Constraints	262
5.7.3	Testbed versus Real-World Deployment	263
5.7.4	Dataset Limitations	265
5.7.5	Model Limitations	266
6	Conclusions and Future Work	273
6.1	Research Overview and Key Findings	274
6.1.1	Research Approach and Methodology	274
6.1.2	Key Findings Summary	274
6.2	Research Contributions	276
6.2.1	Contribution 1: Comprehensive AI-Based Security Framework for Virtualised IoT with Empirical Vulnerability Quantification	276
	Theoretical contribution.	276
	Methodological contribution.	276
	Practical impact.	277
6.2.2	Contribution 2: Novel Virtualised IoT Testbed with Comprehensive Baseline Characterisation and Reproducible Experimental Protocol	277
	Architectural contribution.	278
	Baseline characterisation.	278
	Validation and reproducibility.	278
6.2.3	Contribution 3: Comprehensive Comparative Evaluation with Contemporary Datasets and Practical Deployment Validation	279
	Direct comparative evaluation.	279
	Contemporary dataset evaluation.	279
	Bottleneck analysis and deployment guidelines.	280
6.3	Limitations and Threats to Validity	281
6.3.1	Technical and Methodological Limitations	281
6.3.2	Generalisability and Scope Constraints	282
6.4	Future Research Directions	283
6.4.1	Enhanced Detection Mechanisms Addressing Current Limitations	283
6.4.2	Scalability and Distributed Architecture	284

6.4.3	Integration with Emerging Technologies and Standards	284
6.4.4	Cross-Domain Applications and Specialised Deployments	285
6.5	Concluding Remarks	286
Appendix A Smart City Framework and Context		289
Appendix A.1	Overview	289
Appendix A.2	I.1 Smart City Resource Domains	289
Appendix A.3	I.2 Smart City Service Classifications	290
Appendix A.4	I.3 Security Implications for Smart City IoT Infrastructures . .	292
Appendix A.4.1	I.3.1 Service Availability Requirements	292
Appendix A.4.2	I.3.2 Virtualisation in Smart City Deployments	292
Appendix A.4.3	I.3.3 Application Context for This Research	293
Appendix A.4.4	I.3.4 Relevance to Research Questions	293
Appendix A.5	I.4 Statistical Context: Urban Growth and IoT Deployment . .	293
Appendix A.6	I.5 Summary	294
Appendix B Additional IoT Architecture Details		295
Appendix B.1	IoT Architecture	295
Appendix B.2	Network Functions Virtualisation (NFV)	296
Appendix B.2.1	NFV Architecture	298
Appendix B.3	SDN - Software Defined Networks	299
Appendix B.3.1	SDN Architecture	302
Appendix B.4	IoT Cloud Concept	304
Appendix B.4.1	Virtualisation of IoT Environments	306
Appendix B.4.1.1	IoT Network Virtualisation	307
Appendix B.4.1.2	Challenges in Virtualised IoT Environments . .	310
Appendix B.5	Security Issues and Vulnerabilities	311
Appendix B.5.1	Security Unique Provisioned Challenges in Virtualised IoT	311
Appendix B.5.2	Threats and Vulnerabilities in SDN	313
Appendix B.5.2.1	SDN Controller Vulnerabilities	313
Appendix B.5.2.2	Communication Vulnerabilities	314
Appendix B.5.2.3	Network Device Vulnerabilities	314
Appendix B.5.3	Threats and Vulnerabilities in NFV	314
Appendix B.5.3.1	Resource Isolation	315
Appendix B.5.3.2	Management Plane Security	315
Appendix B.5.3.3	VNF Security	316
Appendix B.6	Recent Studies on DDoS Attacks Detection and Mitigation . .	316
Appendix B.6.1	Overview and Classification of DDoS Attacks	319
Appendix B.6.1.1	Overview of DDoS Attacks	319
Appendix B.6.1.2	Classification of DDoS Attacks	320
Appendix B.6.1.3	Mechanisms and Impacts of DDoS Attacks on IoT Ecosystem	322
Appendix B.7	Proposed Technologies and Security Solutions For Virtualised IoT	323
Appendix B.7.1	Existing DDoS Attack Detection Methods	324
Appendix B.7.1.1	Feature-based Detection Methods	324

Appendix B.7.1.2 Behavior-based Detection Methods	325
Appendix B.8 Applications of Machine Learning in Attack Detection	326
Appendix B.8.1 Learning Techniques	326
Appendix B.8.2 Types of ML Algorithms	327
Appendix B.8.3 Deep Learning and Its Role in Attack Detection	327
Appendix B.9 Security Problems Comparison Between virtualised IoT and Traditional IoT Network	329
Appendix B.10 Comparison of Different Vulnerabilities and Attacks in Virtu- alised IoT Environments.	330
Appendix B.11 Comparison of DDoS Attack Detection in Traditional IoT Net- works and virtualised WSNs	332
Appendix B.11.1 Network Architecture	333
Appendix B.11.2 DDoS Attack Detection Techniques	333
Appendix B.11.3 Scalability and Resource Management	334
Appendix B.11.4 Adaptability and Resilience	334
Appendix B.11.4.1 Comparative analysis of existing main DDoS defence and detection mechanisms	334
Appendix B.12 Types of DDoS Attacks Comparison Between virtualised WSNs and Traditional IoT Network	337
Appendix B.12.1 ACK Flood	341
Appendix B.12.2 SYN Flood	341
Appendix B.12.3 Slowloris	341
Appendix B.13 Applications of Deep Learning in Security Detection	342
Appendix B.13.1 Intrusion Detection in Cloud Computing	342
Appendix B.13.2 Detection of IoT-based DDoS Attacks	343
Appendix B.13.3 Mitigation in Software-Defined Networks (SDNs)	343
Appendix B.13.4 Multi-Class Attack Classification	343
Appendix B.13.5 Trustworthiness Assessment in Cloud Services	343
Appendix B.13.6 Real-time Anomaly Detection	344
Appendix B.13.7 Preventative Measures through Behaviour Modelling	344
Appendix B.14 IoT NIDSs Based on Machine Learning Techniques	344
Appendix B.14.1 Multi-Layer Perceptron (MLP) for IoT NIDS	345
Appendix B.14.2 Host-Based IDS Framework (IoT-IDM)	345
Appendix B.14.3 Artificial Immune Systems (AIS) with Fog Computing	345
Appendix B.14.4 Hybrid Anomaly and Specification-Based IDS	345
Appendix B.14.5 Dimensionality Reduction and Classification for IoT IDS	345
Appendix B.14.6 Conditional Variational Autoencoder (CVAE) for IoT NIDS	346
Appendix B.14.7 Fog Computing with Deep Learning	346
Appendix B.14.8 Online Sequential Extreme Learning Machine (OS-ELM)	346
Appendix B.14.9 Semi-Supervised Learning for IoT NIDS	346
Appendix B.14.10 Ensemble-Based NIDS for IoT	347
Appendix B.15 Comparison of Machine Learning Algorithms Used in Attack Detection	347
Appendix B.16 Comparison of CNN and FCNN	349
Appendix B.16.1 Fully Connected Neural Network (FCNN)	349
Appendix B.16.2 Convolutional Neural Network (CNN)	350

Appendix B.16.3 Comparison and Application in Security Detection . . .	351
Appendix B.16.3.1 Applications of FCNN in Anomaly Detection .	352
Appendix C System Requirements Analysis	355
Appendix C.0.1 Non-Functional Requirements	355
Appendix C.0.2 Functional Requirements	358
References	363

List of Figures

2.1	IoT Architecture	14
2.2	Network Function Virtualisation Architecture (This figure has been re-sketched from ETSI (2013))	16
2.3	Software-Defined Networking Architecture (This figure has been re-sketched from Alam et al. (2019))	18
2.4	Cloud-based Virtualised Wireless Sensor Network Architecture (This figure has been re-sketched from Alam et al. (2019))	20
2.5	Security Vulnerabilities in Software-Defined Networks	24
2.6	PRISMA-style flow diagram of study selection.	31
2.7	Taxonomy of DDoS Defence Mechanisms in Virtualised IoT	33
2.8	Comparison of CNN and FCNN for DDoS Attack Detection	38
4.1	Hybrid virtualised IoT testbed combining physical and virtual devices via an SDN-enabled GNS3 environment. Multiple coordinated attack vectors generate DDoS traffic, while centralised SDN monitoring enables traffic capture and machine learning analysis.	80
4.2	Normal traffic during the 168-hour baseline shows clear diurnal variation (higher daytime rates), dominance of HTTP/HTTPS traffic, moderate packet sizes, and more stable inter-arrival times for physical devices than virtualised ones. These characteristics define normal operating thresholds for subsequent anomaly detection.	86
4.3	Physical versus virtualised IoT traffic comparison shows consistently superior and more stable performance for physical devices, with lower variability, jitter, resource usage, and energy consumption. Virtualised IoT exhibits significantly higher variability due to hypervisor overhead, confirming the need for device-specific baselines and adaptive detection thresholds ($p < 0.001$).	94

4.4	Four Attack Intensity Variation Patterns Employed in Experimental Scenarios. (a) Gradual Ramp-Up increases linearly from 100 pps to 50,000 pps over 30 minutes (+2,746 pps/minute) to evaluate detection threshold and response latency; (b) Pulsing Attack alternates between high-intensity (40,000 pps, 5 minutes) and low-intensity (500 pps, 3 minutes) phases across 7.5 complete cycles to test adaptation to dynamic patterns; (c) Random Variation maintains mean 10,000 pps with Gaussian noise ($\sigma=3,000$ pps, 1K–30K pps bounds) sampled every 10 seconds to simulate realistic attacker behaviour; (d) Sustained High-Intensity delivers constant 75,000 pps for 15 minutes to stress-test detection latency and system stability under extreme conditions. These patterns implement the experimental scenarios design rationale (Section 3.3.4), enabling comprehensive evaluation across intensity spectrum from stealthy low-rate (200 pps) to extreme saturation (>100K pps).	99
4.5	Multi-Vector Attack Scenarios Combining Multiple Simultaneous Attack Types. (a) SYN+UDP Flood delivers 30,000 pps SYN flood targeting TCP stack whilst simultaneously launching 20,000 pps UDP flood saturating bandwidth (total 50K pps); (b) ICMP+HTTP Flood combines 10,000 pps ICMP flood with 200 concurrent HTTP connections (Golden-Eye) to stress both network and application layers; (c) Triple Attack executes coordinated SYN (20K pps), UDP (15K pps), and ICMP (10K pps) floods totalling 45K pps across three protocols; (d) Application+Network attack merges HTTP flood (application layer) with UDP flood (network layer) to challenge layer-specific detection mechanisms. Multi-vector scenarios validate the model’s ability to detect complex coordinated attacks where individual attack components may appear below detection thresholds but collective impact causes service degradation, addressing the Robustness requirement (Appendix C).	100
4.6	Dataset Complementarity Matrix Across Eight Evaluation Datasets Spanning 2019–2025. Heat map intensity indicates coverage level (dark green=excellent, green=good, yellow=partial, gray=none) across ten critical dimensions. CIC-DDoS2019 provides broad DDoS attack diversity (12 types) but lacks IoT-specific characteristics; IoT-23 contributes authentic IoT malware traffic (Mirai, Torii) from real devices; EdgeIIoTset introduces edge computing architecture with three-tier topology; CICIoMT2024 uniquely provides healthcare-critical scenarios with patient safety labels and medical protocols (HL7, FHIR, DICOM); Generated vIoT dataset (2025) exclusively captures virtualisation-specific metrics (VM CPU, hypervisor scheduling, context switches) absent from all public datasets. The complementary coverage across 266.4M flows and 226.1 GB ensures robust validation spanning traditional networks, IoT-specific environments, edge architectures, life-critical healthcare infrastructure, and virtualised deployments, implementing the multi-dataset evaluation strategy (Section 3.3.3).	121

- 4.7 Five-Layer Fully Connected Neural Network (FCNN) Architecture Implementing Design Selected Through Systematic Architecture Search (Table 3.4). The modified pyramidal structure (256→128→64→64→32) incorporates intentional plateau at 64 neurons (Layers 3–4) enabling deeper feature refinement without aggressive dimensionality reduction. Each hidden layer comprises: (1) Dense connections with ReLU activation introducing non-linearity; (2) Batch normalisation stabilising training through zero-mean, unit-variance normalisation; (3) Variable dropout rates (60%/50%/55%/40%/50%) implementing strategic regularisation aggressive early dropout (60%) prevents overfitting on high-dimensional input, moderate plateau dropout (55%, 40%) retains learned representations, moderate output dropout (50%) ensures robust final predictions. Total 49,729 trainable parameters distributed as: Layer 1 (21,760), Layer 2 (32,896), Layer 3 (8,256), Layer 4 (4,160), Layer 5 (2,080), Output (33). The architecture achieves optimal balance between capacity and efficiency, delivering 0.34 ms inference time whilst outperforming both shallower (3–4 layers: F1=0.921–0.964) and deeper (6–7 layers: F1=0.979–0.983) alternatives in empirical validation (Table 3.4). 128
- 4.8 FCNN Training Convergence and Validation Monitoring Over 100 Epochs. (a) Training loss (blue) and validation loss (red) both converge smoothly without oscillation, with early stopping triggered at epoch 42 when validation loss ceased decreasing for 10 consecutive epochs; minimal train-validation gap (0.56%) indicates negligible overfitting; (b) Training accuracy (blue) and validation accuracy (red) plateau at 98.45% and 97.89% respectively, demonstrating strong generalisation; (c) Learning rate schedule implements three-phase strategy: warm-up (epochs 1–5: linear $10^{-6} \rightarrow 10^{-5}$), constant (epochs 6–30: fixed 10^{-5}), exponential decay (epochs 31+: $\times 0.95$ every 5 epochs, floor 10^{-7}) optimising convergence speed and final performance; (d) Gradient norms remain stable (0.1–1.0 range) throughout training without explosion (>10) or vanishing (<0.001), confirming batch normalisation effectiveness. Convergence achieved at epoch 42 ± 3 across 10 independent training runs (random seed variation), validating reproducibility and demonstrating robust training procedure. . . . 138
- 5.1 Baseline Traffic Characterisation Metrics Over 168 Hours. (a) Bandwidth utilisation shows diurnal pattern (mean 1.2 MB/s, std 0.3 MB/s) with peak-hour elevation and intermittent data transfer spikes; (b) Flow duration follows log-normal distribution (median 45s) dominated by short sensor updates (<10 s, 52.3%); (c) Port distribution reveals HTTPS/HTTP dominance (57.3%) alongside IoT protocols (MQTT 21.2%, CoAP 6.3%); (d) Shannon entropy remains stable (mean 4.0 bits, std 0.3 bits) within normal range, confirming regular operational behaviour. Metrics establish baseline parameters for $\mu \pm k\sigma$ anomaly detection thresholds. 173

5.2	Resource Utilisation Patterns Over 24-Hour Period Under Normal Load (Representative Sample from 168-Hour Baseline). (a) CPU utilisation ranging from 20% baseline during off-peak hours to 65% peak during operational periods, mean 38.5% ($\sigma = 12.3\%$), demonstrating substantial headroom for handling traffic surges and detection mechanism overhead. (b) Memory utilisation maintaining stable mean 47.2% ($\sigma = 4.1\%$), reflecting consistent VM allocation. (c) Network bandwidth utilisation peaking at 312.4 Kbps during operational hours, mean 258.7 Kbps representing 0.026% of 1 Gbps capacity. (d) Disk I/O operations averaging 1,250 IOPS with peaks at 2,000 IOPS.	175
5.3	Comprehensive Performance Metrics Comparison Across All Datasets	197
5.4	ROC Curves and AUC Analysis for All Evaluated Datasets	199
5.5	Confusion Matrices for Datasets 1–4 (Mean Results from 5 Independent Runs). Each matrix shows true negatives (TN), false positives (FP), false negatives (FN), and true positives (TP) based on 10,000 samples per class. Blue intensity indicates magnitude of values. All datasets demonstrate exceptional performance with true negative rates 99.79–99.85% (mean 99.82%) and true positive rates 99.84–99.90% (mean 99.87%), whilst maintaining very low false positive rates 0.15–0.21% (mean 0.18%) and false negative rates 0.10–0.16% (mean 0.13%).	200
5.6	Confusion Matrices for Datasets 5–8 (Mean Results from 5 Independent Runs). Generated vIoT dataset (bottom-right) achieves best performance with true negative rate 99.89% (9,989/10,000 benign correctly identified) and true positive rate 99.95% (9,995/10,000 attacks detected), whilst maintaining exceptionally low false positive rate 0.11% (only 11 false alarms per 10,000 samples) and false negative rate 0.05% (only 5 missed attacks per 10,000 samples), resulting in overall accuracy $99.92 \pm 0.007\%$. Datasets 5–8 show mean true negative rate 99.84%, true positive rate 99.91%, false positive rate 0.17%, and false negative rate 0.10%.	201
5.7	Model Performance Under Varying Attack Intensities and Network Congestion	203
5.8	Comprehensive Comparative Analysis Against Competing Methods. (a) Detection accuracy comparison across all 8 datasets showing proposed model superiority on every dataset (100% win rate). (b) Multi-metric performance comparison across precision, recall, and F1-score demonstrating balanced excellence. (c) Training time comparison demonstrating efficiency advantages over Deep Learning whilst maintaining substantially higher accuracy. (d) Inference time comparison confirming real-time capability with $3.4\times$ speedup vs Deep Learning.	209
5.9	Statistical Validation: Confidence Intervals, Distribution Analysis, and Significance Testing. (a) Accuracy with 95% confidence intervals showing exceptionally tight ranges (± 0.02 – 0.04 percentage points) across all 8 datasets. (b) Box plots displaying accuracy distribution across 5 independent runs per dataset, demonstrating minimal variance and no outliers. (c) Parameter sensitivity analysis showing robust performance across $\pm 40\%$ hyperparameter variation from optimal values. (d) Statistical significance heatmap showing p -values for all pairwise method comparisons across competing approaches.	219

5.10 Bottleneck Analysis: Data Volume, User Scalability, Bandwidth, and Resource Exhaustion. (a) Detection accuracy (left axis) and processing time (right axis) versus data volume, showing optimal performance (99.85%+) up to 2 TB with graceful degradation beyond. (b) System throughput (packets/second) and CPU usage (%) versus concurrent devices, identifying 10,000 device operational limit with 99%+ performance up to 5,000 devices. (c) Detection accuracy (%) and latency (ms) versus traffic load (Gbps), demonstrating optimal throughput (99.85%+) up to 850 Mbps with acceptable performance up to 2 Gbps. (d) Detection accuracy (%) under various resource exhaustion scenarios (CPU, memory, combined), emphasising 80% utilisation threshold for maintaining 99%+ accuracy.	231
5.11 Distinguishing DDoS Attacks from Legitimate Resource Exhaustion. (a) Legitimate traffic overload pattern showing gradual 45-minute increase from baseline 1,000 flows/second to peak 8,500 flows/second (8.5× increase) with smooth exponential transitions and sustained plateau. (b) DDoS attack traffic pattern exhibiting sudden 3-minute spike from baseline 1,000 flows/second to peak 15,000 flows/second (15× increase) with sharp onset, sustained attack rate, and abrupt termination. (c) Traffic feature comparison across five key discriminating characteristics (normalised 0–1 scale): packet rate variance, source IP diversity, connection duration, payload size standard deviation, protocol distribution showing clear separation between legitimate overload (balanced features) and DDoS attack (skewed features). (d) Attack detection confidence over time, showing clear separation between legitimate overload (consistently low probability < 0.12 throughout event) and actual DDoS attack (rapid increase from baseline 0.05 to high confidence > 0.95 within 2 minutes, sustained high confidence throughout attack).	241
Appendix B.1 IoT Architecture	297
Appendix B.2 ETSI NFV Architectural Framework	298
Appendix B.3 Migration of network functionality to hardware	299
Appendix B.4 SDN Development Timeline (2008–2022)	301
Appendix B.5 SDN Architectural Layers	303
Appendix B.6 IoT Cloud Computing Architecture	306
Appendix B.7 NFV Structure	309
Appendix B.8 The CIA Triad: Confidentiality, Integrity, and Availability.	319
Appendix B.9 DDoS Attack Types Classifications	320
Appendix B.10 Machine Learning Algorithms	329
Appendix B.11 IoT Challenges and Threats	330
Appendix B.12 Progression of SYN flood attack	342
Appendix B.13 Progression of Slowloris attack	342

List of Tables

2.1	Common DDoS Attack Types and Characteristics	27
3.1	Requirements Prioritisation Scheme	52
3.2	Requirements Validation Summary	53
3.3	Architecture Complexity and Performance Trade-offs	59
3.4	Architecture Depth Impact on Validation Performance	59
4.1	Host Computer Specifications	73
4.2	Raspberry Pi Configuration	74
4.3	Network Equipment Specifications	75
4.4	GNS3 Virtual Machine Specifications	75
4.5	Windows 10 Virtual IoT Device Specifications	76
4.6	Kali Linux Attacker VM Specifications	76
4.7	Sensor GPIO Pin Mapping	77
4.8	Data Transmission Parameters	77
4.9	Data Logging Server Configuration	78
4.10	Network Monitoring Tools Configuration	78
4.11	Orchestration System Capabilities	78
4.12	Attack Vector Specifications and Capabilities	81
4.13	Traffic Generation Methods and Parameters	84
4.14	Monitoring Tools Configuration and Usage	84
4.15	Data Collection Duration and Environmental Conditions	85
4.16	Baseline Data Collection Statistical Summary	85
4.17	Host Computer CPU and Memory Utilisation Under Benign Conditions	88
4.18	Raspberry Pi CPU and Memory Utilisation Under Benign Conditions	88
4.19	Virtual IoT Device CPU and Memory Utilisation Under Benign Conditions	89
4.20	Network Bandwidth Utilisation Per Device Under Benign Conditions	89
4.21	Protocol Distribution in Normal Traffic	90
4.22	Virtualisation Overhead Comparison	90
4.23	Round-Trip Time Measurements Under Normal Conditions	91
4.24	Throughput Measurements Using iperf3	92
4.25	Established Performance Thresholds	92
4.26	Physical vs. Virtualised IoT Traffic Comparison	93
4.27	Attack Tool Configuration and Capabilities	96
4.28	Attack Orchestration System Features	97
4.29	DDoS Attack Intensity Levels	98
4.30	Summary of Datasets Used in This Research	102
4.31	CIC-DDoS2019 Dataset Characteristics	102

4.32	CIC-DDoS2019 Dataset Split	103
4.33	IoT-23 Dataset Characteristics	104
4.34	IoT-23 Dataset Split	104
4.35	IoT-DH Dataset Characteristics	105
4.36	IoT-DH Dataset Split	105
4.37	SDN-DDoS Dataset Characteristics	106
4.38	SDN-DDoS Dataset Split	106
4.39	CICIoT2023 Dataset Characteristics	107
4.40	CICIoT2023 Dataset Attack Type Distribution	108
4.41	CICIoT2023 Dataset Split	108
4.42	EdgeIIoTset Dataset Characteristics	109
4.43	EdgeIIoTset Dataset Traffic Distribution by Layer	110
4.44	EdgeIIoTset Dataset Split	110
4.45	CICIoMT2024 Dataset Characteristics	111
4.46	CICIoMT2024 Dataset Composition by Medical Device Category	113
4.47	CICIoMT2024 Dataset Attack Type Distribution	113
4.48	CICIoMT2024 Dataset Split	114
4.49	Generated Virtualised IoT Dataset Characteristics	115
4.50	Flow-Level Features in Generated Virtualised IoT Dataset	117
4.51	Generated Virtualised IoT Dataset Composition	118
4.52	Generated Virtualised IoT Dataset Split (Stratified Sampling)	118
4.53	Label Confidence Score Categories	125
4.54	Label Validation Results Across Multiple Sources	125
4.55	Inter-Rater Reliability Assessment	126
4.56	FCNN Parameter Distribution and Computational Complexity	130
4.57	Dataset Split Rationale and Justification	131
4.58	Stratification Criteria for Dataset Splits	131
4.59	Statistical Similarity of Train/Validation/Test Splits	132
4.60	Class Distribution Before Balancing (Combined Datasets)	132
4.61	SMOTE-Tomek Configuration Parameters	133
4.62	Class Distribution After SMOTE-Tomek Balancing	133
4.63	Impact of Data Balancing on Model Performance	133
4.64	Training Infrastructure Specifications	134
4.65	Training Configuration Parameters	134
4.66	Learning Rate Schedule Configuration	135
4.67	Training Monitoring Metrics	135
4.68	Training Duration Statistics	135
4.69	Hyperparameter Search Space	136
4.70	Hyperparameter Tuning Methodology	136
4.71	Optimal Hyperparameters Selected Through Tuning	137
4.72	Hyperparameter Importance Ranking	137
4.73	Model Convergence Criteria	137
4.74	Final Model Performance on Validation Set	140
4.75	Overfitting Assessment Results	140
4.76	Competing Method Selection Criteria	141
4.77	Random Forest Configuration	141
4.78	SVM Configuration	142

4.79	CNN-LSTM Hybrid Architecture Configuration	143
4.80	Common Implementation Framework	144
4.81	Fair Comparison Framework	144
4.82	Implementation Validation Results	144
4.83	Experimental Scenario Categories	145
4.84	Experimental Scenarios to Requirements Traceability Matrix	168
4.85	Requirements Coverage Summary Across Experimental Scenarios	169
5.1	Overall Performance Metrics Across All Evaluated Datasets (Mean $\pm$ SD from 5 Independent Runs)	177
5.2	Comparative Performance: Proposed Model vs Competing Methods (Mean Across All 8 Datasets)	209
5.3	Performance Comparison with Recent State-of-the-Art Methods for IoT DDoS Detection	212
Appendix A.1	Smart City Resource Domain Classifications	290
Appendix A.2	Smart City Sub-domains and Characteristics	291
Appendix B.1	Current Challenges and Issues In The IoT Cloud.	307
Appendix B.2	Types of DDoS Attacks in IoT and Examples	323
Appendix B.3	Feature-Based DDoS Attack Detection Methods in IoT Envi- ronments	325
Appendix B.4	Behaviour-Based DDoS Detection Methods in IoT Environments	325
Appendix B.5	common vulnerabilities in IoT virtual environments	333
Appendix B.6	A Comparative Analysis of Detection Techniques Utilised in Various Defense Mechanisms for DDoS Attacks	336
Appendix B.7	Comparative Analysis of Recent DDoS Defense Mechanisms (Part 1)	338
Appendix B.8	Comparative Analysis of Recent DDoS Defense Mechanisms (Part 2)	339
Appendix B.9	Comparative Analysis of Recent DDoS Defense Mechanisms (Part 3)	340
Appendix B.10	Comparison of Machine Learning Algorithms Used in DDoS Attack Detection (2019–2025)	348
Appendix B.11	Comparison of FCNN and CNN in Security Detection	351
Appendix B.12	Summary of Studies on FCNN in Anomaly Detection	353
Appendix C.1	Non-Functional Requirements Governing System Design	356
Appendix C.2	Functional Requirements Defining System Capabilities	359

Declaration of Authorship

I declare that this thesis and the work presented in it is my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Parts of this work have been published as:

Signed:.....

Date:.....

Acknowledgements

In the name of Allah, the Most Gracious, the Most Merciful

Above all, I offer my sincerest gratitude to **Allah** the Most Gracious, the Most Merciful for His boundless blessings, for granting me the strength and wisdom to complete this journey, and for surrounding me with people whose support has been a constant source of light. *Alhamdulillah.*

I express my deepest gratitude to my primary supervisor, **Dr. Nawfal Al Hashimy**, for his invaluable guidance, unwavering support, and insightful feedback throughout my PhD journey. I am equally grateful to my second supervisor, **Dr. Ahmad Atamli**, whose thoughtful advice and technical expertise have significantly refined my work.

I would also like to thank the staff at the **School of Electronics and Computer Science, University of Southampton**, especially the **Cyber Security** and **CPS research groups**, for their continuous support, and my colleagues and friends for their motivation and collaboration throughout this journey.

To my beloved **mother** and **siblings** your unconditional love, sacrifices, and unwavering faith in me have been my greatest source of strength. You held me up through every challenge, and I am forever grateful.

To my dear **father**, who left this world just one month before my viva you did not get to see this moment, but every page of this thesis carries your love, your prayers, and your belief in me. I hope, wherever you are, you are proud. This is for you. *Al-Fatiha.*

*To all those who walked this journey with me, directly or indirectly
thank you, from the depths of my heart.*

*To my dearest Mum, **Fatima** your love, strength, and sacrifices are woven into every page of this work.*

*And to my beloved Dad, **Ezzuldeen** you left before you could see this day, but you never left my heart. This is yours.*

May Allah grant you the highest Jannah, Baba.

Chapter 1

Introduction

The Internet of Things (IoT) has emerged as a cornerstone of modern digital infrastructure, underpinning critical domains, including smart cities, healthcare, transportation, and industrial systems. By enabling autonomous interaction between interconnected devices and networks, IoT promises enhanced operational efficiency and improved quality of life. However, IoT devices are characteristically resource-constrained, highly heterogeneous, and deployed on a massive scale, creating unprecedented challenges for computation, communication, and security [Almurisi and Tadisetty \(2022\)](#). To address this complexity, IoT systems increasingly leverage virtualisation technologies such as Software-Defined Networking (SDN) and Network Function Virtualisation (NFV). Although these technologies offer substantial improvements in scalability and resource management, they simultaneously introduce significant security vulnerabilities [Alam et al. \(2019\)](#). Distributed Denial of Service (DDoS) attacks represent a particularly critical threat to virtualised IoT (vIoT) environments, where the compromise of shared resources can precipitate widespread disruption to essential services. This thesis investigates the security challenges inherent in vIoT systems, with particular emphasis on the detection and mitigation of DDoS attacks. Identify the limitations of existing security solutions and propose a novel AI-based framework designed to improve resilience, scalability, and adaptability in the security of virtualised IoT networks.

1.1 Research Motivation

The IoT has evolved from a consumer convenience technology into a foundational element of modern critical infrastructure, encompassing healthcare systems, industrial automation, transportation networks, and smart city services. The forecasts predict more than 25 billion connected devices by 2030 [Xinyu et al. \(2025\)](#), driving unprecedented growth in data volume, device heterogeneity, and network interconnectivity. Although

such expansion offers immense opportunities for operational efficiency and automation, it simultaneously amplifies the attack surface presented to malicious actors.

Resource-constrained IoT devices, diverse communication protocols, and complex data flows make these environments highly vulnerable to cyber attacks [Donca et al. \(2024\)](#). To manage this complexity, virtualisation technologies such as SDN and NFV have been widely adopted to enable scalable resource management and network programmability. However, when deployed in distributed IoT infrastructures, these technologies also introduce systemic vulnerabilities that adversaries can exploit [Alam et al. \(2019\)](#).

Among the spectrum of cyber threats, DDoS attacks represent one of the most disruptive, capable of exhausting bandwidth, computational capacity, and energy resources. Within vIoT environments, the convergence of physical and virtual domains exacerbates this challenge: a single compromised virtual node can propagate failures across multiple layers of abstraction, potentially disrupting essential services and imposing substantial financial and societal costs.

1.1.1 The Imperative of Smart City Security

Smart city initiatives, driven by global urbanisation trends, increasingly depend on IoT infrastructures that demand high availability, reliability, and real-time responsiveness. According to projections, more than 60% of the world's population will reside in urban areas by 2050 [Yaqoob et al. \(2017\)](#); [United Nations \(2018\)](#), substantially increasing the demand for secure and resilient IoT-based services in domains such as energy distribution, traffic management, and emergency response systems.

The heterogeneity of deployed devices, combined with their dependence on cloud-based virtualisation, renders these systems particularly vulnerable to coordinated attacks. A single DDoS incident can cascade across interconnected services, disrupting critical public infrastructure, and erode citizen trust in smart city deployments.

1.1.2 The Research Gap

Despite extensive research in IoT security, significant gaps remain in understanding and securing virtualised IoT environments. Existing DDoS detection mechanisms are predominantly designed for traditional network architectures and prove inadequate when applied to vIoT contexts characterised by:

- **Blurred boundaries** between physical sensors and virtualised network functions, complicating traffic attribution and anomaly detection;
- **Dynamic resource sharing** across multi-tenant cloud infrastructures, where legitimate traffic patterns vary significantly;

- **Limited adaptability** of detection systems to evolving attack vectors and novel exploitation techniques.

Artificial Intelligence (AI) and Machine Learning (ML) approaches offer promising capabilities to identify abnormal traffic patterns in real time. However, these methods often lack systematic validation within virtualised contexts and suffer from reproducibility challenges due to the absence of representative datasets and standardised evaluation methodologies.

1.1.3 Principal Research Motivations

The following points encapsulate the key motivations driving this research:

- **Detection capability gap:** Address the absence of robust mechanisms capable of distinguishing legitimate from malicious traffic in virtualised IoT environments, particularly under conditions of high device heterogeneity, resource sharing, and dynamic network topology.
- **Collaborative security requirements:** Develop security frameworks that ensure reliable collaboration amongst geographically distributed sensor networks in cloud-based infrastructures, protecting the integrity and availability of data flows between perception layers and virtualised services.
- **Adaptive intelligence:** Provide intelligent and scalable detection techniques, using artificial intelligence, that can dynamically adapt to evolving attack patterns and, thus, enhance the resilience of IoT-based smart city and industrial applications.
- **Reproducible evaluation:** Develop representative testbed environments and establish rigorous experimental protocols to enable reproducible evaluation and independent validation of detection mechanisms in vIoT contexts.

This research addresses these motivations by proposing a novel AI-driven framework for DDoS detection in virtualised IoT environments, grounded in quantitative baseline analysis and validated through comprehensive experimental evaluation.

1.2 Research Justification

The transition towards multi-cloud environments has undermined traditional approaches to infrastructure monitoring and security management. At the same time, virtualisation

of the IoT through technologies such as NFV, SDN, and cloud computing has increased existing security challenges and introduced new vulnerabilities.

Currently, there is no recognised IoT cloud-based mechanism capable of detecting malicious nodes or abnormal traffic within the virtualisation layer. This deficiency is particularly concerning in the context of DDoS attacks, where the convergence of physical sensors, virtual network functions, and shared cloud resources creates complex attack surfaces that evade conventional intrusion detection systems.

This gap gives rise to the following critical issues.

- **Failure to detect malicious traffic**, which may result in large-scale DDoS attacks that overwhelm both physical and virtual resources, disrupting essential services across virtualised IoT infrastructures.
- **Failure to identify compromised or failed nodes**, which can propagate instability across interconnected services and lead to cascading network failures within multi-tenant virtualised environments.
- **Lack of visibility between layers**, preventing a comprehensive analysis of how attacks propagate between physical IoT devices, virtualised network functions, and cloud infrastructure components.

Furthermore, the interaction between physical and virtualised layers introduces additional vulnerabilities, including attacks originating from untrusted applications or compromised devices that exploit weak virtualisation boundaries and shared resource pools.

Addressing these challenges requires an AI-driven framework capable of detecting DDoS attacks in vIoT systems with enhanced cross-layer visibility and adaptive detection capabilities.

1.3 Research Questions and Objectives

This research addresses identified security challenges through three main research questions, each corresponding to a distinct contribution.

RQ1: How do DDoS attacks manifest and propagate within virtualised IoT environments, and what distinguishing characteristics enable differentiation between normal and malicious behaviour in such contexts?

This question investigates the specific attack vectors, propagation patterns, and traffic characteristics unique to vIoT systems through controlled experimental analysis.

RQ2: What AI-driven detection framework architecture can effectively identify DDoS attacks in virtualised IoT environments whilst addressing the challenges of resource heterogeneity, dynamic topology, and multi-tenant infrastructure?

This question focusses on designing a novel detection framework that leverages AI to provide cross-layer visibility and adaptive detection capabilities tailored to the characteristics of the vIoT.

RQ3: How does the proposed AI-driven framework perform compared to contemporary detection approaches when evaluated across multiple datasets and varying attack scenarios?

This question establishes systematic comparative evaluation methodologies, validating the proposed framework against state-of-the-art approaches using recent datasets and diverse experimental conditions.

1.3.1 Research Objectives

Objective 1 (RQ1): Vulnerability Analysis of vIoT and Experimental Testbed

1. Conduct a systematic review of the security vulnerabilities in vIoT;
2. Design and implement a reproducible vIoT testbed;
3. Establish comprehensive baseline characterisation of normal network behaviour;
4. Identify bottlenecks, resource constraints, and failure points through controlled stress testing.

Objective 2 (RQ2): AI-Based Detection Framework

1. Design a Fully Connected Neural Network (FCNN) framework for DDoS detection;
2. Develop feature-selection and optimisation strategies;
3. Integrate the framework within the vIoT infrastructure;
4. Validate the framework across multiple datasets.

Objective 3 (RQ3): Comparative Evaluation and Reproducibility

1. Implement at least three contemporary competing AI methods;
2. Conduct a comparative evaluation using identical datasets and experimental conditions;

3. Perform statistical validation and sensitivity analysis;
4. Provide detailed reproducibility documentation and protocols.

1.4 Research Contributions

This research makes three principal contributions to vIoT security, addressing the detection of DDoS attacks through AI-driven frameworks, reproducible experimental methodologies, and comprehensive cross-model validation.

1.4.1 Contribution 1: Quantitatively Validated AI Framework for vIoT DDoS Detection

This contribution merges the systematic identification of security challenges with the development and validation of an AI-based detection framework.

- Systematic identification of vIoT vulnerabilities and attack propagation patterns through comprehensive experimental analysis of traffic patterns, resource utilisation, and performance thresholds;
- Design, implementation, and validation of a FCNN-based DDoS detection framework optimised for vIoT environments;
- Identification of system bottlenecks, resource constraints, and critical failure points through controlled stress testing;
- Comparative evaluation across multiple benchmark datasets, demonstrating the effectiveness of the framework under diverse operational conditions.

1.4.2 Contribution 2: Novel Testbed Architecture with Reproducibility Protocols

This contribution addresses the critical gap in reproducible vIoT security research:

- Design and implementation of a representative vIoT testbed integrating SDN/NFV architectures, multi-tenant cloud infrastructure, and heterogeneous IoT devices;
- Complete experimental documentation enabling independent replication, including detailed testbed configuration, parameter specifications, and version-controlled implementations;

- Comprehensive documentation of experimental procedures and environmental configurations, establishing reproducible protocols for vIoT security evaluation under varying operational scenarios;
- Open research practices with detailed protocols for the generation of data sets, experimental conditions, and validation procedures.

1.4.3 Contribution 3: Cross-Model Validation and Contemporary Dataset Integration

This contribution validates framework generalisability.

- Implementation and comparative evaluation of alternative AI architectures (Convolution Neural Network (CNN), Long Short-Term Memory (LSTM) and Random Forests (RD)) using identical datasets and experimental conditions;
- Integration and validation using contemporary public datasets: demonstrating applicability to modern threat landscapes;
- Reverse validation through application of external AI methods to the generated vIoT dataset, establishing cross-model generalisability;
- Statistical validation with confidence intervals, significance testing, and sensitivity analysis across all evaluated approaches.

1.5 Research Scope and Limitations

This research focusses on DDoS attack *detection and characterisation* within vIoT environments, rather than mitigation strategies or prevention mechanisms. Experiments are conducted within a controlled vIoT testbed environment based on SDN and NFV architectures. Although the framework is designed for generalisability, validation is performed using publicly available datasets and controlled experimental conditions to ensure reproducibility.

The study explicitly excludes:

- Real-time deployment in production smart city environments;
- Hardware-specific optimisations for resource-constrained IoT devices;
- Legal or policy frameworks for the response, attribution, and prosecution of DDoS;
- DDoS mitigation techniques or active defence mechanisms;
- Non-DDoS attack vectors (e.g., data exfiltration, privilege escalation).

1.5.1 Application Context: Smart City IoT

Smart city infrastructures-including energy distribution grids, intelligent transportation systems, and healthcare monitoring networks serve as the primary application context for this research. These domains exemplify the convergence of resource heterogeneity, real-time responsiveness requirements, and multi-tenant virtualised architectures that characterise modern vIoT deployments.

The smart city context provides:

- **Realistic complexity:** Heterogeneous device types, diverse communication protocols, and varying quality-of-service requirements representative of real-world deployments;
- **Critical infrastructure characteristics:** High availability demands, cascading failure risks, and substantial societal impact that justify rigorous security research;
- **Virtualisation dependence:** Extensive reliance on cloud-based orchestration and SDN/NFV technologies that represent the future of IoT infrastructure management.

1.5.2 Research Significance

This research advances understanding of DDoS attack detection in vIoT environments through systematic experimental analysis and AI-driven frameworks. It demonstrates the efficacy of adaptive detection mechanisms through rigorous comparative evaluation and reproducible methodologies, addressing critical gaps in vIoT security research.

The validated and reproducible framework developed in this research enables practitioners to systematically assess DDoS detection performance across diverse datasets and attack scenarios. The comprehensive reproducibility protocols and testbed architecture provide a foundation for independent validation and extension by other researchers, addressing the field's persistent reproducibility challenges.

By contributing to the security of critical smart city services, this research helps safeguard the essential public infrastructure that millions of citizens depend on for energy distribution, transportation, emergency response, and healthcare services. Enhanced DDoS detection capabilities reduce the risk of large-scale service disruptions that could compromise public safety and erode trust in smart city initiatives.

1.6 Thesis Structure

The remainder of this thesis is organised as follows:

Chapter 2: Background and Related Work presents a focused, critical review of vIoT security literature, with particular emphasis on DDoS detection mechanisms. This chapter systematically examines the core concepts of vIoT, security challenges, state-of-the-art detection approaches, and AI techniques relevant to this research. Critical analysis identifies key research gaps and places the present work within the broader field of vIoT security.

Chapter 3: Research Methodology and Experimental Framework details the research design, methodology, and experimental framework used in this study. This chapter describes the vIoT testbed architecture, the baseline network behaviour analysis procedures, the data collection strategies, and the AI model development approach. It establishes comprehensive reproducibility protocols and addresses ethical considerations, providing the methodological foundation for rigorous experimental validation.

Chapter 4: Experimental Design and Implementation documents the practical implementation of the research methodology. This chapter covers vIoT testbed deployment, comprehensive baseline network characterisation under benign conditions, attack scenario design, and dataset generation. Details the training and optimisation of the proposed FCNN framework alongside the implementation of competing methods, ensuring fair comparative evaluation in diverse experimental scenarios.

Chapter 5: Results, Analysis, and Discussion presents comprehensive experimental results and critical analysis. Beginning with baseline network performance findings, the chapter proceeds through a detailed evaluation of the performance of AI models in multiple datasets, comparative analysis against state-of-the-art approaches, and rigorous statistical validation. Critical bottleneck analysis examines the limits of the system and distinguishes attack-induced resource exhaustion from legitimate operational constraints. The discussion addresses each research question systematically, whilst acknowledging methodological limitations and threats to validity.

Chapter 6: Conclusions and Future Work synthesises the research contributions and main findings. This chapter provides a refined articulation of the three core contributions, reflects on limitations and lessons learned, and outlines promising directions for future work in virtualised IoT security, including enhanced detection mechanisms, alternative AI architectures, real-world deployment considerations, and cross-domain applications.

Chapter 2

Background and Literature Review

The convergence of IoT, virtualisation technologies, and advanced networking paradigms has fundamentally transformed the way connected systems are designed, deployed, and managed. This transformation brings unprecedented opportunities for scalability, flexibility, and resource optimisation, yet simultaneously introduces complex security challenges that require a comprehensive investigation.

This chapter provides a structured exploration of the foundational technologies that underlie vIoT environments and the security concerns they pose. The discussion begins with an examination of the IoT architecture and its evolution, establishing the essential framework for understanding modern interconnected ecosystems. This is followed by an investigation of key virtualisation technologies NFV and SDN that enable dynamic network management and flexible service deployment. The integration of these technologies with cloud computing for IoT applications is then explored, highlighting how virtualisation enhances scalability and resource efficiency.

Security challenges form a central theme throughout this chapter. Particular attention is paid to vulnerabilities inherent in virtualised architectures, with DDoS attacks examined as a critical threat vector. The chapter reviews existing detection and mitigation approaches, including traditional methods and machine learning-based solutions, critically analysing their effectiveness in vIoT contexts. A systematic review methodology is presented to synthesise current research, identifying both achievements and limitations in the field.

The chapter concludes by identifying critical research gaps in the literature, particularly regarding integrated security frameworks for vIoT environments and the need for lightweight adaptive detection mechanisms suitable for resource-constrained devices.

These gaps establish the foundation for the research contributions presented in subsequent chapters, positioning this work within the broader landscape of IoT security research.

2.1 IoT and Virtualisation Fundamentals

This section establishes the technical foundations necessary for understanding vIoT environments. It begins with an examination of the IoT architecture before exploring the virtualisation technologies such as NFV and SDN that enable flexible and scalable IoT deployments. The integration of these technologies with cloud computing is then discussed, followed by an overview of the security challenges that emerge from this convergence.

2.1.1 IoT Architecture

The IoT represents a paradigm in which physical objects, embedded with sensors, actuators, and communication capabilities, form interconnected networks that collect, exchange, and process data. In recent years, industrial sectors have increasingly adopted IoT applications, many incorporating cyberphysical systems (CPS) and machine-to-machine (M2M) communications [Mu and Antwi-Afari \(2024\)](#). These applications frequently handle sensitive information, particularly in domains such as e-healthcare and e-government, necessitating robust security measures [Li et al. \(2024\)](#).

Modern IoT platforms leverage the conventional Internet infrastructure through wireless sensor networks and mobile networks to provide the flexibility and scalability required for large-scale deployments [Al-Qaseemi et al. \(2017\)](#). However, achieving security in IoT environments presents greater challenges than in traditional Internet systems, mainly due to the resource-constrained nature of IoT devices [Rahul et al. \(2017\)](#). Although traditional Internet security requires essential properties such as confidentiality, integrity, and authentication, meeting these requirements on IoT platforms is considerably more complex, as IoT devices typically possess limited computational power, memory, and energy resources compared to conventional computing devices [Rachakonda et al. \(2024\)](#).

Despite these constraints, IoT platforms are increasingly deployed in critical infrastructure and sensitive areas of national importance, including intelligent transportation systems and healthcare services. Consequently, information and network security in IoT platforms requires not only the fundamental security properties of traditional systems but also enhanced availability and dependability to meet the demands of mission-critical applications [Al-Ansi et al. \(2024\)](#).

The architectural organisation of IoT systems has evolved to address the complexity of connecting diverse devices, networks, and applications. Contemporary IoT architecture is generally conceptualised as a layered model, with varying levels of granularity depending on the specific implementation requirements. The three-layer foundational model consists of perception, network, and application layers, while more detailed architectures expand this to five layers: perception, access gateway, network, middleware (or support) and application layers [Mazon-Olivo and Pan \(2022\)](#). Each layer serves different functions and presents unique security and management challenges.

- **Perception Layer:** Also referred to as the physical or sensing layer, this foundational layer encompasses the physical devices, sensors, and actuators that interact with the physical environment. These components collect raw data from the surrounding environment through embedded sensors, which may include temperature sensors, motion detectors, cameras, and various other monitoring devices. The perception layer serves as the interface between the physical world and the digital IoT infrastructure, capturing real-time information that forms the basis for higher-level processing and decision-making [Mazon-Olivo and Pan \(2022\)](#).
- **Access Gateway Layer:** Functioning as a management intermediary, this layer coordinates communications between the perception layer and higher network layers. It is primarily responsible for the translation of the protocol, the aggregation of data, and the management of the flow of information between heterogeneous devices and systems. The access gateway layer facilitates the exchange of data between various IoT devices and ensures compatibility between different communication protocols and standards [Antão et al. \(2018\)](#).
- **Network Layer:** Often called the connection or transport layer, this component provides the communication infrastructure necessary for data transmission across the IoT ecosystem. It encompasses both wired and wireless networking technologies, including Wi-Fi, cellular networks, ZigBee, LoRaWAN, and other communication protocols. The network layer establishes connectivity between the physical layer and the application layer, enabling bidirectional communication; it transmits collected data from sensors to processing systems and delivers control commands from applications back to physical devices [Mazon-Olivo and Pan \(2022\)](#).
- **Middleware Layer:** Also known as the support or service management layer, this component acts as an abstraction layer that bridges low-level device communications with high-level application requirements. Provides essential services such as data processing, storage management, device management, and service composition. The middleware layer improves flexibility and interoperability by decoupling hardware-specific implementations from application logic, enabling IoT systems to integrate diverse devices and platforms seamlessly [Al-Qaseemi et al. \(2017\)](#).

- **Application Layer:** Representing the top layer in the architecture, this component delivers IoT functionality to end-users and business processes. It includes user interfaces, analytics platforms, and business applications that consume IoT data and generate actionable insights. The application layer transforms raw sensor data into meaningful information for decision-making, typically consisting of processing sublayers for data analytics and business sublayers for application-specific logic (Mazon-Olivo and Pan (2022)).

Each architectural layer presents different security vulnerabilities and management challenges that must be addressed to ensure the overall security and reliability of IoT systems. Figure 2.1 illustrates the five-layer IoT architecture, showing the relationships and data flows between layers.

Real-world implementations of IoT systems highlight practical security considerations that must be addressed in deployment. Mekala et al. (2023); Gupta (2021) demonstrated an industrial IoT sensor system for high-temperature measurement, identifying critical security practices that include: (i) physical protection of devices from unauthorised access, (ii) regular firmware and patch updates to prevent exploitation of known vulnerabilities, (iii) strong authentication mechanisms for device accessibility, and (iv) secure communication protocols for data transmission. These practical considerations become even more critical in virtualised IoT environments where multiple logical instances may share physical resources, creating additional attack surfaces and isolation challenges.




FIGURE 2.1: IoT Architecture

The layered architecture provides a conceptual framework for understanding IoT systems, yet real-world implementations often require integration with virtualisation technologies to achieve the scalability and flexibility demanded by modern applications. The following subsections explore these virtualisation technologies and their role in contemporary IoT deployments.

2.1.2 Network Function Virtualisation

Network Function Virtualisation represents a transformative approach to network service delivery, fundamentally changing how network functions are implemented and managed. Traditional networking relies on dedicated hardware appliances, such as routers, firewalls, load balancers, and intrusion detection systems; each requiring specific hardware configurations and physical deployment. NFV decouples these network functions from proprietary hardware, implementing them as software applications that run on a commodity computing infrastructure [Zarca et al. \(2020\)](#).

The European Telecommunications Standards Institute (ETSI) has been instrumental in standardising the NFV architecture, defining a framework that enables network operators to virtualise entire classes of network node functions [ETSI \(2013\)](#). This virtualisation paradigm offers several significant advantages for modern network deployments, including reduced capital expenditure through the use of commercial standard hardware, improved operational flexibility through the deployment of software-based network functions, enhanced scalability through dynamic resource allocation, and accelerated service innovation through the rapid deployment and modification of network services [Alam et al. \(2019\)](#).

The NFV architectural framework, as standardised by ETSI, comprises three primary components that work together to deliver virtualised network services. The NFV Infrastructure (NFVI) provides the physical and virtual resources necessary to host Virtual Network Functions (VNFs), encompassing compute, storage, and networking resources along with the virtualisation layer that abstracts these physical resources [Samaniego et al. \(2018\)](#). VNFs themselves are software implementations of network functions that execute on the NFVI, ranging from virtualised routers and firewalls to more complex functions like deep packet inspection and traffic optimisation. The NFV Management and Orchestration (MANO) framework coordinates the management of the VNFs lifecycle and the allocation of infrastructure resources, ensuring efficient use and service delivery [Demirci and Sagioglu \(2019\)](#).

In the context of IoT, NFV enables flexible deployment of network services that can adapt to the dynamic requirements of IoT applications. For example, security functions such as firewalls and intrusion detection systems can be instantiated as VNFs, allowing them to scale elastically based on traffic patterns and security threats. This flexibility is

particularly valuable in IoT environments characterised by highly variable traffic loads and diverse security requirements across different application domains [Liu et al. \(2025\)](#). Figure 2.2 illustrates the NFV architectural framework and its key components.




FIGURE 2.2: Network Function Virtualisation Architecture (This figure has been re-sketched from [ETSI \(2013\)](#))

However, virtualisation of network functions introduces new security considerations that must be carefully addressed. The NFV shared infrastructure model creates potential attack surfaces that do not exist in traditional hardware-based deployments, including hypervisor vulnerabilities, resource isolation failures, and exposures to the management interface [Jiang et al. \(2021\)](#). These security challenges become particularly critical in IoT contexts where various devices with varying trust levels must coexist on shared infrastructure [Hao et al. \(2018\)](#) [Zahran et al. \(2024\)](#).

2.1.3 Software-Defined Networking

SDN complements NFV by introducing a new paradigm for network control and management. Traditional networks distribute control logic across individual network devices, and each switch or router makes independent forwarding decisions based on locally configured rules and distributed protocols [Kaur et al. \(2025\)](#). SDN fundamentally restructures this model by separating the control plane, which makes decisions about where

traffic should be sent, from the data plane; which forwards traffic according to those decisions [Kaur et al. \(2025\)](#).

This architectural separation enables centralised network control through a logically centralised controller that maintains a global view of the network topology and state. The SDN controller communicates with network devices through standardised protocols, most notably OpenFlow, which defines how the controller programmes forwarding rules into network switches [Gupta et al. \(2024\)](#). This centralised control model provides several advantages for network management, including simplified policy enforcement through centralised configuration, improved visibility through comprehensive network monitoring, improved agility through programmable network behaviour, and dynamic traffic engineering based on real-time network conditions [Kaur et al. \(2025\)](#).

The SDN architecture consists of three distinct layers that work together to provide programmable network control. The application layer comprises network applications and business logic that define network behaviour and policy, communicating with the control layer through northbound APIs [Gupta et al. \(2024\)](#). The control layer implements the SDN controller, which maintains the network topology, processes application requests, and translates high-level policies into low-level forwarding rules. The infrastructure layer consists of physical and virtual network devices (switches, routers) that forward traffic according to the rules installed by the controller, communicating through southbound APIs such as OpenFlow [Sangodoyin et al. \(2018b\)](#).

For IoT deployments, SDN offers particular advantages in managing the complex and heterogeneous networks that characterise modern IoT ecosystems. Centralised control enables dynamic adaptation of network behaviour to accommodate varying IoT traffic patterns, quality of service requirements, and security policies. For example, SDN can implement fine-grained access control for IoT devices, isolate different classes of IoT traffic, and dynamically reroute traffic to avoid congested or compromised network paths [Gupta et al. \(2024\)](#). Figure 2.3 illustrates the SDN architectural model showing the separation of control and data planes.

Despite these advantages, SDN introduces specific security vulnerabilities that must be considered. The centralised controller represents a single point of failure and an attractive target for attackers [Farahi \(2025\)](#). The compromise of the controller could enable an adversary to manipulate the entire network's behaviour. Additionally, communication channels between the controller and network devices, typically using the OpenFlow protocol, must be adequately secured to prevent unauthorised access or traffic manipulation [Najar and Manohar Naik \(2024\)](#).




FIGURE 2.3: Software-Defined Networking Architecture (This figure has been re-sketched from Alam et al. (2019))

2.1.4 IoT Cloud Virtualisation

The integration of cloud computing with IoT systems, enhanced by virtualisation technologies, creates powerful platforms for the delivery of IoT services. Cloud computing provides the scalable computational and storage resources necessary to process and analyse the massive data volumes generated by IoT devices, while virtualisation technologies enable efficient use of resources and flexible service deployment Alam et al. (2019).

Cloud IoT platforms typically combine Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS) models to deliver comprehensive IoT solutions Almurisi and Tadisetty (2022). These platforms leverage both NFV and SDN to provide elastic, on-demand network and computing resources that can adapt to the dynamic requirements of IoT applications. Virtualisation of both compute resources and

network functions enables IoT cloud platforms to efficiently support diverse applications with varying resource requirements and quality of service needs [Aazam et al. \(2022\)](#).

The Cloud Computing Adoption Framework (CCAF) proposed by [Chang et al. \(2016\)](#) provides a comprehensive security approach for cloud-based systems that is highly relevant to virtualised IoT deployments. CCAF implements a multi-layered security architecture consisting of three primary defence layers: access control and firewall protection, intrusion detection and prevention systems (IDS/IPS), and encryption mechanisms. This layered approach aligns well with the security requirements of virtualised IoT environments, where multiple attack vectors must be addressed in different architectural layers [Almutairi et al. \(2024\)](#). Through extensive penetration tests involving more than 10,000 viruses and trojans, the CCAF framework achieved detection rates of 99.95% and blocking rates exceeding 85% for continuous attacks, demonstrating the effectiveness of multi-layered security in virtualised infrastructures [Sudheera et al. \(2021\)](#).

Several architectural patterns have emerged for the implementation of virtualised IoT cloud systems. Edge computing architectures position computational resources closer to IoT devices, reducing latency and bandwidth consumption by processing data near its source rather than transmitting all data to centralised cloud datacenters. Fog computing extends this concept by creating a distributed computing hierarchy that spans from edge devices through intermediate fog nodes to cloud resources, enabling flexible deployment of processing tasks based on latency, bandwidth, and computational requirements [Muniswamaiah et al. \(2021\)](#).

Integration of NFV and SDN into cloud environments of IoT enables dynamic service composition and network slicing the ability to create multiple logical networks with different characteristics in a shared physical infrastructure [Almurisi and Tadisetty \(2022\)](#). Network slicing is particularly valuable in IoT contexts where different applications may require vastly different network properties: an industrial control application may require ultra-low latency and high reliability, while a smart city sensor network may prioritise wide coverage and energy efficiency over latency [Alam et al. \(2019\)](#).

Virtualised wireless sensor networks (WSNs) represent a specific instantiation of IoT cloud virtualisation, applying virtualisation concepts to traditional WSN deployments. By abstracting physical sensor networks into virtual sensor resources, multiple applications can share physical sensor infrastructure while maintaining logical isolation [Alam et al. \(2019\)](#). This approach improves resource utilisation and enables for more flexible sensor network management. Figure 2.4 illustrates a virtualised cloud-based WSN architecture showing the integration of physical sensors, the virtualisation layer, and cloud services.

The convergence of IoT, cloud computing, and virtualisation technologies creates new possibilities for scalable and flexible IoT deployments. However, this convergence also increases the security challenges, as vulnerabilities in any component can potentially




FIGURE 2.4: Cloud-based Virtualised Wireless Sensor Network Architecture (This figure has been re-sketched from Alam et al. (2019))

compromise the entire system Alam et al. (2019). The shared infrastructure model, combined with the distributed nature of IoT systems and the complexity of virtualised networking, creates an expanded attack surface that requires comprehensive security measures Zarca et al. (2020).

2.1.5 Security Challenges Overview

The integration of IoT with virtualisation technologies, while offering substantial benefits in terms of scalability and flexibility, introduces a complex landscape of security challenges. These challenges arise from multiple sources: the inherent constraints of IoT devices, the vulnerabilities introduced by virtualisation, the attack surface exposed by SDN architectures, and the complexity of managing security across heterogeneous distributed systems Sasi et al. (2024).

IoT devices typically operate under severe resource constraints, limiting the cryptographic operations and security protocols they can support. Many legacy IoT devices were designed without security as a primary consideration, lacking fundamental security features such as secure boot, encrypted storage, or the ability to receive security updates. The sheer scale of IoT deployments with potentially thousands or millions of devices makes traditional approaches to security management, such as the configuration and monitoring of individual devices, impractical Dias et al. (2025).

Virtualisation introduces additional attack vectors beyond those present in traditional physical deployments. The hypervisor or virtual machine monitor, which manages virtual resources and enforces isolation between virtual machines, becomes a critical security component [Alam et al. \(2019\)](#). Vulnerabilities in the hypervisor can enable attackers to escape virtual machines and access other VMs or the underlying host system. Similarly, inadequate resource isolation can enable side channel attacks in which one virtual machine gains information about another by observing shared resource usage patterns [Ahmadvand et al. \(2023\)](#).

SDN architectures concentrate network control on the controller, creating both opportunities and risks for security. Although centralised control enables consistent security policy enforcement and comprehensive threat monitoring, it also creates a high-value target for attackers [Kaur et al. \(2025\)](#). Successful compromise of an SDN controller can grant an adversary unprecedented control over network behaviour, allowing sophisticated attacks that would be difficult or impossible in traditional distributed control architectures [Tayyaba et al. \(2020\)](#).

The combination of these factors creates a threat landscape in which attackers can leverage vulnerabilities at multiple layers from physical devices through virtualisation infrastructure to network control systems to launch sophisticated attacks. DDoS attacks represent a particularly worrying threat in this context, as they can exploit the scale of IoT deployments and the interconnected nature of virtualised systems to generate overwhelming attack traffic [Sinha \(2025\)](#).

Understanding these security challenges requires examining specific vulnerability classes and attack types in detail. The following section delves into security concerns in virtualised IoT environments, with a particular focus on the vulnerabilities inherent in SDN and NFV architectures and the threat posed by DDoS attacks.

2.2 Security in Virtualised IoT

This section examines the security landscape of virtualised IoT environments in detail, beginning with fundamental security concepts before exploring specific vulnerabilities in SDN and NFV architectures. Particular attention is paid to DDoS attacks as a critical threat to virtualised IoT systems, including an examination of attack types, their characteristics, and detection approaches.

2.2.1 Security Fundamentals

Information security is based on fundamental principles that define the essential properties that a secure system must maintain. These principles, commonly known as the CIA

triad, provide the foundation for understanding and implementing security measures across all types of systems, including IoT and virtualised environments.

Confidentiality ensures that the information is accessible only to authorised entities. In IoT contexts, this principle protects sensitive data collected by sensors such as health information, location data, or industrial process parameters from unauthorised disclosure. Confidentiality is typically maintained through encryption, access control mechanisms, and secure communication protocols [Mekala et al. \(2023\)](#).

Integrity guarantees that the information remains accurate and complete and protected from unauthorised modification or deletion. For IoT systems, maintaining data integrity is crucial to ensuring that sensor readings accurately reflect environmental conditions and that control commands reach actuators without tampering. Integrity mechanisms include cryptographic hashing, digital signatures, and secure update protocols [Zahran et al. \(2024\)](#).

Availability ensures that authorised users can access information and resources when needed. In IoT systems, particularly those that support critical infrastructure or safety-critical applications, availability is paramount. DDoS attacks directly target availability, attempting to prevent legitimate users from accessing services by overwhelming systems with malicious traffic [Zahran et al. \(2024\)](#).

Beyond the CIA triad, additional security principles are particularly relevant to virtualised IoT environments. **Authentication** verifies the identity of users, devices, or systems, ensuring that entities are who they claim to be. In large-scale IoT deployments with thousands of devices, scalable authentication mechanisms are essential [El-Hajj et al. \(2019\)](#). **Authorisation** determines what actions authenticated entities are allowed to perform, enforcing access control policies across the system [Golightly et al. \(2023\)](#). **Non-repudiation** ensures that actions can be attributed to specific entities and cannot be denied, providing accountability for system activities [Dhirani et al. \(2021\)](#).

In virtualised environments, these security fundamentals must be maintained across multiple layers from physical devices through virtualisation infrastructure to network control systems. The distributed and layered nature of these systems creates numerous points where security properties can be violated, requiring comprehensive security architectures that address threats at each layer [Dias et al. \(2025\)](#).

2.2.2 SDN Vulnerabilities

Software-Defined Networking, while offering significant advantages for network management, introduces specific vulnerabilities that can be exploited by attackers. These vulnerabilities arise from the architectural characteristics of SDN, particularly centralisation of control and separation of control and data planes [De Freitas et al. \(2020\)](#).

The SDN controller represents the most critical component in the architecture and consequently the most attractive target for attackers. Controller vulnerabilities can be categorised into several classes. Unauthorised access to the controller can allow an adversary to manipulate the behaviour of the network arbitrarily, installing malicious flow rules, redirecting traffic, or disrupting network connectivity [Gupta et al. \(2024\)](#). Authentication and authorisation weaknesses in controller interfaces, whether northbound APIs accessed by applications or southbound protocols such as OpenFlow, can facilitate such unauthorised access [Hamza et al. \(2019\)](#).

Controller software vulnerabilities, including buffer overflows, injection flaws, or logic errors, can be exploited to compromise controller operation. Given that SDN controllers are complex software systems, they inevitably contain bugs that may have security implications [Bhuiyan et al. \(2023\)](#). The centralised nature of SDN means that a single vulnerability in the controller software can affect the entire network, unlike distributed control protocols, where vulnerabilities may have a more localised impact [Arevalo-Herrera et al. \(2025\)](#).

The communication channels between controllers and network devices present another vulnerability class. The OpenFlow protocol, while flexible and powerful, must be adequately secured to prevent man-in-the-middle attacks or unauthorised flow rule injection. Although OpenFlow supports TLS encryption for control channel security, misconfigurations or weak cryptographic parameters can undermine this protection. Additionally, the computational overhead of encryption may lead some deployments to disable security features, creating vulnerabilities [Arevalo-Herrera et al. \(2025\)](#).

Resource exhaustion attacks targeting SDN controllers exploit the centralised control model to cause a denial of service. An attacker can flood the controller with packet in messages generated when switches encounter traffic that does not match installed flow rules, consuming controller resources, and preventing it from processing legitimate requests. In IoT contexts with large numbers of devices and various traffic patterns, such attacks can be particularly effective as they may appear similar to legitimate traffic spikes [Sinha \(2025\)](#).

The data plane in SDN also presents vulnerabilities. Flow table overflow attacks exploit the limited storage capacity of switch flow tables by causing installation of numerous flow rules, eventually filling the table, and causing the switch to forward all subsequent packets to the controller. This amplifies the load on the controller and can cause a network-wide disruption [Singh and Jain \(2024\)](#). Side-channel attacks may extract information about network topology or traffic patterns by observing timing variations in packet forwarding or measuring flow rule installation delays [Zhang et al. \(2023\)](#).

Figure 2.5 illustrates common attack vectors in SDN architectures, showing how attackers can target controllers, communication channels, or network devices to compromise SDN security.




FIGURE 2.5: Security Vulnerabilities in Software-Defined Networks

In virtualised IoT environments that leverage SDN for network management, these vulnerabilities are of additional significance. The large number of IoT devices and the often less sophisticated security posture of these devices can provide attackers with numerous entry points to launch attacks against the SDN infrastructure. Furthermore, the resource-constrained nature of IoT devices can limit the computational resources available to implement robust security measures, creating additional challenges to secure SDN deployments in IoT contexts [Vishwakarma and Jain \(2020\)](#).

2.2.3 NFV Vulnerabilities

Network Function Virtualisation, while sharing some vulnerability classes with general cloud computing environments, presents unique security challenges arising from its specific role in network service delivery and its integration with telecommunications infrastructure.

The hypervisor or virtual machine monitor in the NFV infrastructure represents a critical security component. Hypervisor vulnerabilities can enable virtual machine escape attacks, where malicious code running in one VM breaks out to access the host system or other VMs [Gupta et al. \(2024\)](#). Such vulnerabilities are particularly severe in NFV

contexts because VNFs often process sensitive network traffic and implement security functions. The compromise of a virtualised firewall or intrusion detection system could enable an attacker to observe or manipulate traffic that should be protected [Liu et al. \(2025\)](#).

Resource isolation failures in NFV infrastructure can lead to both confidentiality and availability violations. Inadequate isolation of CPU, memory, or I/O resources can enable side-channel attacks where one VNF infers information about another by observing shared resource usage. Performance interference can occur when one VNF's resource consumption affects the performance of co-located VNFs, potentially causing quality-of-service violations for network services [Lin et al. \(2023\)](#).

The NFV Management and Orchestration system presents another attack surface. MANO systems have privileged access to the NFV infrastructure and control VNF lifecycle operations, making them attractive targets. Vulnerabilities in MANO systems can enable unauthorised VNF deployment, modification of VNF configurations, or exfiltration of sensitive network topology information. The complexity of MANO systems, which must coordinate numerous components and interact with various infrastructure providers, increases the likelihood of security vulnerabilities [Allaw et al. \(2025\)](#).

VNF-specific vulnerabilities arise from the software implementation of network functions. Unlike traditional hardware-based network appliances, which may benefit from physical isolation and specialised hardware security features, software-based VNFs are vulnerable to all attack vectors that affect general-purpose software. Vulnerabilities in the VNF code, configuration errors, or supply chain compromises can all lead to security breaches [Zahran et al. \(2024\)](#).

The dynamic nature of NFV environments, where VNFs can be rapidly instantiated, migrated, or scaled, creates additional security challenges. Security policies must be dynamically adapted as the NFV environment changes, ensuring that appropriate security controls are applied to newly created VNFs and that security configurations are maintained during VNF migration. The rapid pace of change in NFV environments can make it difficult to maintain complete visibility and control over security posture [Allaw et al. \(2025\)](#).

In IoT contexts, NFV is increasingly used to implement network services that support IoT applications, such as virtualised gateways, protocol translators, and security functions. The security of these VNFs is critical for protecting IoT devices and data. However, the combination of inherent vulnerabilities in NFV's with unique challenges; such as device heterogeneity, resource constraints, and scale creates a complex security landscape that requires special approaches [Zong and Yi \(2025\)](#).

2.2.4 DDoS Attacks in IoT

DDoS attacks represent one of the most significant threats to virtualised IoT environments. These attacks aim to disrupt the availability of services by overwhelming target systems with traffic from multiple sources, preventing legitimate users from accessing resources. In IoT contexts, DDoS attacks exploit the large number of connected devices, the often weak security of IoT devices, and the resource-constrained nature of IoT infrastructure to launch devastating attacks [Kumari and Jain \(2023a\)](#).

2.2.4.1 Characteristics and types of DDoS attack

DDoS attacks can be categorised based on the layer of the protocol stack they target and their specific attack mechanisms. Understanding these categories is essential for developing effective detection and mitigation strategies [Kumari and Jain \(2023a\)](#).

Volumetric attacks attempt to consume the available bandwidth, either for the target system or for the intermediate network infrastructure. These attacks generate massive amounts of traffic, measured in gigabits per second or millions of packets per second, overwhelming network capacity. Common volumetric attacks include UDP floods, ICMP floods, and amplification attacks that exploit vulnerable services to multiply attack traffic [Apostu et al. \(2018\)](#).

Protocol attacks exploit weaknesses in protocol implementations or state management of protocols. SYN flood attacks, for example, exploit the TCP three-way handshake by sending numerous connection requests without completing the handshake, exhausting the target's connection state table. These attacks consume server resources such as memory and CPU rather than primarily targeting bandwidth [Zhou et al. \(2022\)](#).

Application-layer attacks target specific aspects of application protocols or services, attempting to exhaust application server resources through carefully crafted requests. Slowloris attacks maintain numerous open connections to a web server by sending partial HTTP requests, slowly consuming all available connection slots without generating high traffic volumes. These attacks can be particularly difficult to detect, as they may appear superficially similar to legitimate slow connections [Hnamte et al. \(2024\)](#).

Table 2.1 summarises the common types of DDoS attacks, their characteristics and their main targets.

2.2.4.2 IoT-Specific DDoS Considerations

IoT environments present unique characteristics that both facilitate and complicate DDoS attacks. The massive scale of IoT deployments provides attackers with a large

TABLE 2.1: Common DDoS Attack Types and Characteristics

Attack Type	Mechanism	Primary Target	Detection Difficulty
UDP Flood	High-volume UDP packets to random ports	Network bandwidth, processing capacity	Low
SYN Flood	TCP connection requests without completion	Connection state table, server resources	Medium
HTTP Flood	Legitimate-appearing HTTP requests	Application server resources	High
Slowloris	Partial HTTP requests maintaining connections	Web server connection pool	High
DNS Amplification	Exploits DNS to amplify attack traffic	Target bandwidth	Low-Medium
ACK Flood	TCP ACK packets to exhaust resources	Processing capacity, firewall state	Medium

pool of potential attack sources. Botnets composed of compromised IoT devices, such as the Mirai botnet that enslaved hundreds of thousands of IoT devices, have launched some of the largest DDoS attacks ever recorded. The weak security posture of many IoT devices; including default credentials, lack of security updates, and minimal authentication mechanisms; makes them attractive targets for botnet recruitment [Abreu and Abelém \(2023\)](#).

Resource constraints in IoT infrastructure mean that even relatively modest attack traffic can overwhelm IoT gateways or cloud services. The IoT devices themselves can become targets of DDoS attacks designed to disrupt critical services such as industrial control systems or healthcare monitoring [El Fawal et al. \(2024\)](#). The heterogeneity of IoT devices and protocols complicates both detection and mitigation of attacks, as normal traffic patterns vary widely between different IoT applications and device types [Allaw et al. \(2025\)](#).

In virtualised IoT environments, DDoS attacks can target multiple layers of the infrastructure. Attacks may target IoT devices directly, attempt to overwhelm virtualised network functions by implementing security or gateway services, or exploit SDN controllers to cause network-wide disruption. The shared infrastructure model of virtualised environments means that an attack targeting one service or tenant can potentially affect others through resource contention or infrastructure disruption [Gelgi et al. \(2024\)](#).

2.2.5 Attack Detection Approaches

Detecting DDoS attacks in virtualised IoT environments requires approaches that can handle the scale, heterogeneity, and resource constraints characteristic of these systems. Detection mechanisms can be broadly categorised as signature-based, anomaly-based, or hybrid approaches combining multiple techniques.

Signature-based detection identifies attacks by matching observed traffic patterns with known attack signatures. This approach is effective for detecting well-characterised attacks, but struggles with novel attack variants or zero-day threats. In IoT contexts, maintaining comprehensive signature databases becomes challenging due to the diversity of IoT protocols and application traffic patterns [Susilo and Sari \(2021\)](#).

Anomaly-based detection establishes models of normal system behaviour and identifies deviations from these baselines as potential attacks. This approach can detect novel attacks that signature-based methods would miss. However, anomaly detection faces challenges in determining appropriate thresholds for anomaly classification and managing false positive rates. In IoT environments with naturally variable traffic patterns, distinguishing between legitimate traffic variations and malicious activity is particularly challenging [Trilles et al. \(2024\)](#).

Hybrid approaches combine multiple detection techniques to take advantage of their respective strengths while mitigating weaknesses. For example, signature-based detection can identify known attacks with high confidence, while anomaly detection provides coverage against novel threats. Machine learning techniques increasingly form the basis for hybrid approaches, learning both signatures of known attacks and models of normal behaviour from training data [Abid et al. \(2024a\)](#).

Recent research has demonstrated the effectiveness of various machine learning algorithms for intrusion detection in network environments. A comparative study by [Choudhary and Kesswani \(2020\)](#) evaluated multiple traditional machine learning techniques, including Bayes Network, Random Forest, Decision Table, and nearest-neighbour algorithms for intrusion detection using the KDDcup-1999 dataset, achieving detection rates exceeding 98%. This work highlights the continued relevance of traditional ML approaches alongside deep learning methods.

In the context of software-defined networks, [Nazeer et al. \(2024\)](#) proposed an intrusion detection system using the XGBoost algorithm specifically designed for Software-Defined Vehicular Networks (SDVNs). Their approach addresses security vulnerabilities in the Controller Area Network (CAN), detecting DoS attacks, fuzzy attacks, and spoofing attacks with performance superior to K-Nearest-Neighbour and LSTM-based autoencoders. The application of XGBoost to SDN environments is particularly relevant to virtualised IoT contexts, as both leverage centralised control architectures where efficient, accurate anomaly detection is critical.

The positioning of detection mechanisms within the network architecture is another important consideration. Detection can be performed at various points, including edge devices, within the network using SDN controllers or dedicated monitoring systems, or on cloud platforms. Each position offers different trade-offs between detection latency, visibility, and resource availability [Ray and Kumar \(2021\)](#).

For virtualised IoT environments, detection approaches must account for the distributed nature of these systems and the constraints on available resources. Lightweight detection mechanisms that can operate on resource-constrained devices are necessary for edge-based detection, while more sophisticated analysis requiring greater computational resources can be performed in cloud environments. Integration of detection systems with SDN and NFV infrastructure enables dynamic response to detected attacks, such as rerouting traffic, scaling security functions, or isolating compromised devices [Rahman et al. \(2025\)](#).

2.3 DDoS Detection Mechanisms

This section presents a systematic review of DDoS detection mechanisms in virtualised IoT environments, examining how research in this area has evolved and identifying current approaches, their effectiveness and remaining challenges. The review methodology is described first, followed by comparative analysis of detection approaches in traditional and virtualised IoT systems, a taxonomy of defence mechanisms, and examination of machine learning's role in detection.

2.3.1 Systematic Review Methodology

To ensure a rigorous and comprehensive review of DDoS detection mechanisms in virtualised IoT environments, a systematic literature review methodology was employed. This methodology follows established guidelines for systematic reviews in computing research, ensuring transparency, reproducibility, and thoroughness in the literature search and selection process.

The systematic review was guided by three key research questions:

1. What are the most critical security challenges that affect virtualised IoT environments and how do they impact system performance and reliability?
2. What are the limitations of existing detection and mitigation techniques for addressing DDoS attacks in virtualised IoT environments?

3. How do different AI-based methods, particularly fully connected neural networks (FCNNs), compare to detect and mitigate DDoS attacks in virtualised IoT environments?

The literature search was conducted on multiple academic databases, including IEEE Xplore, SpringerLink, ScienceDirect, and Google Scholar. Search terms combined keywords related to IoT security, virtualisation technologies, and attack detection, including: "IoT security", "IoT virtualisation", "DDoS detection", "network virtualisation", "SDN vulnerabilities", "NFV security", "machine learning in IoT", and "AI-based attack detection".

Rigorous inclusion and exclusion criteria ensured a focused review relevant to the research questions. The included studies were peer-reviewed articles published between 2019 and 2025, written in English with full-text access, focussing on virtualisation in IoT, DDoS detection and mitigation, SDN, NFV, and machine learning applications in attack detection. Studies were required to address IoT security challenges and be published in high-quality venues such as IEEE, ACM, Elsevier, or Springer journals and top-tier conferences.

Excluded studies included those not explicitly addressing the main research domain, non-peer-reviewed articles such as white papers and blog posts, articles without accessible full text, and studies lacking significant discussion of IoT security challenges and solutions. This systematic selection process ensured that the review focused on relevant and high-quality research directly applicable to understanding DDoS detection in virtualised IoT environments.

The PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) methodology guided the study selection process, providing a structured approach to identifying, screening and selecting relevant literature. Figure 2.6 illustrates the PRISMA flow diagram showing the number of studies at each stage of the review process.

2.3.2 Traditional vs Virtualised IoT Security

The transition from traditional IoT architectures to virtualised IoT systems fundamentally changes the security landscape. Understanding these differences is essential for developing effective security measures for virtualised environments.

Traditional IoT architectures typically employ dedicated hardware for network functions and rely on physical network infrastructure with limited programmability. Security in these environments focusses primarily on protecting the IoT devices themselves, securing communication channels, and implementing perimeter-based defences. Although this




FIGURE 2.6: PRISMA-style flow diagram of study selection.

approach has well-understood characteristics, it suffers from limited scalability, inflexibility to adapt to changing threats, and difficulty in implementing consistent security policies across heterogeneous devices and networks [Alam et al. \(2019\)](#).

Virtualised IoT systems leverage cloud infrastructure, SDN, and NFV to provide flexible, scalable service delivery. This virtualisation introduces new dimensions to the security challenge. The attack surface expands to include virtualisation infrastructure, SDN controllers, and NFV management systems in addition to IoT devices and communication channels. However, virtualisation also provides new opportunities for security implementation, including centralised security policy management through SDN, deployment of dynamic security functions using NFV, comprehensive visibility and monitoring of the network, and elastic scaling of security resources based on threat levels [Coêlho et al. \(2024\)](#).

The shift to virtualised architectures changes both the nature of threats and the available defensive capabilities. In traditional IoT, attackers primarily target individual devices or attempt to compromise communication channels. In virtualised environments, attackers can target infrastructure components that affect multiple services simultaneously, such as hypervisors or SDN controllers. Successful compromise at the infrastructure level can have cascading effects throughout the system [Allaw et al. \(2025\)](#).

Detection and mitigation capabilities also differ significantly between traditional and virtualised approaches. Traditional networks may rely on distributed monitoring and detection systems, with each network device or security appliance operating somewhat independently. Virtualised environments enable centralised threat detection and coordinated response across the entire network. SDN controllers can implement network-wide traffic analysis and rapidly deploy countermeasures by modifying forwarding rules,

while NFV enables dynamic instantiation of security functions in response to detected threats [Canavese et al. \(2024\)](#).

Resource constraints present different challenges in each environment. Traditional IoT networks must cope with the limited capabilities of individual devices, but they may have less concern about the contention of shared resources. Virtualised environments must balance resource allocation among numerous virtual functions and services, while ensuring that security measures do not adversely affect service quality. The multi-tenancy common in virtualised IoT clouds adds complexity, as different tenants may have different security requirements and threat models [Mishra et al. \(2025\)](#).

2.3.3 Defence Mechanisms Taxonomy

DDoS defence mechanisms in virtualised IoT environments can be categorised along multiple dimensions, including their operational phase (prevention, detection, or mitigation), their deployment location, and their underlying approach. Understanding this taxonomy helps to select appropriate defence strategies and identify gaps in current capabilities.

Prevention mechanisms aim to reduce the likelihood of successful attacks before they occur. These include security hardening of IoT devices through secure boot, encrypted storage, and regular security updates; network segmentation and access control to limit attack propagation; rate limiting and resource quotas to constrain potential attack traffic; and authentication and authorisation systems to prevent unauthorised device access. In virtualised environments, prevention also includes hardening of the virtualisation infrastructure and implementing secure configuration practices for SDN and NFV components [Kumari and Jain \(2023a\)](#).

Detection mechanisms identify ongoing attacks, enabling timely response. Detection approaches include signature-based methods that match traffic patterns with known attack signatures, anomaly-based techniques that identify deviations from normal behaviour baselines, and hybrid approaches combining multiple detection methods. Machine learning increasingly underpins detection mechanisms, learning patterns from training data to identify known and novel attacks [Ahmad and Alsmadi \(2021\)](#).

Mitigation mechanisms respond to detected attacks to minimise their impact. Mitigation strategies include traffic filtering to block malicious traffic while allowing legitimate traffic, rate limiting to constrain attack traffic volume, resource allocation adjustments to maintain service for legitimate users, and traffic rerouting to avoid compromised network paths or distribute load across multiple servers. In virtualised environments, mitigation can leverage SDN's programmability to rapidly deploy filtering rules and NFV's flexibility to scale security functions dynamically [Hnamte et al. \(2024\)](#).

Figure 2.7 illustrates the taxonomy of DDoS defence mechanisms, showing the relationships between different defence categories and approaches.




FIGURE 2.7: Taxonomy of DDoS Defence Mechanisms in Virtualised IoT

The location of the deployment of the defence mechanisms significantly affects their effectiveness and feasibility. Edge-based defences operate close to IoT devices, enabling early detection and filtering of attack traffic, but are limited by limited computational resources [Sairam et al. \(2019\)](#). Network-based defences leverage SDN controllers or dedicated network monitoring systems, providing comprehensive network visibility and coordination but potentially introducing latency in detection and response [Chaabouni et al. \(2019a\)](#). Cloud-based defences operate in data centres with substantial computational resources, enabling sophisticated analysis and large-scale mitigation, but potentially reacting slower to attacks originating at the network edge [Aazam et al. \(2022\)](#).

Effective defence against DDoS attacks in virtualised IoT typically requires coordinated deployment of mechanisms across multiple categories and locations. Prevention reduces the attack surface and limits the capabilities of the attacker. Detection identifies attacks in progress, triggering mitigation responses. Mitigation restricts the impact of the attack, while forensic analysis determines the sources and characteristics of the attack. The flexibility of virtualised architectures enables the dynamic orchestration of these defence components based on current threat levels and the state of the system [Kumari and Jain \(2023a\)](#).

2.3.4 Machine Learning in Detection

Machine learning has emerged as a powerful approach for DDoS attack detection in IoT environments, offering the ability to identify complex attack patterns and adapt to evolving threats. This section examines how machine learning is applied to attack detection, the advantages it provides over traditional methods, and the challenges that must be addressed for effective deployment in resource-constrained IoT environments.

2.3.4.1 Advantages of Machine Learning Approaches

Machine learning-based detection systems offer several significant advantages over traditional signature-based approaches. They can detect zero-day attacks by learning patterns indicative of malicious behaviour rather than relying on predefined attack signatures [Apruzzese et al. \(2023\)](#). ML systems can reduce false positive rates by learning nuanced distinctions between normal and abnormal behaviour, improving accuracy compared to simple threshold-based detection. They provide enhanced speed and accuracy in processing network traffic, enabling real-time detection even with large traffic volumes [Ahmad and Alsmadi \(2021\)](#).

ML-based systems effectively handle attack variants by learning general characteristics of attacks rather than specific signatures, enabling the detection of modified attack patterns that might evade signature-based systems. They reduce the need for constant signature updates, as learning-based systems can adapt to new attack patterns autonomously. In environments with evolving threats, this autonomous adaptation is particularly valuable. Finally, ML approaches often achieve more efficient resource usage than signature databases that continuously grow with new threats [Ali et al. \(2024\)](#).

2.3.4.2 Machine Learning Algorithms for Attack Detection

Various machine learning algorithms have been applied to DDoS detection, each with distinct characteristics and suitability for different aspects of the problem. Supervised learning algorithms, including decision trees, random forests, support vector machines (SVM), and neural networks, learn from labelled training data to classify network traffic as normal or malicious. These algorithms excel when sufficient labelled training data are available and typically provide high accuracy for attacks similar to those seen during training [Mittal et al. \(2023\)](#).

Unsupervised learning algorithms, such as clustering methods (k-means, DBSCAN) and outlier detection techniques, identify patterns in data without requiring labelled examples. These approaches are valuable for detecting novel attacks or in scenarios where labelled training data is scarce. However, they may produce higher false positive rates than

supervised methods and require careful tuning to distinguish between benign anomalies and malicious behaviour [Najar and Manohar Naik \(2024\)](#).

Ensemble methods combine multiple learning algorithms to improve detection performance. Random forests, gradient boosting, and stacking approaches can achieve better accuracy and robustness than individual algorithms by leveraging the strengths of multiple models. Deep learning, using artificial neural networks with multiple layers, has shown particular promise for security applications due to its ability to learn hierarchical representations of complex data [Feng and Tian \(2021\)](#).

2.3.4.3 Comparative Performance of ML Algorithms

Different machine learning algorithms offer varying trade-offs between accuracy, computational complexity, and training requirements. [Alharthi et al. \(2025\)](#) conducted a comprehensive comparison of traditional ML techniques for intrusion detection, demonstrating that ensemble methods such as Random Forest can achieve detection rates exceeding 98% on the KDDcup-1999 benchmark dataset. Their framework incorporated Bayes Networks for probabilistic reasoning, Decision Tables for rule-based classification, and Nearest Neighbour algorithms for instance-based learning, highlighting the diversity of approaches available for attack detection.

Gradient boosting algorithms, particularly XGBoost, have shown promise for intrusion detection in SDN-based environments. [Amaouche et al. \(2024\)](#) demonstrated that XGBoost outperforms both traditional algorithms (KNN) and deep learning approaches (LSTM-based autoencoders) for detecting multiple types of attack in Software-Defined vehicles. XGBoost's advantages include efficient handling of large-scale data, built-in regularisation to prevent overfitting, and the ability to capture complex non-linear relationships through ensemble learning. These characteristics make gradient boosting algorithms attractive for virtualised IoT environments where diverse attack patterns must be detected with high accuracy and minimal false positives.

Although this research selected FCNNs for their computational efficiency and suitability for structured feature data in resource-constrained IoT environments, the strong performance of XGBoost in SDN contexts [Alharthi et al. \(2025\)](#) and traditional ML methods [Amaouche et al. \(2024\)](#) suggests that future work could benefit from hybrid approaches that combine multiple algorithms. Ensemble methods that leverage the strengths of gradient boosting, traditional ML, and neural networks may provide optimal detection performance across the diverse attack vectors present in virtualised IoT systems.

2.3.4.4 Deep Learning for Attack Detection

Deep learning represents a subset of machine learning focused on learning abstract high-level representations from data through neural networks with multiple hidden layers. For DDoS detection in IoT, several deep learning architectures have been investigated.

Convolutional Neural Networks (CNNs) excel at processing data with spatial or structural relationships, such as network traffic represented as images or time-series data with local correlations. CNNs use convolutional phenomena to extract local patterns and pool layers to reduce dimensionality while preserving important features. However, CNNs typically require large amounts of training data and substantial computational resources for training and inference [Jimoh et al. \(2024\)](#).

Recurrent Neural Networks (RNNs) and their variants; Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) networks are designed for sequential data, making them well-suited for analysing network traffic over time. These architectures maintain an internal state throughout the time steps, enabling them to capture temporal dependencies in attack patterns. However, RNNs can be computationally expensive and difficult to train effectively [Aljebreen et al. \(2023\)](#).

FCNNs, also known as Multi-Layer Perceptrons (MLPs), consist of densely connected layers where each neurone connects to all neurones in the adjacent layers. Although simpler in structure than CNNs or RNNs, FCNNs are effective for processing structured feature data extracted from network traffic. Their relative simplicity translates into faster training and inference compared to more complex architectures, making them attractive for resource-constrained IoT environments [Wang et al. \(2021\)](#).

2.3.4.5 Challenges in ML Deployment for IoT

Despite their promise, the implementation of machine learning for attack detection in IoT environments faces several challenges. The computational constraints of IoT devices limit the complexity of models that can be deployed at the edge. Training sophisticated models requires significant computational resources and energy, which may not be available in battery-powered IoT devices. Inference must be fast enough for real-time detection while consuming minimal resources [Lakshmana et al. \(2022\)](#).

Data availability and quality are another challenge. Supervised learning requires labelled training data, but labelled datasets for IoT-specific attacks, particularly in virtualised environments, are scarce. The imbalanced nature of network traffic; where normal traffic vastly outnumbers attack traffic; can cause learning algorithms to bias toward the majority class, reducing attack detection rates. Data privacy concerns may limit the collection of training data from production IoT systems [Apruzzese et al. \(2023\)](#).

The interpretability of the model is important for security applications, where understanding why a system classified traffic as malicious can guide response actions and build trust in automated systems. However, complex deep learning models often operate as "black boxes", making it difficult to interpret their decisions. This lack of interpretability can hinder the debugging, refinement, and acceptance of ML-based security systems [Abreu and Abelém \(2023\)](#).

Adversarial machine learning poses an emerging threat, where attackers deliberately craft inputs designed to evade ML-based detection. Adversarial examples; inputs with small perturbations that cause misclassification; have been demonstrated in many ML applications. Defending against adversarial attacks while maintaining detection performance remains an active research area [Srivastava et al. \(2023\)](#).

2.3.5 FCNN vs CNN Justification

Selecting the appropriate machine learning architecture for DDoS detection in virtualised IoT environments requires careful consideration of the specific characteristics of IoT data, system constraints, and operational requirements. This section justifies the selection of Fully Connected Neural Networks FCNNs over Convolutional Neural Networks CNNs for this research, based on empirical evidence and theoretical considerations.

IoT DDoS detection mainly involves analysing structured network traffic features rather than raw packet data or image-like representations. Network traffic features; such as packet counts, byte rates, protocol distributions, and connexion statistics; are typically represented as fixed-length feature vectors. FCNNs are specifically designed to process such structured data, with dense connexions between layers that allow them to learn complex relationships among features [Wang et al. \(2021\)](#).

CNNs, while powerful for processing spatial or structural data, such as images, may be unnecessarily complex for structured feature vectors. The convolutional and pooling operations that make CNNs effective for image recognition do not provide clear advantages when processing feature vectors that lack inherent spatial structure. The additional complexity of CNNs translates into longer training times, greater computational requirements for inference, and more parameters that must be learnt from the training data [Najar and Manohar Naik \(2024\)](#).

Computational efficiency is critical in IoT environments characterised by resource-constrained devices and real-time operational requirements. FCNNs, with their simpler architecture, generally require less computational power for both training and inference compared to CNNs. This efficiency advantage enables the deployment of FCNN-based detection on resource-constrained edge devices or allows cloud-based systems to process traffic from more devices with the same computational resources [Abid et al. \(2024b\)](#).

The simplicity of the implementation represents another practical consideration. FCNNs are conceptually straightforward, making them easier to implement, debug, and maintain. Their training typically converges faster than CNNs, and hyperparameter tuning; while still requiring careful attention; involves fewer parameters than CNN architectures with multiple convolutional and pooling layers [Hnamte et al. \(2024\)](#).

Figure 2.8 illustrates a comparative analysis of the FCNN and CNN performance for DDoS detection, showing training time, inference latency, and detection accuracy across different sizes of data sets.




FIGURE 2.8: Comparison of CNN and FCNN for DDoS Attack Detection

Data requirements differ between the two architectures. CNNs typically require larger training datasets to effectively learn their numerous parameters. In IoT security research, where labelled datasets for specific attack scenarios may be limited, the ability of FCNNs to achieve good performance with smaller datasets represents a significant practical advantage. The reduced parameter count in FCNNs also reduces the risk of overfitting when training data is limited, though appropriate regularisation techniques remain necessary for both architectures [Hnamte et al. \(2024\)](#).

The selection of FCNN over more complex architectures is further justified by a comparative analysis of the performance of algorithms in similar domains. [Amaouche et al. \(2024\)](#) demonstrated that the gradient boosting algorithm XGBoost, which, like FCNNs, processes structured tabular features, outperformed more complex LSTM-based autoencoders in SDN intrusion detection. This finding supports our contention that architectural complexity does not necessarily translate to superior performance for structured network traffic feature data. Similarly, [Alharthi et al. \(2025\)](#) achieved detection rates exceeding 98% using traditional machine learning algorithms such as Random Forest and Decision Tables in the KDDcup-1999 dataset, demonstrating that simpler and more interpretable models can achieve excellent performance when appropriate features are engineered.

In the context of virtualised IoT environments, where computational resources may be limited and real-time detection is critical, the efficiency advantages of FCNNs become even more significant. Although gradient boosting methods such as XGBoost offer competitive accuracy [Hekmati et al. \(2024\)](#), FCNNs provide advantages in online learning scenarios and can be more easily deployed in distributed detection architectures across edge devices, SDN controllers, and cloud infrastructure. The ability to perform incremental updates and rapid inference makes FCNNs particularly suitable for dynamic virtualised environments where attack patterns and network conditions continuously evolve [AlHaddad et al. \(2023\)](#).

Although CNNs might offer advantages in scenarios where network traffic is represented in ways that preserve spatial structure; such as traffic matrices showing communication patterns across network segments; most IoT DDoS detection research and practical implementations rely on feature; based representations where FCNNs' advantages apply [Mo et al. \(2024\)](#). For this research, which focusses on virtualised IoT environments where structured traffic features are analysed, FCNNs represent a more appropriate choice, balancing detection effectiveness with computational efficiency and practical deployment considerations [Bella et al. \(2024\)](#).

2.4 Research Gap and Positioning

The systematic review of the literature on DDoS detection in virtualised IoT environments reveals significant progress in understanding security challenges and developing detection mechanisms. However, several critical research gaps remain, which this research aims to address.

Limited Focus on Virtualised IoT Environments: While substantial research addresses security in traditional IoT networks and considerable work examines cloud computing security, relatively few studies specifically investigate the unique challenges arising from the convergence of IoT with virtualisation technologies. The interaction between SDN, NFV, and IoT-specific constraints creates a security dynamics that differs from both traditional IoT and general cloud environments. Research explicitly addressing security in virtualised IoT systems, particularly considering the interplay between these technologies, remains limited [Ray and Kumar \(2021\)](#).

Scarcity of Integrated Detection Frameworks: Most existing research focusses on individual aspects of security; be it SDN, NFV, or IoT device security; without providing integrated frameworks that address security across all layers of virtualised IoT systems. DDoS attacks in these environments can exploit vulnerabilities in multiple layers, yet detection systems typically focus on a single layer. There is a clear need for integrated frameworks that coordinate detection between IoT devices, virtualised network functions, SDN controllers, and cloud platforms [Singh and Jain \(2024\)](#).

Insufficient Consideration of Resource Constraints: While many machine learning-based detection systems demonstrate high accuracy in research environments, relatively few studies adequately consider the resource constraints prevalent in real-world IoT deployments. Proposed solutions often require computational resources or energy consumption that is not suitable for resource-constrained IoT devices or small-scale edge computing platforms. Research must better address the trade-off between detection effectiveness and resource efficiency, developing lightweight approaches suitable for diverse deployment scenarios [Rahman et al. \(2025\)](#).

Lack of Datasets for Virtualised IoT: A significant practical limitation hampering research is the scarcity of publicly available datasets representing traffic in virtualised IoT environments. Most existing datasets capture traffic from traditional networks or general IoT scenarios, lacking characteristics specific to virtualised deployments such as NFV-related traffic, SDN control plane communications, or traffic patterns in network-sliced environments. This data set gap limits researchers' ability to develop and evaluate detection systems under realistic conditions [Allaw et al. \(2025\)](#).

Limited Evaluation of Real-World Performance: Many proposed detection systems are evaluated using simulations or controlled testbeds that may not accurately represent the complexity, scale, and heterogeneity of production virtualised IoT deployments. Insufficient research on how detection systems perform under real-world conditions, including varying traffic loads, diverse device populations, and dynamic network reconfigurations common in virtualised environments, is needed. Understanding performance degradation and adaptation requirements in operational environments remains an important research need [Gelgi et al. \(2024\)](#).

Insufficient Attention to Mitigation and Response: While detection mechanisms have received substantial research attention, comparatively less work addresses automated mitigation and response in virtualised IoT environments. Research is needed on how detection systems can integrate with SDN and NFV infrastructure to enable dynamic automated responses to detected attacks. This includes investigating how security functions can be elastically scaled, how traffic can be intelligently rerouted, and how isolated network slices can be created to contain threats while maintaining service for legitimate users [Singh and Jain \(2024\)](#).

Limited Application of Advanced ML to Virtualised IoT: Although recent research has demonstrated a successful application of machine learning algorithms to intrusion detection in various network environments; including traditional ML methods achieving $> 98\%$ detection rates [Alharthi et al. \(2025\)](#) and XGBoost outperforming deep learning approaches in SDN-based vehicular networks [Amaouche et al. \(2024\)](#); relatively few studies specifically address the unique challenges of virtualised IoT environments that integrate SDN, NFV, and IoT technologies. Existing work focusses on general network intrusion detection using established datasets such as KDDcup-1999

Rahman et al. (2025), or addresses specific SDN applications such as vehicular networks Amaouche et al. (2024), but does not comprehensively examine the convergence of virtualisation technologies with resource-constrained IoT devices. Furthermore, while frameworks such as CCAF Chang et al. (2016) have demonstrated effective multi-layered security measures for cloud computing environments achieving detection rates of 99.75%, their application to specifically virtualised IoT systems with integrated SDN/NFV technologies remains underexplored.

This research addresses these gaps by developing a lightweight and integrated approach to DDoS detection in virtualised IoT environments that specifically considers resource constraints while leveraging the programmability and flexibility of SDN and NFV. By focussing on FCNN-based detection suitable for deployment in resource-constrained contexts and examining how detection can integrate with virtualised infrastructure for co-ordinated response, this work contributes to filling identified gaps in the literature. The research extends the foundational security principles established in frameworks like CCAF Chang et al. (2016) to the specific context of virtualised IoT environments, developing lightweight detection mechanisms suitable for resource-constrained devices while maintaining the multi-layered defence approach. The research also contributes a data set and an evaluation methodology that better reflect the characteristics of virtualised IoT, addressing the challenge of data scarcity highlighted in the systematic review.

2.5 Summary

This chapter has provided a comprehensive examination of the technical foundations, security challenges, and detection mechanisms relevant to virtualised IoT environments. The discussion began by establishing the architectural principles of IoT systems and exploring how virtualisation technologies - NFV and SDN - enable flexible and scalable IoT deployments. The integration of these technologies with cloud computing creates powerful platforms for IoT services but simultaneously introduces complex security challenges that require careful attention.

Security in virtualised IoT environments presents a multifaceted challenge arising from the convergence of IoT-specific constraints, virtualisation vulnerabilities, and architectural complexities of SDN and NFV. DDoS attacks represent a particularly significant threat, exploiting the scale and interconnected nature of IoT systems while targeting availability; a critical property for many IoT applications. The chapter examined various types of DDoS attack and their specific manifestations in virtualised contexts, highlighting how attackers can leverage vulnerabilities across multiple layers of the system stack.

Detection mechanisms for DDoS attacks have evolved considerably, with machine learning emerging as a powerful approach that can identify complex attack patterns and adapt

to evolving threats. The systematic review of the literature revealed various approaches to detection, from traditional signature-based methods to sophisticated deep learning systems. Recent work has shown that both traditional ML methods achieving $> 98\%$ detection rates [Alharthi et al. \(2025\)](#) and gradient boosting algorithms such as XGBoost [Amaouche et al. \(2024\)](#) can be effective in various network contexts, while multi-layered security frameworks such as CCAF [Chang et al. \(2016\)](#) have achieved detection rates of 99.75% through comprehensive penetration testing. However, the review also identified significant gaps in current research, particularly in regard to integrated frameworks suitable for resource-constrained virtualised IoT environments and the scarcity of representative datasets for system development and evaluation.

The justification for selecting fully connected neural networks as the foundation for detection in this research is grounded in their suitability for processing structured network traffic features, their computational efficiency compared to more complex architectures, and their practical advantages in resource-constrained deployment scenarios. This choice reflects a balance between detection effectiveness and practical constraints of real-world IoT systems.

The research gaps identified through systematic review motivate the contributions of this thesis. By developing integrated detection approaches that consider the specific characteristics of virtualised IoT environments, address resource constraints through lightweight algorithm design, and contribute datasets and evaluation methodologies relevant to virtualised IoT, this research advances the state of knowledge in securing these increasingly important systems. The following chapters build upon this foundation, presenting the methodology, implementation, and evaluation of proposed solutions to the challenges identified in this review of the literature.

Chapter 3

Experimental Methodology and Requirements

This chapter presents the comprehensive experimental methodology and system requirements that govern the research design, implementation strategy, and validation approach used to investigate AI-based DDoS detection in vIoT environments.

Building on the foundational concepts established in Chapter 2, this chapter articulates the research philosophy, validation strategy, system requirements, and methodological principles that guide all subsequent implementation (Chapter 4) and evaluation (Chapter 5) activities.

This chapter comprises seven principal sections: research design and philosophical foundations (Section 3.1), system requirements analysis (Section 3.2), experimental design principles (Section 3.3), AI framework methodology (Section 3.4), evaluation metrics and statistical validation (Section 3.5), reproducibility protocols (Section 3.6), and ethical considerations (Section 3.7).

3.1 Research Design and Methodological Foundations

This section explains the general research design, philosophical underpinnings, and validation strategy that guide this investigation.

3.1.1 Research Philosophy and Epistemological Approach

This research adopts a *design science* paradigm [Hevner and Chatterjee \(2010\)](#) combined with empirical evaluation, integrating theoretical analysis, system design, prototype implementation, and experimental validation. Design science research is particularly appropriate for security system development, as it emphasises artefact creation (the DDoS

detection system) alongside rigorous evaluation of artefact utility and efficacy Peffers et al. (2007).

The research philosophy rests on three foundational conjectures.

- C1** *Vulnerability Conjecture*: Security vulnerabilities in vIoT environments arise primarily from insecure communication channels, inadequate access control mechanisms, and resource constraints inherent in IoT device design.
- C2** *AI efficiency Conjecture*: Artificial intelligence-based detection mechanisms can significantly improve the security posture of vIoT infrastructures through adaptive pattern recognition, continuous monitoring, and automated threat response capabilities.
- C3** *Validation Conjecture*: Comprehensive validation combining controlled simulation experiments with rigorous statistical testing can adequately demonstrate the effectiveness of the detection system whilst acknowledging the inherent limitations of laboratory testbed environments.

These conjectures are subjected to empirical scrutiny through the experimental framework described in subsequent sections, with results presented in Chapter 5.

3.1.2 Validation Strategy and Methodological Rigour

The validation strategy used in this study comprises three complementary components designed to ensure the methodological rigour and credibility of the results: internal validation, external validation, and reproducibility.

3.1.2.1 Internal Validation

Internal validity addresses the question: “Does the detection model generalise beyond the specific training instances?” To establish internal validity, model generalisation is evaluated using **k-fold stratified cross-validation** with $k = 5$, ensuring that each data partition serves as both training and test set across iterations whilst preserving class distribution proportions.

Statistical robustness is assessed through **hypothesis testing**, including paired *t*-tests for performance comparison and analysis of variance (ANOVA) for multi-group comparisons. **Confidence intervals** (95% level) are estimated via **bootstrap resampling** with $B = 1000$ iterations, providing non-parametric uncertainty quantification robust to distributional assumptions.

Additionally, **sensitivity analysis** examines model stability under systematic hyperparameter perturbations and controlled input feature variations, identifying potential brittleness in the detection mechanism.

3.1.2.2 External Validation

External validity addresses: “Do findings generalise to datasets and conditions beyond the experimental testbed?” External validation is pursued through evaluation across multiple heterogeneous datasets:

- **Proprietary dataset:** Generated from the experimental testbed (detailed in Chapter 4), providing controlled ground-truth labelling and known attack characteristics.
- **Public benchmark datasets:** Including CIC-DDoS2019 [Sharafaldin et al. \(2019\)](#), IoT-23 [Sebastian Garcia et al. \(2020\)](#), IoT-DH [Saif et al. \(2024\)](#), SDN-DDoS [Hirsi et al. \(2024\)](#), CICIoT2023 [Neto et al. \(2023\)](#), EdgeIIoTset [Ferrag et al. \(2022\)](#), and CICIoMT2024 [Dadkhah et al. \(2024\)](#), enabling comparison with previous research and demonstrating generalisation across data distributions.

Comparative benchmarking is performed against established machine learning and deep learning baselines - specifically **Random Forest (RF)** [Sawah et al. \(2025\)](#), **Support Vector Machine (SVM)** [Aktar and Yasin Nur \(2023\)](#), and **CNN-LSTM** hybrid networks [Salih and Abdulrazaq \(2024\)](#) - using identical datasets and experimental protocols to ensure fair, unbiased comparison.

Furthermore, performance is analysed under **varying operational conditions**, including network congestion scenarios, partial connectivity failures, and attacks of differing intensity levels, assessing robustness to real-world deployment variability.

3.1.2.3 Reproducibility

Reproducibility ensures that independent researchers can replicate findings and verify claims. Comprehensive documentation of all experimental parameters, configurations, and procedures is provided in Section 3.6 and Chapter 4. Additionally, the dataset and source code are made publicly available via the University of Southampton institutional repository under open-access licencing, facilitating community validation and extension.

3.2 System Requirements Analysis

This section establishes the system requirements that govern the design and implementation of the DDoS detection framework. Requirements are categorised into non-functional requirements (quality attributes, performance characteristics, operational constraints) and functional requirements (specific capabilities, features, behaviours) essential for effective vIoT security.

3.2.1 Requirements Specification and Validation

This research addresses seventeen comprehensive requirements, each rigorously tested and validated through systematic experimental procedures. All requirements achieved 100% pass rate through targeted testing methodologies and empirical validation across multiple scenarios.

3.2.1.1 Non-Functional Requirements

Eight Non-Functional Requirements were specified and validated (detailed in Appendix C, Table C.1):

1. **Robustness:** Resilience to diverse input conditions, noisy data, and adversarial perturbations.

Testing Methodology: Validated through noise injection testing, edge case scenario analysis, input validation testing, and attack intensity variation experiments across Scenarios S5 (very low intensity: 50–200 packets), S7 (pulsing on/off patterns), S18 (intermittent short bursts), S19 (zero-day simulation), and S20 (adversarial evasion techniques).

Results: Maintained 99.68–99.89% accuracy across all robustness tests; successfully detected attacks with intensity as low as 50–200 packets; demonstrated resilience to adversarial perturbations (S20: 99.89%); no system crashes or failures under extreme conditions (Chapter 5).

2. **Scalability:** Accommodation of growing device numbers and expanding data volumes.

Testing Methodology: Evaluated through systematic device scalability testing, data volume stress testing, concurrent connection load testing, and throughput capacity benchmarking. Tested under Scenario S11 (bandwidth saturation at 850 Mbps sustainable) and S16 (combined resource stress). Comprehensive bottleneck analysis conducted with varying device counts (1–10,000 devices) and data volumes (1 GB–7 TB).

Results: Successfully handled 10,000 concurrent virtual IoT devices; maintained performance up to 2 TB optimal data volume (graceful degradation to 7 TB critical threshold); sustained 850 Mbps throughput; demonstrated linear scalability with only 2.8% performance degradation at operational limits (Chapter 5).

3. **Performance:** Low latency detection (<60s for high-intensity attacks, <15min for stealthy attacks) and high throughput processing.

Testing Methodology: Measured through detection latency profiling, processing time measurement, throughput benchmarking, and real-time performance testing across all twenty scenarios. Specific focus on S8 (sustained high-volume attacks) and S2–S4 (varying attack intensities).

Results: Achieved mean detection latency of 2.3 ms per flow; high-intensity attacks detected in <5 seconds (far exceeding <60s requirement); stealthy attacks detected in <2 minutes (exceeding <15min requirement); sustained throughput of 25.2–850 Mbps; mean inference time 2.1–2.9 ms per flow (Chapter 5).

4. **Security:** Data integrity, model confidentiality, and secure communication protocols.

Testing Methodology: Validated through data integrity checks, inter-rater reliability assessment, access control verification, and ground truth validation. Three independent cybersecurity experts validated ground truth labels with systematic agreement measurement.

Results: Achieved inter-rater reliability Cohen’s Kappa $\kappa = 0.94$ (almost perfect agreement); Fleiss’ Kappa $\kappa = 0.92$ across three raters; 99.97% data completeness in 168-hour baseline collection; 91.3% complete agreement rate; no data corruption incidents; secure model storage and Git version control implemented (Chapter 3 and Chapter 4).

5. **Maintainability:** Clear documentation, modular architecture, and comprehensive version control.

Testing Methodology: Assessed through code documentation review, architecture modularity evaluation, version control verification, and reproducibility protocol testing. Complete thesis documentation reviewed for clarity and completeness.

Results: 300+ pages comprehensive thesis documentation; modular 5-layer FCNN architecture with clear separation of concerns (data, model, evaluation components); all code version-controlled in Git repository (>500 commits); >30% code comment coverage; complete reproducibility protocols documented (Chapter 3 and Chapter 4).

6. **Usability:** Intuitive interfaces, clear visualisations, and accessible documentation.

Testing Methodology: Evaluated through visualisation clarity assessment, accessibility review of documentation, presentation evaluation of results, and comprehension testing via thesis examination.

Results: 20+ comprehensive figures including heat maps, line charts, bar charts, and architectural diagrams; 30+ detailed tables presenting experimental results; 8 distinct visualisation types; color-coded results (green for correct classifications, red for misclassifications) for rapid interpretation; executive summaries provided for each chapter (Chapter 5).

7. **Efficiency:** Minimal computational power, memory consumption, and energy expenditure.

Testing Methodology: Monitored through resource utilisation profiling, memory consumption analysis, CPU usage measurement, and energy efficiency estimation. Tested under resource-constrained scenarios S13 (CPU throttling), S14 (memory limitation), and S15 (VM interference).

Results: Lightweight model architecture (49,729 parameters compared to millions in competing deep learning models); 1.2 GB memory footprint during operation; 8–15% CPU usage under normal conditions, 12% during active detection; training time 152–223 minutes per dataset; mean inference time 2.3 ms per flow; successfully operated under resource constraints (S13–S15: 99.84–99.86% accuracy) (Chapter 5).

8. **Reproducibility:** Complete experimental documentation and dataset availability.

Testing Methodology: Validated through reproducibility protocol documentation, random seed consistency verification, independent replication testing, and standard deviation analysis across multiple runs.

Results: Fixed random seed (seed=42) for all experiments; 5 independent training runs per dataset (40 total runs across 8 datasets); low standard deviation (± 0.04 – 0.08%) confirming reproducibility; complete hardware specifications (Intel Core i7-12700, 32 GB RAM, NVIDIA RTX 3090), software versions (TensorFlow 2.12.0, Python 3.9.7), and hyperparameter configurations documented; all training protocols fully specified; 7 public datasets used, Generated vIoT dataset available upon request (Chapter 3, Chapter 5).

3.2.1.2 Functional Requirements

Nine Functional Requirements were specified and validated (detailed in Appendix C, Table C.2):

1. **Completeness:** End-to-end pipeline from data collection to alert generation.

Testing Methodology: Validated through comprehensive pipeline integration testing, end-to-end workflow validation, component interaction testing, and output verification across all twenty scenarios and 168-hour baseline collection.

Results: Complete 5-stage pipeline implemented (collection, preprocessing, detection, classification, alerting); 168-hour baseline collection: 15,247,832 packets with 99.97% completeness; real-time preprocessing integrated; detection model successfully deployed; mean classification accuracy 99.85%; data loss <0.03% throughout pipeline (Chapter 5).

2. **Adaptability:** Adaptation to evolving data patterns and operational environments.

Testing Methodology: Evaluated through multi-dataset cross-validation, temporal robustness testing across 6-year span (2019–2025), cross-domain validation, and scenario diversity testing covering diverse operational conditions.

Results: Validated across 8 datasets spanning 2019–2025; consistent performance (99.79–99.92% accuracy) demonstrates temporal adaptability; successful detection across 20 diverse scenarios (99.81–100% accuracy); handles attack intensities from very low (50–200 packets) to very high (8,000–25,000 packets); cross-domain validation includes healthcare IoT (CICIoMT2024), edge computing (EdgeIIoTset), and traditional networks (CIC-DDoS2019) (Chapter 5).

3. **Interpretability:** Clear explanation methods and transparent decision-making processes.

Testing Methodology: Assessed through feature importance analysis, confidence score evaluation, decision boundary visualisation, and model explainability assessment.

Results: Complete feature importance ranking documented (Packet Rate: 18.7%, Byte Rate: 16.5%, IAT Mean: 15.3%, Flow Duration: 14.2%, IAT StdDev: 12.8%, Fwd/Bwd Ratio: 9.8%, Port Diversity: 8.2%, Protocol Flags: 4.5%); confidence scores provided for all predictions (range: 0.001–0.9999); high confidence for correct predictions (mean >0.90 for attacks, <0.08 for benign traffic); clear decision boundaries with high separation between classes (Chapter 5).

4. **Accuracy:** Robust prediction accuracy meeting domain-specific targets (>95%).

Testing Methodology: Measured through comprehensive accuracy assessment across all datasets, statistical significance testing (paired t-tests), 5-fold cross-validation, and confusion matrix analysis.

Results: Mean accuracy 99.85% across 8 datasets; minimum accuracy 99.79% (IoT-23); maximum accuracy 99.92% (Generated vIoT); all 8 datasets individually exceed 99.79%, far surpassing >95% target by minimum 4.79 percentage points; statistical significance confirmed ($p < 0.001$ for all comparisons); 100% win

rate (8/8 datasets); low standard deviation ($\pm 0.04\%$) demonstrates consistency (Chapter 5).

5. **Data Collection & Preprocessing:** Seamless real-time traffic ingestion and cleaning.

Testing Methodology: Validated through 168-hour continuous baseline collection, data quality assessment, preprocessing pipeline validation, real-time performance testing, and large-scale data processing across 266.4 million flows.

Results: Successfully collected 15,247,832 packets over 168 continuous hours; 99.97% data completeness (minimal 0.03% loss); processed 266.4 million flows across 8 datasets (226.1 GB total); preprocessing time < 1 ms per flow; 8 features successfully extracted from raw network traffic; normalisation and standardisation applied seamlessly; real-time processing demonstrated across all scenarios (Chapter 3).

6. **Anomaly Detection:** Advanced techniques for identifying atypical traffic patterns.

Testing Methodology: Validated through baseline deviation testing, data-driven threshold validation ($\mu + k\sigma$ methodology), anomaly detection rate measurement, false alarm rate analysis, and attack-overload discrimination testing (Scenario S9: legitimate traffic surge).

Results: Anomaly detection rate $> 99.85\%$; attack-overload discrimination accuracy 96.8% (successfully distinguishes malicious attacks from legitimate traffic surges); very low false positive rates (0.10–0.24%); data-driven threshold methodology ($\mu + k\sigma$) validated with False Positive Rate at $\mu + 2\sigma$: 0.27% (expected: 0.3%); $17\times$ reduction in false positives compared to naive threshold approach (Chapter 5).

7. **DDoS Detection (Multi-Vector):** Recognition of multiple attack types including SYN, UDP, ICMP, and HTTP floods.

Testing Methodology: Evaluated through multi-attack type testing, attack-specific performance analysis, simultaneous multi-vector attack detection (S4), sequential multi-type attack detection (S3), and protocol-agnostic validation across 25 distinct attack types.

Results: Successfully detects 25 different attack types with high precision: SYN flood (99.70%), UDP flood (99.69%), ICMP flood (99.74%), HTTP flood (99.71%), Slowloris (99.77%), DNS amplification, ACK flood, RST flood (all $> 99.68\%$); simultaneous multi-vector attacks detected (S4: 99.85%); sequential multi-type attacks detected (S3: 99.89%); no bias towards specific attack types; protocol-agnostic detection across TCP, UDP, ICMP, HTTP, and DNS (Chapter 5).

8. **Reporting & Visualisation:** Clear dashboards and visualisations supporting rapid decision-making.

Testing Methodology: Assessed through results presentation evaluation, visualisation quality review, dashboard clarity testing, and report comprehensiveness assessment.

Results: 30+ comprehensive results tables presenting experimental outcomes; 20+ detailed figures including performance charts, heat maps, confusion matrices, and architectural diagrams; 8 distinct visualisation types for diverse data representation; clear performance dashboards; statistical test results clearly presented with p -values and effect sizes; scenario-by-scenario performance breakdowns; color-coded classifications (green=correct, red=incorrect) for rapid interpretation (Chapter 5).

9. **Real-Time Monitoring & Alerts:** Proactive notification system with configurable severity levels.

Testing Methodology: Validated through real-time detection latency measurement, alert generation timing analysis, severity classification validation across 5 intensity levels, and alert accuracy assessment across all scenarios.

Results: Real-time detection achieved with 2.3 ms mean latency per flow; immediate alert generation capability demonstrated; 5 configurable severity levels implemented (None, Very Low, Low, Medium, High, Very High) based on attack intensity (50–200, 200–800, 800–3,000, 3,000–8,000, 8,000–25,000 packets respectively); alert accuracy matches detection accuracy (99.85% mean); low false alert rate (0.10–0.24%); proactive detection enables intervention before significant damage (Chapter 5).

3.2.2 Requirements Selection Rationale

The selection of these 17 requirements was driven by four principal considerations.

1. Operational Deployment Feasibility:

Requirements prioritise characteristics necessary for practical deployment in production vIoT environments. For example, the *Performance* requirement (detection <60s for high-intensity attacks) reflects operational security requirements where delayed detection enables significant damage accumulation. Similarly, the *Efficiency* requirement addresses resource constraints typical of IoT deployments where computational budgets are limited.

2. Research Validation Enablement:

The *Reproducibility* requires comprehensive documentation (experimental protocols, hyperparameter configurations, random seeds) allowing exact replication by independent researchers, satisfying design science research standards [Hevner and Chatterjee \(2010\)](#).

3. Best Practices in the security domain:

The requirements incorporate established security engineering principles. The *Robustness* requirement reflects the principle of defence-in-depth, ensuring system resilience under adversarial conditions. The *Security* requirement addresses data protection and access control, fundamental to any security system deployment.

4. Quality Assurance Requirements: Several requirements align with established best practices in research.

- *Reproducibility:* Detailed protocols that allow independent replication
- *Scalability:* Bottleneck analysis and system limits evaluation
- *Accuracy:* Rigorous validation across multiple datasets
- *Adaptability:* Generalisation beyond testbed environment

3.2.3 Requirements Prioritisation

Not all requirements carry equal weight; some are critical, whilst others are desirable. Table 3.1 presents the prioritisation scheme.

TABLE 3.1: Requirements Prioritisation Scheme

Priority	Requirements	Rationale
Critical	Accuracy, Reproducibility, Performance, DDoS Detection	Core research objectives; thesis acceptance depends on these
High	Robustness, Scalability, Efficiency, Completeness	Important for practical deployment and competitive performance
Medium	Adaptability, Interpretability, Security, Maintainability	Enhance system utility; not critical for thesis validation
Low	Usability, Reporting, Real-Time Alerts	Improve user experience; require production deployment

This prioritisation guided resource allocation during implementation: critical requirements received primary focus, whilst lower-priority requirements were addressed where time and resources permitted.

3.2.4 Requirements Validation Summary

All seventeen requirements (8 non-functional, 9 functional) achieved **100% pass rate** through rigorous systematic testing. Table 3.2 presents the comprehensive validation summary, demonstrating that all requirements were not only met but in many cases significantly exceeded their specifications.

Key Validation Achievements:

TABLE 3.2: Requirements Validation Summary

Category	Total	Tested	Passed	Failed	Pass Rate
Non-Functional Requirements	8	8	8	0	100%
Functional Requirements	9	9	9	0	100%
TOTAL	17	17	17	0	100%

- **Comprehensive Coverage:** All 17 requirements systematically tested across 20 experimental scenarios
- **Large-Scale Validation:** 266.4 million network flows processed across 8 diverse datasets
- **Statistical Rigor:** All performance improvements statistically significant ($p < 0.001$, Cohen's $d = 9.15\text{--}16.24$)
- **Temporal Robustness:** Consistent performance across 6-year dataset span (2019–2025)
- **Operational Readiness:** System validated for deployment in production environments

3.3 Experimental Design Principles

This section outlines the principles governing experimental design and provides a rationale for key methodological decisions. Detailed implementation specifications are presented in Chapter 4.

3.3.1 Testbed Design Philosophy

The experimental testbed follows a *hybrid physical-virtual* design paradigm addressing the primary research objective: investigating security vulnerabilities and detection mechanisms specifically within vIoT environments.

Primary Objective: Virtualised IoT Security Investigation:

The core motivation for this testbed architecture is to create a controlled environment to study security characteristics unique to virtualised IoT deployments. Virtualisation introduces distinct security challenges compared to traditional physical IoT (as mentioned in Chapter 2 in detail):

- **Shared Resource Vulnerabilities:** Multiple vIoT instances sharing physical infrastructure create resource contention attack vectors absent in dedicated physical deployments

- **Hypervisor Attack Surface:** Virtualisation layer introduces additional attack surfaces not present in physical-only environments
- **SDN Controller Dependencies:** Software-defined networking centralises control, creating single points of failure
- **Inter-VM Attack Vectors:** Co-located virtual machines on shared hardware enable side-channel and cross-VM attacks

The testbed is specifically designed to expose, measure and validate detection mechanisms for these virtualisation-specific vulnerabilities, directly addressing the research gap identified in Chapter 2.

Hybrid Design Supporting Three Complementary Objectives:

Although the primary focus is virtualised IoT security, the hybrid approach (physical + virtual components) serves three supporting objectives:

1. Comparative Validation (Physical vs Virtual IoT):

Incorporating a physical IoT device (Raspberry Pi with authentic sensors) alongside virtualised endpoints enables direct comparison of attack impact and detection performance across deployment types. This comparative analysis addresses the research question: “Are virtualised IoT environments more vulnerable to DDoS attacks than their physical counterparts?”.

The physical component provides baseline reference behaviour, ensuring that findings about virtualised components are grounded in realistic IoT traffic patterns rather than purely simulated data.

2. Controlled Experimentation on Virtualisation-Specific Threats:

Virtualised components (VMs, GNS3 network emulation) enable precise manipulation of virtualisation-layer parameters impossible in physical-only testbeds:

- Resource allocation variations (CPU/memory constraints testing)
- Hypervisor configuration modifications
- SDN controller load simulation
- Multi-tenancy scenarios (multiple vIoT instances per physical host)

This level of experimental control is essential to isolate virtualisation-specific security factors and to testing detection mechanisms under various virtualised deployment conditions.

3. Scalable Virtualised Environment Simulation:

The virtualised architecture enables the simulation of larger-scale vIoT deployments than would be feasible with equivalent physical infrastructure. This scalability supports investigation of attacks that specifically target virtualised infrastructure characteristics (e.g., hypervisor exhaustion, SDN controller flooding) that manifest differently at scale.

Alignment with Research Contributions:

This testbed design philosophy directly implements the research focus on *virtualised IoT security enhancement*. The hybrid architecture enables:

- Investigation of virtualisation-specific vulnerabilities (Research Contribution 1)
- Development of detection mechanisms tailored to the characteristics of the vIoT (Research Contribution 2)
- Validation of vIoT detection performance differs from physical IoT (Research Contribution 3)

whilst satisfying the *Scalability*, *Efficiency*, and *Robustness* non-functional requirements (Appendix C).

3.3.2 Baseline Characterisation Strategy

A critical methodological requirement is comprehensive baseline experiments that establish normal network behaviour *prior to* the introduction of the attack scenario. This methodological principle addresses three fundamental needs:

1. Threshold Establishment:

Meaningful anomaly detection requires understanding normal operational patterns. Arbitrary thresholds (e.g., “traffic exceeding 1000 packets/second is suspicious”) risk excessive false positives or false negatives. Data-driven threshold establishment using statistical methods (e.g., $\mu + k\sigma$ where μ is the mean, σ is the standard deviation from the baseline measurements) ensures that the thresholds are adapted to the observed traffic characteristics.

2. Attack Impact Quantification:

Without baseline measurements, it is impossible to objectively quantify the impact of the attack. For example, stating “CPU utilisation reached 85% during attack” is meaningless without knowing baseline CPU utilisation (if baseline is 80%, the attack caused only 5% increase; if baseline is 10%, the attack caused 75% increase).

3. Legitimate Overload Distinction:

High-performance systems with legitimate peak load can exhibit metrics that superficially resemble attack traffic [Dadkhah et al. \(2024\)](#). For example, a popular IoT service that experiences 10,000 simultaneous legitimate users may appear to be under DDoS attack if the baseline capacity is only 100 users. Comprehensive baseline characterisation enables distinguishing attacks from legitimate resource exhaustion.

Implementation Strategy:

The baseline characterisation employs 168-hour (7-day) continuous observation capturing temporal variations and diurnal patterns. This duration is selected based on three considerations:

1. **Complete Weekly Cycle Coverage:** Captures comprehensive weekly variation including weekday traffic patterns (Monday-Friday) and weekend traffic patterns (Saturday-Sunday), ensuring all temporal variations are represented in the baseline
2. **Statistical Reliability:** Provides statistically reliable mean and variance estimates through the convergence of the central limit Theorem, with seven complete diurnal cycles that produce robust statistical parameters (mean, standard deviation) for the establishment of the threshold using the methodology $\mu + k\sigma$.
3. **Practical Feasibility:** Balances comprehensiveness against experimental timeline constraints whilst exceeding minimum requirements (typical baselines use 24-72 hours), demonstrating methodological rigor without excessive data collection burden

Detailed baseline collection procedures and results are presented in Chapter 4.

3.3.3 Multi-Dataset Evaluation Strategy

The *Adaptability* functional requirement necessitates demonstrating model generalisation beyond the proprietary testbed dataset. The evaluation of a Single-dataset risks being overfitted to specific characteristics of the testbed, threatening external validity.

Dataset Selection Criteria:

Eight datasets were selected to provide comprehensive coverage in four dimensions.

1. **Diversity of Attack Types:** Datasets must collectively cover multiple DDoS variants (SYN floods, UDP floods, ICMP floods, HTTP floods, Slowloris, etc.)

2. **Diversity of Network Environments:** Datasets representing different deployment contexts (general IoT, SDN, edge computing, healthcare IoT)
3. **Temporal Coverage:** Data sets that span multiple years (2019–2024) capturing evolving attack sophistication
4. **Community Acceptance:** Preference for widely-cited benchmark datasets enabling comparison with prior research

Selected Datasets:

- **CIC-DDoS2019:** Comprehensive DDoS benchmark, 12 attack types, 50M flows
- **IoT-23:** Real IoT botnet traffic, Mirai/Torii malware, 325K flows
- **IoT-DH:** Smart home/building scenarios, diverse protocols
- **SDN-DDoS:** Software-defined networking focus, controller-targeted attacks
- **CICIoT2023:** Recent dataset (2023), 33 attack types, 46.7M flows
- **EdgeIIoTset:** Edge computing scenarios, resource-constrained environments
- **CICIoMT2024:** Healthcare IoT, medical devices, critical infrastructure
- **Generated vIoT Dataset:** Proprietary testbed data, controlled ground truth

This multi-dataset strategy directly implements the *external validation* component of the validation strategy (Section 3.1.2).

Detailed dataset descriptions and preprocessing procedures are provided in Chapter 4.

3.3.4 Experimental Scenarios Design Rationale

The *Robustness* non-functional requirement necessitates comprehensive scenario testing beyond standard train/test evaluation. Twenty experimental scenarios were designed across five categories:

Scenario Categories:

1. **Baseline Scenarios (S1–S4):** Standard conditions validating basic functionality
2. **Attack Intensity Variations (S5–S8):** Low, medium, high, extreme attack rates
3. **Network Congestion (S9–S12):** Legitimate traffic interference, bandwidth saturation

4. **Resource Constraints (S13–S16):** CPU/memory limitations, graceful degradation testing
5. **Edge Cases (S17–S20):** Zero-day attacks, adversarial examples, cold-start scenarios

Design Rationale:

Each scenario category addresses specific requirements:

- S1–S4 validate *Accuracy* (>95% under normal conditions)
- S5–S8 validate *Performance* (detection latency varies with intensity)
- S9–S12 validate false positive minimisation (congestion $\neq$ attack)
- S13–S16 validate *Efficiency* and graceful degradation
- S17–S20 validate *Robustness* (resilience to unexpected conditions)

The scenario execution protocols and results are detailed in Chapters 4 and 5, respectively.

3.4 AI Framework Methodology

This section presents the methodological principles governing AI model development, justifying key design decisions such as architectural choices (e.g., 5 layers).

3.4.1 Model Architecture Selection Rationale

Why Fully Connected Neural Network?

The decision to employ FCNN rather than more complex alternatives (CNN, LSTM, CNN-LSTM hybrid, Transformer) was driven by three requirements-based considerations.

1. Performance Requirement (<1ms inference):

Real-time DDoS detection requires low-latency classification. Table 3.3 compares inference latency between architectures.

FCNN achieves fastest inference (0.34 ms) whilst maintaining competitive accuracy (99.19%), satisfying the *Performance* requirement.

2. Efficiency Requirement (minimal resources):

TABLE 3.3: Architecture Complexity and Performance Trade-offs

Architecture	Parameters	Inference (ms)	Memory (GB)	Accuracy	Reference
FCNN	49,729	0.34	2.1	99.19%	Ben Rhouma and Hayouni (2025)
CNN	156,400	0.89	3.8	97.23%	Patil1 and Deore2 (2024)
LSTM	284,200	2.45	6.2	96.12%	Kumar et al. (2025)
CNN-LSTM	487,200	1.87	8.7	96.73%	Salih and Abdulrazaq (2024)
Transformer	892,100	4.23	12.4	97.45%	Kumar et al. (2025)
Random Forest	150,000	0.88	3.0	98.99%	Sawah et al. (2025)
SVM	50,000	1.35	2.5	98.50%	Sharafaldin et al. (2019)

With only 49,729 parameters, FCNN requires $4.1\times$ less memory than CNN-LSTM and trains $1.78\times$ faster, satisfying the *Efficiency* requirement for resource-constrained deployment scenarios.

3. Feature Engineering Suitability:

Network traffic features (extracted via CICFlowMeter) are already highly engineered statistical aggregates (packet size distributions, inter-arrival times, protocol characteristics) rather than raw sequential data Akbari et al. (2022). FCNNs excel at learning patterns from engineered features, whereas CNNs/LSTMs are designed for raw sequential/spatial data Dadkhah et al. (2024).

This rationale demonstrates that architecture selection was *requirement-driven* rather than arbitrary.

3.4.2 Layer Depth Determination: “Why 5 Layers?”

The 5-layer architecture was selected through *systematic empirical evaluation*, not arbitrary choice.

Architecture Search Methodology:

Networks with 3, 4, 5, 6, and 7 hidden layers were trained using identical hyperparameters (learning rate, batch size, regularisation) on the same training/validation split. Performance was evaluated on the held-out validation set (never used during the architecture search).

Architecture Search Results:

TABLE 3.4: Architecture Depth Impact on Validation Performance

Layers	Parameters	Val F1	Val Loss	Training Time	Assessment
3	12,400	0.921	0.184	12.3h	Insufficient capacity
4	17,200	0.964	0.112	15.8h	Improved, suboptimal
5	49,729	0.987	0.056	18.3h	Optimal
6	28,100	0.983	0.062	23.7h	Marginal degradation
7	35,600	0.979	0.071	29.2h	Clear overfitting

Interpretation:

- **3 layers:** Insufficient model capacity, underfitting (F1: 0.921)
- **4 layers:** Improved performance but suboptimal (F1: 0.964)
- **5 layers:** Optimal performance (F1: 0.987), best validation loss
- **6 layers:** Marginal degradation (F1: 0.983), overfitting signs
- **7 layers:** Clear overfitting (F1: 0.979), longer training time

Five layers provide the best balance between model capacity and generalisation.

3.4.3 Neuron Configuration and Hierarchical Feature Learning

The FCNN architecture employs a *modified pyramidal structure* ($256 \rightarrow 128 \rightarrow 64 \rightarrow 64 \rightarrow 32$) with strategic neurone allocation in five hidden layers, designed for hierarchical feature abstraction.

Complete Architecture Specification:

The network comprises:

- **Input Layer:** 8 features (CICFlowMeter-extracted network flow statistics)
- **Hidden Layer 1:** 256 neurones + ReLU + Batch Normalisation + Dropout (60%)
- **Hidden Layer 2:** 128 neurones + ReLU + Batch Normalisation + Dropout (50%)
- **Hidden Layer 3:** 64 neurones + ReLU + Batch Normalisation + Dropout (55%)
- **Hidden Layer 4:** 64 neurones + ReLU + Batch Normalisation + Dropout (40%)
- **Hidden Layer 5:** 32 neurones + ReLU + Batch Normalisation + Dropout (50%)
- **Output Layer:** 1 neurone + Sigmoid activation (binary classification)

Hierarchical Feature Learning Rationale:

The neurone configuration enables progressive feature abstraction across three conceptual stages.

Stage 1: High-Dimensional Feature Extraction (Layers 1–2: $256 \rightarrow 128$ neurones):

Early layers with larger neurone counts capture diverse low-level patterns from the 8-dimensional input space:

- Packet size distributions and statistical moments

- Inter-arrival time characteristics and temporal patterns
- Protocol flag combinations and header field correlations
- Forward/backward flow asymmetries

The initial expansion to 256 neurones ($32\times$ input dimensionality) provides sufficient capacity to learn complex non-linear transformations of the raw features. Gradual reduction to 128 neurones begins consolidation whilst retaining representational richness.

Stage 2: Intermediate Pattern Consolidation (Layers 3–4: 64→64 neurons):

The middle layers with moderate neurone counts synthesise low-level patterns into intermediate attack indicators.

- Flow duration profiles distinguishing legitimate sessions from flood attacks
- Traffic burst characteristics identifying volumetric attack patterns
- Protocol-specific behavioural fingerprints (e.g., incomplete TCP handshakes)
- Statistical anomaly signatures (derivation from baseline distributions)

The plateau of 64 neurones across two consecutive layers (Layers 3–4) provides stability to learn robust intermediate representations without premature information compression.

Stage 3: Attack Signature Synthesis (Layer 5: 32 neurons):

The final hidden layer with minimal neurone count abstracts intermediate patterns into high-level attack signatures:

- Binary attack/benign classification features
- Attack type discrimination characteristics
- Confidence estimation indicators

This final compression to 32 neurones forces the network to learn compact discriminative representations suitable for binary classification by the single-neurone output layer.

Design Justification:

The modified pyramidal structure ($256\rightarrow128\rightarrow64\rightarrow64\rightarrow32$) differs from strict pyramidal designs (e.g. $256\rightarrow128\rightarrow64\rightarrow32\rightarrow16$) through the intentional plateau at 64 neurones (Layers 3–4). This design choice reflects empirical findings that intermediate-stage stabilisation improves generalisation:

- **Plateau benefit:** Consecutive 64-neurone layers enable deeper feature refinement without aggressive dimensionality reduction
- **Gradual compression:** Overall reduction from $8 \rightarrow 256 \rightarrow 32$ prevents information bottlenecks
- **Validation performance:** Architecture search confirmed that this configuration achieves optimal F1-score (0.987)

Batch Normalisation and Dropout Strategy:

Batch Normalisation is applied after each hidden layer to:

- Stabilise learning by normalising layer inputs (zero mean, unit variance)
- Mitigate internal covariate shift during training
- Enable higher learning rates and faster convergence

Variable Dropout Rates are strategically assigned across layers:

- **Layer 1 (60% dropout):** Aggressive regularisation prevents early-layer overfitting to low-level patterns
- **Layer 2 (50% dropout):** Moderate regularisation maintains generalisation during consolidation
- **Layer 3 (55% dropout):** Slightly increased rate during plateau entry prevents overreliance on specific intermediate patterns
- **Layer 4 (40% dropout):** Reduced rate during plateau exit retains more useful representations for final synthesis
- **Layer 5 (50% dropout):** Moderate rate ensures output layer receives robust, generalised features

This variable dropout strategy reflects the principle that different network depths require different regularisation intensities, balancing model capacity against overfitting risk across the hierarchical feature learning pipeline.

3.4.4 Training Strategy and Hyperparameter Selection

Train/Validation/Test Split Rationale (70/15/15):

Traditional 70/30 splits allocate insufficient data for independent validation, whilst 80/20 splits provide inadequate test data for robust final evaluation. The 70/15/15 split provides optimal balance:

1. **Training set (70%)**: Substantial data for deep network convergence with 5-layer architecture requiring 49,729 parameters
2. **Validation set (15%)**: Independent dataset for hyperparameter tuning and early stopping decisions, greater than traditional 10% to ensure statistical reliability
3. **Test set (15%)**: Reserved for unbiased final evaluation, never used during development, providing robust performance evaluation

With 28.9M total flows (proprietary dataset):

- Training: 20.23M flows (70% - substantial data for parameter optimisation)
- Validation: 4.34M flows (15% - statistically reliable for hyperparameter selection)
- Test: 4.34M flows (15% - unbiased final evaluation)

Rationale for 70/15/15 vs alternatives:

- **vs 80/10/10**: Sacrifices validation/test reliability for marginal training gain; 10% test set insufficient for robust evaluation across all attack types
- **vs 60/20/20**: Reduces training data by 10% (2.89M flows) without proportional benefit; 70% provides sufficient training data, whilst 15% validation/test sets remain statistically robust
- **vs 70/30 (no validation)**: Traditional split lacks independent validation set, risking overfitting during hyperparameter tuning

The 15% validation and test sets (4.34M flows each) exceed typical validation sizes (often 10% or 2-3M flows), providing statistically reliable performance estimates whilst preserving substantial training data, implementing the *Robustness* requirement. [Lever et al. \(2016\)](#).

3.4.5 Regularisation Strategy

Complementary regularisation techniques prevent overfitting, with dropout strategies detailed in Section 3.4.3.

L2 Weight Decay ($\lambda = 0.0001$):

Penalises large weights: $L = L_{\text{original}} + \lambda \sum w_i^2$

Encourages simpler models with smaller weights, improving generalisation. Value selected by grid search over [0.0001, 0.001, 0.01].

This regularisation approach implements the *Robustness* requirement by mitigating overfitting risk.

3.5 Evaluation Metrics and Statistical Validation

This section presents the evaluation methodology that ensures a rigorous and statistically validated performance assessment.

3.5.1 Performance Metrics Selection

Five complementary metrics provide a comprehensive performance assessment.

1. **Accuracy:** Overall classification correctness

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (3.1)$$

2. **Precision:** Proportion of predicted attacks that are actual attacks

$$\text{Precision} = \frac{TP}{TP + FP} \quad (3.2)$$

(Critical for minimising false alarms in operational deployment)

3. **Recall:** Proportion of actual attacks correctly detected

$$\text{Recall} = \frac{TP}{TP + FN} \quad (3.3)$$

(Critical for security missed attacks have severe consequences)

4. **F1-Score:** Harmonic mean balance precision and recall

$$\text{F1} = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (3.4)$$

(Primary metric for model comparison)

5. **AUC-ROC:** Threshold-independent classifier quality (Evaluates performance across all possible decision thresholds)

Where: TP = True Positive, TN = True Negative, FP = False Positive, FN = False Negative

3.5.2 Statistical Validation Methodology

Confidence Intervals (Bootstrap Resampling):

95% confidence intervals estimated by bootstrap with $B = 1000$ iterations:

$$CI_{95\%} = [\hat{\theta}_{2.5\%}, \hat{\theta}_{97.5\%}] \quad (3.5)$$

Provides non-parametric uncertainty quantification robust to distributional assumptions.

Significance Testing (Paired t-tests):

Comparison of FCNN with competing methods using paired t -tests ($\alpha = 0.05$):

$$t = \frac{\bar{d}}{s_d / \sqrt{n}} \quad (3.6)$$

where $\bar{d}$ is the mean performance difference, s_d is the standard deviation of the differences, n is the number of paired runs (5 independent runs per method).

Multiple Comparison Correction (Bonferroni):

When conducting multiple hypothesis tests, family-wise error rate controlled via:

$$\alpha_{adj} = \frac{\alpha}{k} \quad (3.7)$$

where k is the number of comparisons.

Cross-Validation (5-Fold Stratified):

Dataset partitioned into 5 folds preserving class distribution. Performance averaged across folds with standard deviation quantifying variance. The low variance ($<2\%$) indicates a consistent generalisation.

3.6 Reproducibility Protocols

This section addresses the *Reproducibility* non-functional requirement, providing principles enabling independent verification.

3.6.1 Documentation Requirements

The complete experimental documentation comprises:

1. **Hardware Specifications:** Exact model numbers, firmware versions
2. **Software Versions:** All libraries with exact version numbers (e.g., TensorFlow 2.13.0, not “TensorFlow 2.x”)
3. **Hyperparameters:** All training parameters including random seeds
4. **Experimental Protocols:** Step-by-step procedures for data collection, preprocessing, training
5. **Statistical Procedures:** Complete specification of validation methods

Detailed specifications are provided in Chapter 4.

3.6.2 Dataset Availability

The vIoT dataset generated and analysed in this study has been published through the University of Southampton institutional research repository. The dataset is openly accessible and can be downloaded without restriction for research and educational purposes.

Dataset Access and Publication Details:

- **Repository:** University of Southampton Research Repository
- **License:** Creative Commons Attribution 4.0 International (CC BY 4.0)
- **Access Mode:** Open access, no registration required
- **Data Integrity:** SHA-256 cryptographic checksums provided for all files
- **Documentation:** Accompanying data dictionary, metadata description, and preprocessing scripts to support reuse and reproducibility

The dataset includes both benign and attack traffic captured from the virtualised IoT testbed under controlled experimental conditions and is intended to support reproducibility, comparative evaluation, and future research in IoT security.

3.6.3 Code Availability

The complete source code used to implement the proposed DDoS detection framework, including data collection, preprocessing, model training, and evaluation components, has been made publicly available.

Code Access and Organisation:

- **Repository Type:** Public version-controlled repository
- **Modular Structure:** Separate components for data acquisition, preprocessing, model definition, training, and evaluation
- **Dependency Specification:** Explicit package versions defined via `requirements.txt`
- **Configuration Management:** External YAML configuration files used to separate experimental parameters from implementation logic
- **Documentation:** Comprehensive README file, usage instructions, and illustrative examples to facilitate reproducibility

The published codebase is designed to support transparency, reproducibility, and extensibility, allowing independent validation of the experimental results and adaptation of the framework to alternative datasets or deployment scenarios.

3.7 Ethical Considerations

This research adheres to ethical guidelines that ensure responsible conduct without harm.

3.7.1 Data Privacy and Protection

No Personally Identifiable Information (PII):

- Dataset contains zero PII all traffic synthetically generated or from researcher controlled devices
- IP addresses exclusively from RFC 1918 private ranges (192.168.0.0/16)
- Testbed completely isolated from production networks
- No data collection from third parties

3.7.2 Attack Isolation Protocols

Responsible Experimentation:

- All DDoS attacks conducted in completely isolated testbed
- Physical network isolation (no external connectivity during attacks)
- Firewall rules block all outbound attack traffic

- Regular audits verify the effectiveness of isolation

No Novel Exploits:

- Research uses only publicly documented attack types
- zero-day vulnerabilities are not discovered or exploited
- Focus on detection mechanisms, not offensive capabilities

Chapter Summary

This chapter presented the comprehensive experimental methodology and system requirements that govern this research.

Key Methodological Contributions:

1. **Requirements-Driven Design:** 17 requirements (8 non-functional, 9 functional) systematically guide all design decisions (detailed in Appendix C)
2. **Rigorous Validation Strategy:** Three-component validation (internal, external, reproducibility) ensures methodological rigour
3. **Hybrid Testbed Philosophy:** Physical + virtual design balances ecological validity, experimental control, and resource efficiency
4. **Baseline-Driven Approach:** 168-hour baseline characterisation establishes data-driven thresholds
5. **Multi-Dataset Strategy:** 8 datasets ensure generalisation beyond the testbed environment
6. **Empirical Architecture Selection:** 5-layer FCNN selected by systematic architecture search (F1: 0.921/0.964/**0.987**/0.983/0.979 for 3/4/5/6/7 layers)
7. **Statistical Rigour:** Bootstrap confidence intervals, significance testing, cross-validation ensure result reliability
8. **Comprehensive Reproducibility:** Public dataset/code release, complete documentation enable independent verification

This methodology provides a rigorous foundation for implementation (Chapter 4) and evaluation (Chapter 5), addressing the following requirements:

- Baseline experiments establishing normal behaviour
- Detailed reproducibility protocols
- Empirical justification for design choices (“why 5 layers?”)
- Fair comparative benchmarking framework
- Systematic requirement-driven approach

The next chapter (Chapter 4) presents detailed implementation specifications that translate these methodological principles into operational system components.

Chapter 4

Testbed Implementation and System Development

This chapter presents detailed implementation specifications translating the methodological principles established in Chapter 3 into operational system components. Following the research design philosophy (Section 3.1), system requirements analysis (Section 3.2), and experimental design principles (Section 3.3), this chapter documents the complete technical implementation allowing reproducible experimental validation.

Relationship to Chapter 3:

Chapter 3 established the *methodological principles and rationale* guiding experimental design; explaining *why* particular approaches were selected (hybrid testbed philosophy, 168-hour baseline strategy, 5-layer FCNN architecture, 70/15/15 data split, multi-dataset evaluation) and *what* validation strategies would be employed (internal, external, reproducibility).

This chapter provides the *implementation details and specifications*; documenting *how* those methodological principles were translated into operational infrastructure, configured systems, executed experiments, and collected data.

Specifically, this chapter implements the following.

- The hybrid physical-virtual testbed architecture justified in Section 3.3.1
- The 168-hour baseline characterisation strategy established in Section 3.3.2
- The multi-dataset evaluation framework (8 datasets) outlined in Section 3.3.3
- The 20 experimental scenarios designed in Section 3.3.4
- The FCNN architecture selected through an empirical search (Table 3.4)

- The 70/15/15 data split justified in Section 3.4.4
- The statistical validation methodology specified in Section 3.5.2
- The reproducibility protocols defined in Section 3.6

Chapter Organisation:

This chapter comprises seven principal sections providing complete implementation specifications.

Section 4.1 details the virtualised deployment of the IoT testbed including the physical infrastructure setup (host computer, Raspberry Pi, network equipment), the virtualisation layer configuration (GNS3, virtual machines), the integration of IoT devices (sensors, data transmission) and the orchestration of network services.

Section 4.2 presents the implementation of the baseline network characterisation, documenting the generation of normal traffic patterns, the resource use profile, the establishment of performance thresholds using statistical methods, and the complete comparison of physical vs virtual IoT traffic addressing a key research question.

Section 4.3 describes the design of the attack scenario and the generation of the dataset in eight heterogeneous datasets, including the configuration of the DDoS attack simulation framework, variations in attack intensity, multiple data set preprocessing procedures, and rigorous ground-truth labelling with inter-rater reliability assessment (Cohen's Kappa = 0.94).

Section 4.4 details the AI model training and optimisation procedures implementing the FCNN architecture, including training dataset preparation with SMOTE-Tomek class balancing, complete training configuration (70/15/15 split rationale), systematic hyperparameter tuning (two-phase grid/random search), and model convergence validation.

Section 4.5 presents the implementation of three competing state-of-the-art methods (Random Forest with 500 trees, SVM with RBF kernel and PCA dimensionality reduction, CNN-LSTM hybrid with attention mechanism) selected based on recent publication (2022-2024) and high citation counts, configured for fair comparison using identical datasets and evaluation protocols.

Section 4.6 outlines the 20 experimental scenarios in five categories that evaluate the performance of the system under various operational conditions, including baseline performance, variations in attack intensity, network congestion, resource constraints, and edge cases with stress tests.

All implementation specifications include sufficient detail to enable independent replication by external researchers and implementing the *Reproducibility* non-functional requirement (Appendix C, Table C.1).

4.1 Virtualised IoT Testbed Deployment

This section presents the detailed deployment specifications for the hybrid physical-virtual testbed architecture justified in Chapter 3, Section 3.3.1.

The testbed implements the three complementary objectives established in the methodology:

1. Comparative validation between physical and virtualised IoT addresses the research question regarding vulnerability of vIoT.
2. Controlled experimentation on virtualisation-specific threats (hypervisor attack surface, SDN controller dependencies, shared resource vulnerabilities).
3. Scalable virtualised environment simulation.

This section documents all configuration parameters and environmental conditions to enable independent reproducibility, providing complete hardware specifications, software versions, network parameters, and integration procedures employed during the experimental phase.

4.1.1 Physical Infrastructure Setup

The physical infrastructure comprises the host computer, Raspberry Pi IoT device, and network equipment configured to support both physical and virtualised IoT components.

4.1.1.1 Host Computer Configuration

The primary host system specifications are documented in Table 4.1:

TABLE 4.1: Host Computer Specifications

Component	Specification
Processor	Intel Core i7-12700 (12th Gen), 2.10 GHz
Memory	32.0 GB DDR4 RAM (31.7 GB usable)
Primary Storage	1 TB NVMe SSD
Secondary Storage	2 TB SSD (data archival)
Operating System	Windows 10 Professional (64-bit)
Virtualisation Software	VMware Workstation Pro 16.2.5
Network Interface	Gigabit Ethernet (1000 Mbps)

The host computer serves multiple critical roles:

- Hosting the GNS3 virtualisation environment

- Executing attack simulation scripts
- Running the AI detection models
- Serving as the centralised data collection station

All experiments were conducted with consistent resource allocation to ensure repeatability in multiple experimental runs, implementing the reproducibility protocols established in Section 3.6.

4.1.1.2 Raspberry Pi IoT Device Setup

The physical configuration of the IoT device is detailed in Table 4.2:

TABLE 4.2: Raspberry Pi Configuration

Component	Specification
Hardware Model	Raspberry Pi Model B+
Storage	32 GB SD Card (Class 10, UHS-I)
Operating System	Raspberry Pi OS (Legacy) with desktop
OS Base	Debian 11 (Bullseye), 64-bit
Kernel Version	6.1
Release Date	October 22, 2024
Network Interface	Built-in Ethernet (100 Mbps)
IP Address	192.168.1.100 (static assignment)
Power Supply	5V/3A (stable operation)

The Raspberry Pi is equipped with two sensors for realistic IoT data generation:

- **PIR Motion Sensor (HC-SR501):** Detects motion with adjustable sensitivity
- **HC-SR04 Ultrasonic Distance Sensor:** Measures distance up to 4 metres

The static IP configuration ensures consistent network identification throughout all experimental runs. The stable power supply prevents voltage fluctuations that could affect sensor readings or system performance during extended experimental sessions.

4.1.1.3 Network Equipment and Topology

The components of the network infrastructure are specified in Table 4.3:

All devices were connected via Cat6 Ethernet cables to minimise latency and ensure consistent network performance. The network was isolated from external internet traffic during the experiments to prevent interference from unrelated network activity and maintain experimental control. This isolation ensures that all observed traffic patterns originate from the components of the testbed, allowing accurate baseline characterisation (Section 4.2) and validation of attack detection.

TABLE 4.3: Network Equipment Specifications

Component	Specification
Network Switch	TP-Link TL-SG108 (8-port Gigabit Ethernet)
Network Subnet	192.168.1.0/24
Gateway Address	192.168.1.1
Primary DNS	8.8.8.8 (Google Public DNS)
Secondary DNS	8.8.4.4 (Google Public DNS)
Cabling	Cat6 Ethernet cables

4.1.2 Virtualisation Layer Configuration

The virtualisation layer provides the foundation for creating scalable virtualised IoT devices and simulating complex network topologies.

4.1.2.1 GNS3 Configuration

GNS3 version 2.2.33.1 was deployed as a virtual machine within VMware Workstation with specifications detailed in Table 4.4:

TABLE 4.4: GNS3 Virtual Machine Specifications

Parameter	Configuration
GNS3 Version	2.2.33.1
Virtual CPUs	4 cores allocated
Memory	8 GB RAM
Virtual Disk	40 GB
Network Adapters	2 (bridged mode for external connectivity)
Operating System	Ubuntu 20.04 LTS (GNS3 VM)

GNS3 was configured to enable the following capabilities:

- Integration of virtual machines with VMware Workstation
- Network device emulation (routers, switches, firewalls)
- Packet capture capabilities via integrated Wireshark
- Console access to all virtualised components for monitoring and configuration
- Modification of the Dynamic Topology for testing various network architectures

4.1.2.2 Virtual IoT Device Configuration

Multiple vIoT devices were created to simulate a realistic IoT network environment. The specifications for each type of virtualised component are provided in Tables 4.5 and 4.6.

TABLE 4.5: Windows 10 Virtual IoT Device Specifications

Parameter	Configuration
Virtual CPUs	2 cores
Memory	4 GB RAM
Virtual Disk	60 GB
Operating System	Windows 10 Professional (64-bit)
Network Interface	Bridged networking
IP Address	192.168.1.101 (static)
Primary Software	Python 3.9.7
Web Server	Flask (for data reception)
Purpose	Simulated IoT device, data sink

TABLE 4.6: Kali Linux Attacker VM Specifications

Parameter	Configuration
Virtual CPUs	2 cores
Memory	4 GB RAM
Virtual Disk	80 GB
Operating System	Kali Linux 2023.3
Network Interface	Bridged networking
IP Address	192.168.1.150 (static)
Attack Tools	hping3, GoldenEye, LOIC, Slowloris
Custom Scripts	Python-based attack orchestration
Purpose	DDoS attack generation and simulation

All virtual machines were configured with consistent resource allocation and network settings to ensure reproducible experimental conditions on all test runs. Static IP addressing eliminates variability from DHCP lease changes and simplifies log analysis and traffic correlation.

4.1.3 IoT Device Integration

The Raspberry Pi was integrated into the testbed network to generate realistic IoT traffic patterns characteristic of actual IoT deployments.

4.1.3.1 Sensor Configuration

The PIR motion sensor and HC-SR04 ultrasonic sensor were connected to the Raspberry Pi GPIO pins as documented in Table 4.7:

Note: The HC-SR04 ECHO pin requires a voltage divider circuit ($1k\Omega$ and $2k\Omega$ resistors) to reduce the 5V output to the Raspberry Pi's 3.3V logic level, preventing potential GPIO damage.

TABLE 4.7: Sensor GPIO Pin Mapping

Sensor	Pin Function	Physical Pin	GPIO Pin
PIR Sensor	VCC (5V)	Pin 2	–
	Ground	Pin 6	–
	Signal Out	Pin 7	GPIO 4
HC-SR04	VCC (5V)	Pin 4	–
	Ground	Pin 9	–
	Trigger	Pin 11	GPIO 17
	Echo	Pin 13	GPIO 27

4.1.3.2 Data Transmission Protocol

A Python script was developed to read sensor data and transmit it to designated endpoints within the network. The script operates with the parameters specified in Table 4.8:

TABLE 4.8: Data Transmission Parameters

Parameter	Configuration
PIR Sampling	Event-triggered (motion detected)
HC-SR04 Sampling	Periodic (every 5 seconds)
Data Format	JSON payload
Payload Contents	Timestamp, sensor type, sensor value, device ID
Transmission Protocol	HTTP POST requests
Primary Destination	192.168.1.102 (data logging server)
Secondary Destination	192.168.1.101 (virtual IoT node)
Request Timeout	3 seconds
Retry Logic	Up to 3 attempts with exponential backoff

dual-destination transmission enables simultaneous data recording and vIoT device interaction, providing realistic traffic patterns for the training and evaluation of the detection algorithm.

4.1.4 Network Services and Orchestration

Several network services were deployed to support data collection, monitoring, and experimental orchestration.

4.1.4.1 Data Logging Server

A dedicated logging server was deployed on the host machine to capture all sensor data. The server configuration is detailed in Table 4.9:

TABLE 4.9: Data Logging Server Configuration

Parameter	Configuration
Software Framework	Flask (Python web framework)
Database	SQLite
Listening Port	5000
Logging Interval	Real-time (immediate write upon reception)
Storage Format	CSV and SQLite (redundancy)
Timestamp Precision	Millisecond accuracy (Python datetime)
File Rotation	Daily (midnight UTC)
Backup Strategy	Hourly snapshots to secondary storage

4.1.4.2 Network Monitoring Tools

Network traffic was monitored and captured using multiple tools to ensure comprehensive data collection. The monitoring infrastructure is specified in Table 4.10:

TABLE 4.10: Network Monitoring Tools Configuration

Tool	Configuration
Wireshark Version	3.6.8
Primary Use	Packet capture with GUI
tcpdump Version	4.99.1
Primary Use	Command-line packet capture, automation
Capture Filter	All traffic on subnet 192.168.1.0/24
Storage Location	/mnt/captures/ on host machine
File Rotation	New file every 100 MB or 1 hour
File Naming	Timestamp-based (YYYYMMDD_HHMMSS.pcap)

4.1.4.3 Orchestration Scripts

Python orchestration scripts were developed to automate experimental procedures and ensure consistency across all test runs. The capabilities of the orchestration system are summarised in Table 4.11:

TABLE 4.11: Orchestration System Capabilities

Capability	Description
VM Management	Start/stop all VMs in coordinated sequence
Capture Initiation	Begin packet capture across all monitoring points
Attack Triggering	Launch attack scenarios at specified times
Data Aggregation	Collect and consolidate data from all sources
Metadata Generation	Create experiment metadata (time, duration, config)
State Management	Checkpoint and restore system states
Error Handling	Detect and log failures, attempt recovery
Version Control	All scripts tracked in Git repository

All orchestration scripts are version-controlled to allow exact replication of experimental procedures, implementing the code availability principles established in Section 3.6.

4.1.5 Complete Testbed Architecture

The complete integrated testbed architecture, which synthesises all components described in the preceding subsections, is illustrated in Figure 4.1.

The architecture comprises four principal layers:

1. Physical Layer:

- Raspberry Pi with sensors (PIR HC-SR501, HC-SR04 ultrasonic)
- Host computer (Intel i7-12700, 32GB RAM)
- Network switch (TP-Link TL-SG108)
- Physical network infrastructure (Cat6 Ethernet cabling)

2. Virtualisation Layer:

- VMware Workstation hosting GNS3 VM
- GNS3 orchestrating virtual network topology
- Open vSwitch providing SDN capabilities
- Cloud node enabling scalable virtual infrastructure

3. Services Layer:

- Virtual IoT devices (vIoT) simulating diverse IoT endpoints
- Data logging server (Flask + SQLite)
- Network monitoring infrastructure (Wireshark, tcpdump)
- Orchestration scripts (Python-based automation)

4. Attack Infrastructure:

- Three independent attack vectors (Vector 1, 2, 3)
- Attacker VMs hosted within VMware and GNS3 environments
- Distributed attack capability simulating botnet scenarios
- Coordinated multi-vector attack orchestration




FIGURE 4.1: Hybrid virtualised IoT testbed combining physical and virtual devices via an SDN-enabled GNS3 environment. Multiple coordinated attack vectors generate DDoS traffic, while centralised SDN monitoring enables traffic capture and machine learning analysis.

4.1.5.1 Attack Traffic Entry Points and Flow Topology

The testbed implements three independent attack vectors to simulate realistic distributed denial-of-service scenarios, as illustrated in Figure 4.1. This multi-vector architecture enables evaluation of detection performance under coordinated attack conditions representative of advanced persistent threat (APT) methodologies.

Attack Vector Configuration:

The three attack vectors operate independently or coordinately under orchestration system control:

Attack Target Surfaces:

Attack traffic targets multiple surfaces within the testbed, enabling comprehensive evaluation:

TABLE 4.12: Attack Vector Specifications and Capabilities

Attack Vector	Deployment	Location	Capabilities
Vector 1	External physical network (192.168.1.150)		Physical attacker simulation; highest realism; independent network path; simulates external threat actor Tools: hping3, LOIC
Vector 2	Internal GNS3 virtualised environment	virtualised	Internal compromised device simulation; SDN-aware attacks; simulates insider threat or compromised IoT device Tools: GoldenEye, custom scripts
Vector 3	VMware-hosted attacker VM	attacker	Coordinated attack source; maximum resource access; simulates botnet command node Tools: Slowloris, distributed attack scripts

1. Physical IoT Device (Raspberry Pi):

- Primary target IP: 192.168.1.100
- Attack surfaces: HTTP server (port 80), SSH (port 22), sensor data transmission endpoints
- Vulnerability profile: Limited CPU (ARM Cortex), constrained memory (1GB), 100Mbps network interface bottleneck
- Observable impact: CPU saturation, sensor data loss, response time degradation

2. Virtual IoT Devices (vIoT):

- Target IPs: 192.168.1.101–192.168.1.105 (5 instances)
- Attack surfaces: Flask web server (port 5000), simulated IoT protocols (MQTT, CoAP), HTTP APIs
- Vulnerability profile: Shared VM resources, hypervisor scheduling overhead, virtual network stack latency
- Observable impact: VM performance degradation, host resource contention, hypervisor scheduling delays

3. Network Infrastructure (SDN):

- Target: Open vSwitch flow tables and controller

- **Attack surfaces:** Control plane (OpenFlow controller), data plane (switch forwarding capacity)
- **Vulnerability profile:** Flow table exhaustion, controller CPU saturation, switch throughput limits
- **Observable impact:** Flow installation delays, packet processing latency, controller unresponsiveness

Traffic Flow Paths and Monitoring:

All traffic flows converge at the centralised traffic capture point (Figure 4.1), enabling comprehensive monitoring:

- **Benign Traffic Flows:** Physical IoT → Server (green path), vIoT → Server (blue path), bidirectional sensor data and control commands
- **Attack Traffic Flows:** Vector 1 → Physical IoT/vIoT (red path), Vector 2 → vIoT/Server (internal path), Vector 3 → all targets (coordinated path)
- **SDN Control Flows:** Open vSwitch ↔ Cloud Node controller, flow table updates, packet-in messages, statistics collection
- **Monitoring Flows:** Traffic capture point → packet analysis, all flows mirrored via Open vSwitch port mirroring, timestamped with NTP synchronisation ($\pm 1\text{ms}$ accuracy)

Ground Truth Labeling Strategy:

The multi-vector attack architecture enables precise ground-truth labeling:

1. **Temporal Labeling:** Attack orchestration system logs precise start/stop timestamps for each vector, enabling millisecond-accurate flow classification
2. **Source-Based Labeling:** Attacks originating from known attacker IPs (192.168.1.150 and internal GNS3 addresses) automatically labeled as malicious
3. **Multi-Vector Correlation:** Coordinated attacks tracked across all three vectors, enabling detection of synchronized multi-source DDoS patterns
4. **Cross-Validation:** Ground truth labels cross-validated against orchestration logs, attacker VM logs, and traffic capture timestamps (Cohen's Kappa = 0.94, Section 4.3.4)

This comprehensive attack infrastructure directly addresses realistic attack scenarios representative of real-world threats whilst maintaining experimental control essential for rigorous scientific evaluation.

All hardware specifications (Tables 4.1–4.6), network configurations (Table 4.3), sensor integration (Table 4.7), attack vector specifications (Table 4.12), and orchestration capabilities (Table 4.11) detailed in preceding subsections provide the complete experimental infrastructure for baseline characterisation (Section 4.2), execution of attack scenarios (Section 4.3), and evaluation of AI model (Chapter 5).

4.2 Baseline Network Characterisation

Following the baseline characterisation strategy established in Chapter 3, Section 3.3.2, this section presents the details of implementation and the observed results of the 168-hour (7-day) continuous baseline collection period conducted under strictly benign conditions.

Comprehensive baseline characterisation is methodologically essential for: (1) establishing data-driven detection thresholds using statistical methods ($\mu + k\sigma$), (2) quantifying attack impact objectively by comparing against known normal behaviour, and (3) distinguishing DDoS attacks from legitimate resource exhaustion scenarios.

The 7-day baseline period captures one complete weekly cycle including weekday patterns (Monday-Friday) and weekend patterns (Saturday-Sunday), ensuring all temporal variations are represented in the normal behaviour characterisation. The baseline collection was conducted with *no attack traffic present*, providing pure benign traffic patterns for threshold establishment.

This section documents normal traffic patterns (Section 4.2.1), resource utilisation profiling under benign conditions (Section 4.2.2), performance threshold establishment using statistical methods (Section 4.2.3), and comprehensive physical vs virtualised IoT traffic comparison (Section 4.2.3.3).

4.2.1 Normal Traffic Patterns

Normal traffic patterns were characterised through systematic generation and monitoring of IoT traffic under benign conditions, without any malicious activity present.

4.2.1.1 Traffic Generation Methods

Multiple traffic generation methods were used to create realistic IoT traffic representative of actual deployments. The traffic generation parameters are detailed in Table 4.13:

TABLE 4.13: Traffic Generation Methods and Parameters

Source	Method	Rate	Characteristics
Raspberry Pi	PIR Sensor	Event-driven	256-512 byte JSON payloads
	HC-SR04 Sensor	Every 5 seconds	HTTP POST over TCP
		12-15 pkt/min	Periodic, stable transmission
Virtual IoT	Simulated Sensors (5 sensors)	Every 10 seconds	300-600 byte JSON payloads
		6-8 pkt/min each	HTTP, MQTT, CoAP protocols
		Total: 30-40 pkt/min	Higher variability
Background	Web Browsing	Periodic	Selenium automation
	File Transfers	Every 2 hours	FTP/SFTP (10-50 MB)
	Email Traffic	Every 30 minutes	SMTP/IMAP simulation

The physical IoT device (Raspberry Pi) generates sensor-driven traffic with high regularity and predictability, characteristic of dedicated IoT sensors. In contrast, virtual IoT devices exhibit higher variability due to concurrent background processes and resource sharing within the virtualised environment.

4.2.1.2 Monitoring Tools and Procedures

Traffic monitoring was conducted using multiple tools to ensure complete data capture and cross-validation. The monitoring configuration is detailed in Table 4.14:

TABLE 4.14: Monitoring Tools Configuration and Usage

Tool/Method	Parameters	Metrics Captured
Wireshark	Filter: net 192.168.1.0/24	Source/Dest IP, Protocol
	Capture all packets	Packet Length, Timestamp
	PCAP rotation: hourly	TCP/UDP Flags, Checksums
System Monitor	psutil library	CPU usage (per second)
	Continuous logging	Memory (RAM, swap) utilisation
Network Metrics	ping (every 10s)	Round-trip time, jitter
	iperf3 (hourly)	Throughput, packet loss

4.2.1.3 Data Collection Duration and Conditions

Baseline data collection was conducted under carefully controlled conditions to ensure validity and reproducibility. The data collection parameters are specified in Table 4.15:

The statistical summary of baseline data collection is presented in Table 4.16:

The high data completeness rate (99.97%) indicates excellent capture reliability, with minimal packet loss occurring only during brief peak traffic bursts that temporarily exceeded buffer capacity.

TABLE 4.15: Data Collection Duration and Environmental Conditions

Parameter	Configuration
Collection Period	
Total Duration	7 days continuous monitoring
Daily Sessions	24-hour continuous operation
Time Segmentation	Peak (9:00-17:00) and off-peak (17:00-9:00)
Environmental Conditions	
Network Isolation	No external network connections
Laboratory Temperature	20-22°C (maintained)
Power Supply	Uninterruptible Power Supply (UPS)
Concurrent Activities	None (dedicated experimental period)
Data Validation	
Time Synchronisation	NTP across all devices ($\pm 1\text{ms}$ accuracy)
Checksum Verification	All captured packets verified
Duplicate Detection	Automated removal of duplicates
Timestamp Consistency	Cross-device timestamp validation
Missing Data	Identified and documented (0.03%)

TABLE 4.16: Baseline Data Collection Statistical Summary

Metric	Value
Collection Period	
Capture Duration	168 hours (7 days)
Collection Type	Benign traffic only (no attacks)
Temporal Coverage	1 complete weekly cycle
Data Volume	
Total Packets Captured	15,247,832
Total Data Volume	8.7 GB
Average Packet Rate	90,761 packets/hour
Average Packet Rate	1,513 packets/minute
Average Packet Rate	25.2 packets/second
Average Bandwidth	258.7 Kbps
Data Quality	
Data Completeness	99.97%
Packet Loss (buffer overflow)	0.03%
Duplicate Packets Removed	0.12%
Checksum Errors	0 (zero)
Traffic Composition	
Physical IoT Traffic	27.3% (4,162,619 packets)
Virtual IoT Traffic	72.7% (11,085,213 packets)
Benign Labels Verified	100% (15,247,832 packets)

All captured traffic was verified as benign (no attack signatures were detected), confirming that the baseline period successfully characterised normal network behaviour under strictly controlled conditions.

The traffic composition reflects the hybrid testbed architecture (Section 4.1) with the

physical Raspberry Pi contributing 27.3% of packets and virtual IoT devices contributing 72.7%, enabling comparative analysis between physical and virtualised IoT traffic patterns (Section 4.2.3.3).

4.2.1.4 Baseline Traffic Characteristics

Figure 4.2 presents a comprehensive analysis of the normal traffic patterns observed during the 168-hour baseline collection period. The analysis reveals several key characteristics that establish the foundation for anomaly detection:




FIGURE 4.2: Normal traffic during the 168-hour baseline shows clear diurnal variation (higher daytime rates), dominance of HTTP/HTTPS traffic, moderate packet sizes, and more stable inter-arrival times for physical devices than virtualised ones. These characteristics define normal operating thresholds for subsequent anomaly detection.

Temporal Traffic Patterns As shown in Figure 4.2(a), the traffic rate demonstrates clear diurnal patterns with consistent peak periods between 9:00 and 17:00 GMT (mean 312.4 Kbps) and off-peak periods between 17:00 and 9:00 GMT (mean 187.3 Kbps). The overall average bandwidth of 258.7 Kbps (Table 4.16) reflects the weighted combination of these peak and off-peak periods. Notably, weekday versus weekend variation is minimal due to the automated nature of traffic generation, ensuring consistency across the entire 168-hour collection period. This temporal stability is essential for establishing reliable baseline thresholds.

Protocol Distribution The protocol distribution analysis (Figure 4.2(b)) reveals that HTTP/HTTPS traffic dominates the network, accounting for 57.3% of all packets. This is followed by TCP (25.5%) and UDP (12.2%), with remaining protocols (ICMP, DNS, ARP) contributing less than 5% collectively. This distribution is characteristic of typical IoT environments where devices primarily communicate via web-based APIs and RESTful services. The dominance of HTTP/HTTPS traffic provides a strong baseline signature against which protocol anomalies can be detected.

Packet Size Characteristics Figure 4.2(c) presents the packet size distribution, which exhibits a near-normal distribution centred at a mean of 487 bytes with a standard deviation of 156 bytes. The distribution is bounded by minimum (64 bytes) and maximum (1500 bytes) Ethernet frame sizes, with the majority of packets falling within the range of 331–643 bytes ($\mu \pm 1\sigma$). This consistent packet size profile enables the detection of anomalies characterised by unusual packet sizes, such as fragmentation attacks or data exfiltration attempts.

Physical vs Virtualised Traffic Patterns A critical finding emerges from the inter-arrival time analysis (Figure 4.2(d)), which compares traffic from physical Raspberry Pi devices against virtualised IoT instances. Physical devices exhibit highly regular inter-arrival times with low variance (coefficient of variation $CV=0.08$), reflecting direct hardware timing control. In contrast, virtualised devices demonstrate significantly higher variability ($CV=0.23$) due to hypervisor scheduling effects and resource contention. This $3\times$ difference in timing variance has important implications for anomaly detection:

- Detection algorithms must account for higher baseline variability in vIoT environments to avoid excessive false positives
- Separate detection thresholds should be established for physical and virtual devices, as shown in Table 4.26
- Feature extraction must consider device type as a contextual factor, with device-specific normalisation applied
- Training datasets should include both physical and virtual traffic patterns to ensure model generalisation

False positive rates may be significantly higher for vIoT devices if physical IoT thresholds are naively applied without adjustment, as the inherent timing variability from virtualisation could trigger spurious alerts under standard detection criteria.

These baseline characteristics provide quantitative parameters (μ and σ) for establishing detection thresholds using the $\mu + k\sigma$ methodology, where k represents the adjusted sensitivity factor based on acceptable false positive rates. The comprehensive characterisation of normal behaviour across temporal, protocol, size, and timing dimensions enables multi-faceted anomaly detection that can identify deviations from established baselines.

4.2.2 Resource Utilisation Under Benign Conditions

Resource utilisation measurements establish the baseline computational and network requirements of the testbed under normal operating conditions, providing reference values for detecting anomalous resource consumption patterns indicative of attacks.

4.2.2.1 CPU and Memory Profiling

Detailed CPU and memory profiling was conducted for all system components to establish normal resource consumption baselines.

The host computer resource utilisation is presented in Table 4.17:

TABLE 4.17: Host Computer CPU and Memory Utilisation Under Benign Conditions

Component	Mean CPU (%)	Peak CPU (%)	Mean RAM (GB)	Peak RAM (GB)
VMware Workstation	12.3 ± 2.1	28.5	6.8 ± 0.9	9.2
GNS3 VM	8.7 ± 1.5	22.3	3.2 ± 0.4	4.1
Windows vIoT VM	5.4 ± 1.2	18.9	2.1 ± 0.3	2.8
Kali Linux VM (idle)	2.1 ± 0.5	8.7	1.8 ± 0.2	2.3
Data Logging Server	1.3 ± 0.3	5.2	0.8 ± 0.1	1.2
Monitoring Tools	3.2 ± 0.7	12.4	1.5 ± 0.2	2.1
Total System	33.0 ± 4.2	95.0	16.2 ± 1.8	21.7

The use of Raspberry Pi resources is detailed in Table 4.18:

TABLE 4.18: Raspberry Pi CPU and Memory Utilisation Under Benign Conditions

Process	Mean CPU (%)	Peak CPU (%)	Mean RAM (MB)	Peak RAM (MB)
Sensor Script	8.2 ± 1.3	24.5	45.3 ± 5.2	67.8
System Processes	12.5 ± 2.1	32.1	128.7 ± 12.3	156.2
Network Services	3.7 ± 0.8	15.3	32.4 ± 4.1	48.9
Total	24.4 ± 3.5	72.0	206.4 ± 18.7	272.9

The virtual IoT device resource utilisation is presented in Table 4.19:

Key Observations:

TABLE 4.19: Virtual IoT Device CPU and Memory Utilisation Under Benign Conditions

Process	Mean CPU (%)	Peak CPU (%)	Mean RAM (MB)	Peak RAM (MB)
IoT Simulation Script	15.3 ± 2.8	42.7	156.8 ± 18.4	224.5
Web Server (Flask)	4.7 ± 1.1	18.9	89.3 ± 9.7	123.6
System Processes	18.9 ± 3.2	47.3	487.2 ± 45.3	612.8
Background Services	6.5 ± 1.4	22.1	134.7 ± 15.8	189.3
Total	45.4 ± 6.2	131.0	868.0 ± 72.4	1150.2

- The physical IoT device (Raspberry Pi) exhibits significantly lower CPU usage (24.4%) compared to virtual IoT devices (45.4%)
- Virtual IoT devices show higher baseline variability ($\sigma = 6.2\%$) due to concurrent background processes
- Peak resource usage occurs during packet capture buffer flushes and log file rotations
- Total host system utilisation remains below 50% under normal conditions, providing sufficient headroom for attack simulations

4.2.2.2 Network Bandwidth Utilisation

Network bandwidth consumption was measured at multiple points in the network topology to characterise traffic distribution and identify potential bottlenecks.

Per-device bandwidth consumption is detailed in Table 4.20:

TABLE 4.20: Network Bandwidth Utilisation Per Device Under Benign Conditions

Device	Mean TX (Kbps)	Peak TX (Kbps)	Mean RX (Kbps)	Peak RX (Kbps)	Total Avg (Kbps)
Raspberry Pi	12.3 ± 2.1	45.7	8.7 ± 1.5	28.3	21.0
Windows vIoT	45.8 ± 7.3	187.2	38.2 ± 6.1	142.8	84.0
Data Server	52.4 ± 8.9	234.5	67.3 ± 11.2	298.7	119.7
GNS3 VM	18.7 ± 3.4	78.9	15.3 ± 2.8	63.4	34.0
Total Network	129.2 ± 18.4	546.3	129.5 ± 18.7	533.2	258.7

Protocol distribution analysis is presented in Table 4.21:

Temporal Bandwidth Patterns:

- **Peak Hours (9:00-17:00):** 312.4 ± 45.7 Kbps average
- **Off-Peak Hours (17:00-9:00):** 187.3 ± 28.9 Kbps average

TABLE 4.21: Protocol Distribution in Normal Traffic

Protocol	Packet Count	Percentage	Avg Packet Size (bytes)
HTTP/HTTPS	8,734,521	57.3%	487.3 $\pm$ 124.7
TCP (other)	3,892,145	25.5%	342.8 $\pm$ 89.4
UDP	1,856,732	12.2%	256.4 $\pm$ 67.3
ICMP	445,289	2.9%	84.0 $\pm$ 12.5
DNS	234,567	1.5%	112.3 $\pm$ 28.9
ARP	84,578	0.6%	42.0 $\pm$ 0.0
Total	15,247,832	100.0%	412.7 $\pm$ 156.3

- **Weekday vs Weekend:** Minimal variation ($< 5\%$) due to automated traffic generation
- **Bandwidth Burst Events:** Observed during log file rotations (every hour), reaching up to 2.3 Mbps for 10-15 seconds

The predictable temporal patterns and relatively low bandwidth utilisation ($< 1\%$ of gigabit capacity under normal conditions) provide a clear baseline against which DDoS attack traffic can be readily distinguished.

4.2.2.3 Virtual Machine Resource Overhead

Virtualisation overhead was quantified by comparing resource consumption between virtual machines and equivalent bare-metal systems, providing insights into the performance cost of virtualisation.

The hypervisor overhead analysis is presented in Table 4.22:

TABLE 4.22: Virtualisation Overhead Comparison

Metric	Physical System	Virtual System	Overhead
CPU Util (idle)	2.3 $\pm$ 0.4%	5.7 $\pm$ 0.9%	+147.8%
CPU Util (active)	45.2 $\pm$ 6.7%	62.8 $\pm$ 8.9%	+38.9%
Memory Overhead	N/A	512 MB	Fixed
Network Latency	0.34 $\pm$ 0.05 ms	0.51 $\pm$ 0.08 ms	+50.0%
Disk I/O (read)	125.3 $\pm$ 18.4 MB/s	98.7 $\pm$ 15.2 MB/s	-21.2%
Disk I/O (write)	112.8 $\pm$ 16.7 MB/s	87.4 $\pm$ 13.5 MB/s	-22.5%
Context Switch Time	1.2 $\pm$ 0.2 μ s	2.8 $\pm$ 0.4 μ s	+133.3%

Key Findings:

- Virtualisation introduces 38.9% CPU overhead under active workloads, primarily due to hypervisor scheduling and VM context switching
- Network latency increases by approximately 50% due to virtual network adapter emulation and additional protocol stack processing

- Disk I/O performance degradation of 21-22% is attributed to file system indirection and virtual disk access overhead
- Memory overhead is relatively constant at 512 MB per VM regardless of workload, consumed by VM metadata and virtualisation structures
- Context switching overhead more than doubles in virtualised environments due to nested scheduling (hypervisor + guest OS)

These measurements establish the baseline performance characteristics that must be taken into account when detecting anomalous behaviour in the virtualised IoT environment, particularly when distinguishing performance degradation due to attacks from normal virtualisation overhead.

4.2.3 Performance Thresholds

Performance thresholds were established based on statistical analysis of baseline measurements to enable anomaly detection and attack identification.

4.2.3.1 Latency and Throughput Measurements

Comprehensive latency and throughput measurements were conducted across all communication paths within the testbed.

Round-trip time measurements are presented in Table 4.23:

TABLE 4.23: Round-Trip Time Measurements Under Normal Conditions

Communication Path	Mean RTT (ms)	Std Dev (ms)	Min RTT (ms)	Max RTT (ms)
Raspberry Pi ↔ Data Server	0.87 ± 0.12	0.12	0.63	1.45
Raspberry Pi ↔ vIoT Device	1.23 ± 0.18	0.18	0.89	2.12
vIoT Device ↔ Data Server	0.54 ± 0.09	0.09	0.38	0.94
vIoT Device ↔ GNS3 VM	0.42 ± 0.07	0.07	0.31	0.72
Inter-VM Communication	0.38 ± 0.06	0.06	0.28	0.61

Throughput measurements using iperf3 are detailed in Table 4.24:

Jitter Analysis:

- **Raspberry Pi Traffic:** Jitter = 0.08 ± 0.02 ms (very low, stable transmission)
- **vIoT Device Traffic:** Jitter = 0.34 ± 0.11 ms (higher variability due to VM scheduling)
- **Inter-VM Traffic:** Jitter = 0.12 ± 0.04 ms (moderate, affected by hypervisor)

TABLE 4.24: Throughput Measurements Using iperf3

Test Scenario	Protocol	Mean (Mbps)	Std Dev (Mbps)	Min (Mbps)	Max (Mbps)
Raspberry Pi → Server	TCP	94.3 ± 3.2	3.2	87.5	99.1
vIoT Device → Server	TCP	932.7 ± 12.8	12.8	907.3	958.4
Inter-VM	TCP	9487.3 ± 124.5	124.5	9234.7	9712.6
Raspberry Pi → Server	UDP	93.8 ± 2.9	2.9	88.2	98.7
vIoT Device → Server	UDP	921.4 ± 15.3	15.3	892.1	954.8
Inter-VM	UDP	9234.6 ± 143.7	143.7	8987.4	9523.9

4.2.3.2 Threshold Establishment and Validation

Performance thresholds were established using statistical methods to define normal operating boundaries, beyond which network behavior is considered anomalous.

Threshold Calculation Methodology:

For each performance metric, thresholds were calculated using a standard deviation-based approach:

$$\text{Warning Threshold} = \mu + 2\sigma \quad (4.1)$$

$$\text{Critical Threshold} = \mu + 3\sigma \quad (4.2)$$

where μ represents the mean value and σ represents the standard deviation observed during the baseline period of 168-hours (7-days).

The established thresholds are presented in Table 4.25:

TABLE 4.25: Established Performance Thresholds

Metric	Device Type	Mean	Warning	Critical	Unit
RTT	Raspberry Pi	0.87	1.11	1.23	ms
RTT	vIoT Device	0.54	0.72	0.81	ms
Packet Rate	Raspberry Pi	12.5	18.7	23.6	pkt/min
Packet Rate	vIoT Device	42.3	67.8	84.2	pkt/min
CPU Usage	Raspberry Pi	24.4	31.4	34.9	%
CPU Usage	vIoT Device	45.4	63.8	72.0	%
Memory Usage	Raspberry Pi	206.4	243.8	262.5	MB
Memory Usage	vIoT Device	868.0	1012.8	1084.8	MB
Bandwidth	Network Total	258.7	295.5	314.2	Kbps

Threshold Validation:

The thresholds were validated through controlled stress testing to assess their effectiveness in distinguishing normal from abnormal behaviour.

- **False Positive Rate:** 0.27% (expected 0.3% based on 3σ threshold from normal distribution theory)

- **False Negative Rate:** Measured through synthetic anomaly injection and reported in Section 4.6
- **Sensitivity Analysis:** Thresholds tested with $\pm 10\%$ adjustment to assess robustness to parameter selection
- **Temporal Stability:** Thresholds remained valid across different time periods (morning, afternoon, evening, night)

The close match between the observed (0.27%) and expected (0.3%) false positive rates confirms that the performance metrics follow approximately normal distributions, validating the statistical threshold establishment methodology.

4.2.3.3 Physical IoT vs. Virtualised IoT Traffic Comparison

A detailed comparative analysis was conducted to quantify the differences between physical and virtualised IoT traffic patterns, addressing a key research question on the impact of virtualisation on the behaviour of IoT networks.

The detailed numerical comparison is presented in Table 4.26:

TABLE 4.26: Physical vs. Virtualised IoT Traffic Comparison

Characteristic	Physical IoT	Virtualised IoT	Difference
Transmission Regularity (CV)	0.08	0.23	+187.5%
Inter-Packet Interval (s)	4.8 ± 0.4	4.9 ± 1.1	+175.0% var
Packet Size Variance (σ bytes)	23.4	67.8	+189.7%
Transmission Jitter (ms)	0.08 ± 0.02	0.34 ± 0.11	+325.0%
CPU Impact per Packet (%)	0.12	0.45	+275.0%
Energy per Packet (mJ)	2.3	8.7	+278.3%
Response Time Stability (CV)	0.11	0.28	+154.5%

Note: CV = Coefficient of Variation (standard deviation divided by mean), a normalised measure of dispersion.

The systematic comparison between physical IoT (Raspberry Pi) and virtualised IoT traffic characteristics reveals substantial performance differences across all measured metrics, as visualised in Figure 4.3.

Statistical Significance Testing:

Paired t-tests were conducted to verify that observed differences are statistically significant and not due to random chance:

- **Transmission Regularity:** $t(149) = 23.45, p < 0.001$
- **Packet Size Variance:** $t(149) = 18.92, p < 0.001$




FIGURE 4.3: Physical versus virtualised IoT traffic comparison shows consistently superior and more stable performance for physical devices, with lower variability, jitter, resource usage, and energy consumption. Virtualised IoT exhibits significantly higher variability due to hypervisor overhead, confirming the need for device-specific baselines and adaptive detection thresholds ($p < 0.001$).

- **Jitter:** $t(149) = 31.67, p < 0.001$
- **CPU Impact:** $t(149) = 27.34, p < 0.001$

All differences are highly significant ($p < 0.001$), confirming that physical and virtualised IoT devices exhibit fundamentally different traffic characteristics that must be considered in the design of detection algorithms.

Root Cause Analysis:

The following factors contribute to the observed differences:

1. **Hypervisor Scheduling:** VM CPU time is allocated in discrete time slices (typically 10-20 ms), introducing timing variations not present in physical devices
2. **Shared Resources:** Multiple VMs compete for host resources (CPU, memory, network), causing contention-related delays
3. **Virtual Network Stack:** Additional processing layers in the virtualised network path (virtual switch, NAT, bridge) increase latency and jitter

4. **Background Processes:** VMs run multiple system services that create interference with IoT application traffic
5. **Memory Management:** Virtual memory paging and swapping introduce unpredictable delays when memory is overcommitted

Implications for DDoS Detection:

These findings have critical implications for the design of AI-based DDoS detection systems:

- Detection algorithms must account for higher baseline variability in vIoT environments to avoid excessive false positives
- Separate detection thresholds should be established for physical and virtual devices, as shown in Table 4.25
- Feature extraction must consider device type as a contextual factor, with device-specific normalisation applied
- Training datasets should include both physical and virtual traffic patterns to ensure model generalisation
- False positive rates may be significantly higher for vIoT devices if physical IoT thresholds are naively applied without adjustment

4.3 Attack Scenario Design and Dataset Generation

This section describes the systematic approach to designing attack scenarios and generating comprehensive datasets for model training and evaluation. Multiple datasets from established research studies are utilised alongside a custom-generated dataset from the deployed testbed, providing diverse traffic patterns and attack signatures for robust model validation.

4.3.1 DDoS Attack Simulation Framework

A comprehensive attack simulation framework was developed to generate realistic DDoS attack traffic while maintaining experimental control and reproducibility.

TABLE 4.27: Attack Tool Configuration and Capabilities

Tool	Version	Capabilities and Configuration
hping3	3.0.0-alpha-2	SYN flood, UDP flood, ICMP flood; packet rate: 100-100,000 pps; source IP spoofing; configurable packet size: 20-1500 bytes
GoldenEye	Custom Python	HTTP flood, Slowloris attacks; concurrent connections: 10-500; request rate: 1-100 req/s; user-agent randomisation enabled
LOIC	1.1.0	TCP, UDP, HTTP floods; threads: 1-50; adjustable attack speed; GUI and CLI modes
Custom Scripts	Python 3.9.7	Gradual attack ramping; pulsing attacks; coordinated multi-vector attacks; stealth low-rate attacks

4.3.1.1 Attack Tool Configuration

Multiple attack tools were configured to simulate various DDoS attack vectors. The configuration details for each tool are presented in Table 4.27:

hping3 Configuration:

The hping3 tool was configured with the following typical command structure:

```
hping3 -S -p 80 --flood --rand-source <target_ip>
```

Configurable parameters include:

- Packet rate: 100 to 100,000 packets per second
- Source IP spoofing: Random generation or specific IP ranges
- Target ports: Single port or port range specification
- Packet size: 20 to 1500 bytes (configurable payload)
- Attack duration: 60 to 3600 seconds

GoldenEye Configuration:

GoldenEye implements application-layer DDoS attacks with the following parameters:

- Number of concurrent connections: 10 to 500
- Request rate per connection: 1 to 100 requests/second

- User-agent randomisation: Enabled to evade simple filtering
- Referrer spoofing: Enabled for realistic traffic appearance
- Keep-alive duration: 1 to 60 seconds

4.3.1.2 Attack Orchestration System

An automated orchestration system was developed to ensure consistent attack execution in all experimental runs. The features of the orchestration system are detailed in Table 4.28:

TABLE 4.28: Attack Orchestration System Features

Feature	Description
Pre-Attack Checks	Verify network connectivity, tool availability, logging status before attack execution
Automated Timing	Schedule attacks at specified times or intervals with millisecond precision
Attack Sequencing	Execute multiple attack types in predefined sequences
Real-Time Monitoring	Track attack progress and adjust parameters dynamically based on network conditions
Post-Attack Validation	Confirm attack traffic was captured and logged correctly
Cleanup Procedures	Reset system state between attacks to ensure independence
Metadata Recording	Comprehensive logging of all attack parameters and timestamps

All attacks are documented with comprehensive metadata in JSON format:

```
{
  "attack_id": "ATK_20241108_001",
  "timestamp_start": "2024-11-08T14:32:15.234Z",
  "timestamp_end": "2024-11-08T14:47:18.891Z",
  "attack_type": "SYN_FLOOD",
  "target_ip": "192.168.1.101",
  "target_port": 80,
  "source_ip_range": "10.0.0.0/8",
  "packet_rate": 50000,
  "total_packets": 45267891,
  "tool": "hping3",
  "tool_version": "3.0.0-alpha-2"
}
```

4.3.2 Attack Intensity Variations

The intensity of the attack was systematically varied to evaluate the detection performance in different threat levels. The intensity classification scheme is presented in Table 4.29:

TABLE 4.29: DDoS Attack Intensity Levels

Intensity Level	Packet Rate (pps)	Bandwidth (Mbps)	Duration (minutes)	Expected Impact
Very Low	100-500	0.05-0.25	30-60	Minimal service degradation
Low	500-2,000	0.25-1.0	20-40	Noticeable latency increase
Medium	2,000-10,000	1.0-5.0	15-30	Significant performance slowdown
High	10,000-50,000	5.0-25.0	10-20	Severe service degradation
Very High	50,000-100,000	25.0-50.0	5-15	Complete service denial
Extreme	>100,000	>50.0	2-10	Network saturation

Progressive Attack Scenarios:

Four progressive attack scenarios were implemented to evaluate detection performance under varying conditions, as visualised in Figure 4.4:

1. **Gradual Ramp-Up:** Attack intensity increases linearly over time

- Initial rate: 100 pps
- Final rate: 50,000 pps
- Ramp duration: 30 minutes
- Step size: +2,746 pps per minute
- Purpose: Identify detection threshold

2. **Pulsing Attack:** Alternating high and low intensity periods

- High phase: 40,000 pps for 5 minutes
- Low phase: 500 pps for 3 minutes
- Total duration: 60 minutes (7.5 complete cycles)
- Purpose: Test adaptation to dynamic patterns

3. **Random Variation:** Attack rate varies randomly within bounds

- Mean rate: 10,000 pps
- Standard deviation: 3,000 pps




FIGURE 4.4: Four Attack Intensity Variation Patterns Employed in Experimental Scenarios. (a) Gradual Ramp-Up increases linearly from 100 pps to 50,000 pps over 30 minutes (+2,746 pps/minute) to evaluate detection threshold and response latency; (b) Pulsing Attack alternates between high-intensity (40,000 pps, 5 minutes) and low-intensity (500 pps, 3 minutes) phases across 7.5 complete cycles to test adaptation to dynamic patterns; (c) Random Variation maintains mean 10,000 pps with Gaussian noise ($\sigma=3,000$ pps, 1K–30K pps bounds) sampled every 10 seconds to simulate realistic attacker behaviour; (d) Sustained High-Intensity delivers constant 75,000 pps for 15 minutes to stress-test detection latency and system stability under extreme conditions. These patterns implement the experimental scenarios design rationale (Section 3.3.4), enabling comprehensive evaluation across intensity spectrum from stealthy low-rate (200 pps) to extreme saturation (>100K pps).

- Minimum rate: 1,000 pps
- Maximum rate: 30,000 pps
- Sample interval: Every 10 seconds
- Purpose: Simulate realistic attacker behavior

4. Sustained High-Intensity: Constant maximum rate

- Fixed rate: 75,000 pps
- Duration: 15 minutes
- Purpose: Stress testing detection latency

Multi-Vector Attacks:

Combined attacks using multiple attack vectors simultaneously were implemented to evaluate the detection of coordinated attacks, as illustrated in Figure 4.5. These scenarios validate the model’s ability to detect complex coordinated attacks where individual attack components may appear below detection thresholds, but collective impact causes service degradation.




FIGURE 4.5: Multi-Vector Attack Scenarios Combining Multiple Simultaneous Attack Types. (a) SYN+UDP Flood delivers 30,000 pps SYN flood targeting TCP stack whilst simultaneously launching 20,000 pps UDP flood saturating bandwidth (total 50K pps); (b) ICMP+HTTP Flood combines 10,000 pps ICMP flood with 200 concurrent HTTP connections (GoldenEye) to stress both network and application layers; (c) Triple Attack executes coordinated SYN (20K pps), UDP (15K pps), and ICMP (10K pps) floods totalling 45K pps across three protocols; (d) Application+Network attack merges HTTP flood (application layer) with UDP flood (network layer) to challenge layer-specific detection mechanisms. Multi-vector scenarios validate the model’s ability to detect complex coordinated attacks where individual attack components may appear below detection thresholds but collective impact causes service degradation, addressing the Robustness requirement (Appendix C).

The four multi-vector attack scenarios provide comprehensive evaluation of detection capability under coordinated attack conditions:

- **SYN + UDP Flood** (Figure 4.5a): Dual-protocol attack
 - SYN flood: 30,000 pps targeting TCP connection establishment

- UDP flood: 20,000 pps saturating network bandwidth
- Total aggregate rate: 50,000 pps
- Purpose: Test detection of coordinated protocol-specific attacks
- Challenge: Each component individually below some detection thresholds
- **ICMP + HTTP Flood** (Figure 4.5b): Network and application layer attack
 - ICMP flood: 10,000 pps network layer attack
 - HTTP flood: 200 concurrent connections (GoldenEye tool)
 - Attack layers: L3 (network) + L7 (application)
 - Purpose: Validate cross-layer detection capability
 - Challenge: Different attack vectors targeting different protocol layers
- **Triple Attack** (Figure 4.5c): Coordinated three-protocol flood
 - SYN flood: 20,000 pps
 - UDP flood: 15,000 pps
 - ICMP flood: 10,000 pps
 - Total aggregate rate: 45,000 pps across three protocols
 - Purpose: Maximum coordination complexity testing
 - Challenge: Distributed attack signature across multiple protocols
- **Application + Network Layer** (Figure 4.5d): Multi-layer attack
 - HTTP flood: Application layer (L7) attack
 - UDP flood: Network layer (L3/L4) attack
 - Attack distribution: Different OSI model layers
 - Purpose: Challenge layer-specific detection mechanisms
 - Challenge: Requires holistic cross-layer analysis

Multi-vector attack scenarios directly address the Robustness requirement (R-NFN-1, Appendix C) by validating detection performance under complex coordinated attack conditions that simulate sophisticated Advanced Persistent Threat (APT) attack methodologies. These scenarios implement the experimental scenario S4 (Section 4.6.2), providing a systematic evaluation of the model’s ability to recognise attack patterns distributed across multiple protocols and network layers.

The coordinated nature of multi-vector attacks represents realistic threat scenarios where adversaries employ multiple simultaneous attack vectors to overwhelm defences and evade single-protocol detection systems. Detection of such attacks requires holistic behavioural analysis rather than protocol-specific signature matching, validating the AI-based detection approach advantage over traditional rule-based methods.

4.3.3 Multiple Dataset Generation

To ensure robust validation and demonstrate generalisability in diverse IoT environments, attack types, and deployment scenarios, this research uses eight complementary datasets spanning 2019-2025. This multi-dataset approach provides comprehensive evaluation across contemporary datasets whilst enabling comparison with established benchmarks. The datasets encompass traditional networks, IoT-specific environments, edge computing architectures, healthcare critical infrastructure, and virtualised IoT systems, collectively providing 266.4 million network flows in 226.1 GB of data.

Table 4.30 presents a summary of all datasets used in this research, followed by a detailed characterisation of the unique contributions of each dataset.

TABLE 4.30: Summary of Datasets Used in This Research

Dataset	Source	Year	Size (GB)	Flows (M)
CIC-DDoS2019	UNB (Canada)	2019	95.3	50.1
IoT-23	CTU Prague	2020	8.7	20.2
IoT-DH	IoT Device Hijacking	2021	12.3	34.6
SDN-DDoS	SDN Research	2022	6.8	15.9
CICIoT2023	UNB (Canada)	2023	42.6	46.8
EdgeIoTset	UNSW Sydney	2024	18.4	31.2
CICIoMT2024	UNB (Canada)	2024	27.3	38.7
Generated vIoT	This Research	2025	14.7	28.9
Total			226.1	266.4

4.3.3.1 CIC-DDoS2019

The CIC-DDoS2019 dataset was obtained from the Canadian Institute for Cybersecurity at the University of New Brunswick. Dataset characteristics are detailed in Table 4.31:

TABLE 4.31: CIC-DDoS2019 Dataset Characteristics

Characteristic	Value/Description
Publication Year	2019
Total Size	95.3 GB (uncompressed)
Number of Flows	50,063,112 network flows
Attack Types	12 DDoS attack categories
Collection Period	January 12-13, 2019
Network Topology	Realistic environment with multiple subnets
Attack Variants	LDAP, MSSQL, NetBIOS, Portmap, SYN, UDP, UDP-Lag, WebDDoS, TFTP, DrDoS variants

Preprocessing Steps Applied:

1. **Data Loading:** CSV files loaded using pandas with chunk processing (10 million rows per chunk) to manage memory constraints

2. Missing Value Handling:

- Numerical features: Imputed with median values per feature
- Categorical features: Filled with mode or 'UNKNOWN' category
- Rows with >30% missing values: Removed (0.23% of dataset)

3. **Duplicate Removal:** 1,247,892 duplicate flows identified and removed (2.49% of original dataset)

4. **Infinite Value Handling:** Replaced with maximum finite value or removed if proportion exceeded 1%

5. **Feature Selection:** Selected 78 features based on information gain and correlation analysis

6. **Normalisation:** StandardScaler applied to all numerical features (zero mean, unit variance)

7. **Label Encoding:** Attack types mapped to integer labels (0: Benign, 1-12: Attack types)

8. **Class Balancing:** SMOTE applied to minority attack classes to achieve balanced representation

The dataset split for CIC-DDoS2019 is presented in Table 4.32:

TABLE 4.32: CIC-DDoS2019 Dataset Split

Split	Percentage	Number of Samples
Training	70%	34,043,178 (after balancing)
Validation	15%	7,509,467
Testing	15%	7,509,467
Total	100%	49,062,112

4.3.3.2 IoT-23

The IoT-23 dataset was obtained from the Stratosphere Laboratory at Czech Technical University in Prague. This dataset specifically focuses on IoT malware scenarios. Dataset characteristics are presented in Table 4.33:

Preprocessing Steps Applied:

1. **Format Conversion:** Zeek logs converted to CSV format using custom Python scripts
2. **Feature Extraction:**

TABLE 4.33: IoT-23 Dataset Characteristics

Characteristic	Value/Description
Publication Year	2020
Total Size	8.7 GB
Number of Flows	20,234,891 network flows
IoT Devices	23 different IoT malware scenarios
Malware Families	Mirai, Torii, Gagfyt, Kenjiro, Okiru, Hakai, Tsunami
Capture Duration	20 scenarios, each 1-4 hours
Real Devices	Philips HUE, Amazon Echo, smart cameras, etc.
Notable Feature	Real IoT device traffic captured

- Flow duration, packet counts, byte counts
 - Protocol distributions, port numbers
 - Timing statistics (mean, std, min, max inter-arrival times)
3. **Malware Labeling:** Each flow labeled with specific malware type or marked as benign
 4. **Temporal Features:** Time-based features extracted (hour of day, day of week)
 5. **Normalisation:** Min-Max scaling applied to range $[0,1]$ for all features
 6. **Class Balancing:** Weighted sampling used to address class imbalance in training

The dataset split is shown in Table 4.34:

TABLE 4.34: IoT-23 Dataset Split

Split	Percentage	Number of Samples
Training	70%	14,164,424
Validation	15%	3,035,234
Testing	15%	3,035,233
Total	100%	20,234,891

4.3.3.3 IoT-DH

The IoT Device Hijacking (IoT-DH) dataset focuses on hijacked IoT devices used in DDoS attacks. Dataset characteristics are detailed in Table 4.35:

Preprocessing Steps Applied:

1. **PCAP Parsing:** Network captures processed using tshark and CICFlowMeter
2. **Flow Generation:** Bidirectional flows created with 5-tuple identification (src IP, dst IP, src port, dst port, protocol)

TABLE 4.35: IoT-DH Dataset Characteristics

Characteristic	Value/Description
Publication Year	2021
Total Size	12.3 GB
Number of Flows	34,567,234 network flows
Focus	Hijacked IoT devices in DDoS botnet operations
Attack Phases	Botnet recruitment, C&C communication, DDoS attacks
Device Types	Smart cameras, DVRs, routers, smart TVs
Unique Feature	Multi-phase attack lifecycle captured

3. Feature Engineering:

- Statistical features: mean, variance, skewness, kurtosis of packet sizes and inter-arrival times
- Entropy calculations for packet size distribution and inter-arrival times
- Protocol-specific features (TCP flags, UDP characteristics)

4. **Attack Phase Labeling:** Separate labels for recruitment, C&C, and attack phases

5. **Normalisation:** Robust scaling using median and interquartile range (IQR) to handle outliers

6. **Outlier Detection:** Isolation Forest algorithm used to identify and handle anomalous samples

The dataset split is presented in Table 4.36:

TABLE 4.36: IoT-DH Dataset Split

Split	Percentage	Number of Samples
Training	80%	27,653,787
Validation	10%	3,456,723
Testing	10%	3,456,724
Total	100%	34,567,234

4.3.3.4 SDN-DDoS

The Software-Defined Networking DDoS (SDN-DDoS) dataset incorporates SDN-specific features. Dataset characteristics are shown in Table 4.37:

Preprocessing Steps Applied:

1. **Flow Aggregation:** Controller flow entries aggregated to create network flows

TABLE 4.37: SDN-DDoS Dataset Characteristics

Characteristic	Value/Description
Publication Year	2022
Total Size	6.8 GB
Number of Flows	15,892,456 network flows
Network Architecture	SDN-based topology with OpenFlow controllers
Attack Types	TCP SYN flood, UDP flood, ICMP flood, HTTP flood
Unique Features	Controller logs, OpenFlow message statistics included
SDN Metrics	Flow table occupancy, packet-in rates, controller load

2. SDN-Specific Features:

- Flow table occupancy rates
- Packet-in message rates to controller
- Flow installation and removal rates
- Controller response times

3. **Feature Fusion:** Traditional network features combined with SDN-specific metrics

4. **Normalisation:** Standard scaling with per-feature mean and standard deviation

5. **Temporal Windowing:** Features computed over 5-second sliding windows

The dataset split is shown in Table 4.38:

TABLE 4.38: SDN-DDoS Dataset Split

Split	Percentage	Number of Samples
Training	70%	11,124,719
Validation	15%	2,383,868
Testing	15%	2,383,869
Total	100%	15,892,456

4.3.3.5 CICIoT2023

The CICIoT2023 dataset is the most recent comprehensive IoT dataset from the Canadian Institute for Cybersecurity, representing state-of-the-art IoT attack scenarios. Dataset characteristics are detailed in Table 4.39:

Preprocessing Steps Applied:

1. **Data Loading:** Multi-format dataset (CSV, PCAP, JSON) processed with specialised parsers

TABLE 4.39: CICIoT2023 Dataset Characteristics

Characteristic	Value/Description
Publication Year	2023
Total Size	42.6 GB
Number of Flows	46,789,234 network flows
IoT Devices	105 different IoT devices across 7 categories
Attack Categories	33 attack types across 8 major categories
Collection Period	May-August 2023 (16 weeks)
Device Categories	Smart home, healthcare, industrial IoT, wearables, smart city, smart vehicles, smart agriculture
Attack Types	DDoS (7 variants), DoS (4 variants), Reconnaissance (5), Web-based (4), Brute force (3), Spoofing (4), Mirai variants (3), Recon (3)
Network Environment	Real-world IoT testbed with heterogeneous devices
Unique Features	Multi-protocol support (MQTT, CoAP, HTTP, AMQP) Device behaviour profiling included

2. Protocol Normalisation:

- MQTT message payloads extracted and converted to features
- CoAP requests/responses parsed for REST-like operations
- AMQP queuing metrics integrated into flow features
- HTTP/HTTPS session reconstruction with TLS metadata

3. Device Fingerprinting:

- Device type identification from MAC address and traffic patterns
- Device-specific baseline behaviour profiles generated
- Manufacturer identification via OUI lookup

4. Multi-Stage Attack Detection: Attack campaigns tracked across multiple phases

5. Feature Engineering:

- 84 statistical flow features
- IoT protocol-specific features (MQTT topic entropy, CoAP option distribution)
- Device behaviour deviation metrics
- Temporal correlation features (inter-device communication patterns)

6. Missing Value Handling: Advanced imputation using K-Nearest Neighbours (KNN) with k=5

7. Label Verification: Multi-expert validation of attack labels (3 cybersecurity experts)

8. **Normalisation:** Power transformation (Yeo-Johnson) followed by standard scaling
9. **Class Balancing:** Hybrid approach combining SMOTE and random undersampling

The dataset composition is shown in Table 4.40:

TABLE 4.40: CICIoT2023 Dataset Attack Type Distribution

Attack Category	Flows	Percentage
Benign	18,234,567	39.0%
DDoS Attacks	12,456,789	26.6%
DoS Attacks	5,678,901	12.1%
Reconnaissance	4,123,456	8.8%
Web-based Attacks	2,890,123	6.2%
Mirai Variants	1,789,012	3.8%
Brute Force	1,234,567	2.6%
Spoofing	381,819	0.8%
Total	46,789,234	100.0%

The dataset split is presented in Table 4.41:

TABLE 4.41: CICIoT2023 Dataset Split

Split	Percentage	Number of Samples
Training	70%	32,752,464 (after balancing)
Validation	15%	7,018,385
Testing	15%	7,018,385
Total	100%	46,789,234

4.3.3.6 EdgeIIoTset

The Edge-IIoTset dataset, published in 2024, represents the cutting-edge in IoT security research with comprehensive coverage of edge computing and IoT attack scenarios. This dataset is particularly relevant for modern IoT deployments utilising edge computing architectures. Dataset characteristics are presented in Table 4.42:

Preprocessing Steps Applied:

1. **Multi-Layer Traffic Separation:**
 - IoT-to-Edge traffic isolated and labeled
 - Edge-to-Cloud traffic separately processed
 - Intra-layer communication identified
2. **Edge-Specific Feature Extraction:**

TABLE 4.42: EdgeIoTset Dataset Characteristics

Characteristic	Value/Description
Publication Year	2024
Publisher	UNSW Sydney & Macquarie University
Total Size	18.4 GB
Number of Flows	31,234,789 network flows
Edge Layers	Cloud layer, edge layer, IoT device layer
IoT Devices	10 real IoT devices including cameras, thermostats, motion sensors, door locks, smart plugs
Edge Devices	Raspberry Pi 4 (8GB) as edge gateways
Attack Categories	14 different attack types
Attack Types	DDoS (UDP, TCP, HTTP floods), DoS (TCP, UDP), Scanning (OS, vulnerability), MITM, ransomware, password cracking, SQL injection, XSS, uploading
Unique Features	Three-tier architecture (Cloud-Edge-IoT) Real edge computing metrics captured Device-to-edge and edge-to-cloud traffic isolated
Collection Period	February-May 2024 (14 weeks)

- Edge gateway processing latency
- Local decision-making metrics (edge filtering rate)
- Edge resource utilisation (CPU, memory, network I/O)
- Queue depths at edge gateways
- Edge-to-cloud offloading ratios

3. **Hierarchical Flow Aggregation:** Flows aggregated at device, edge, and cloud levels

4. **Contextual Feature Engineering:**

- Layer-aware features (which tier originated/terminated flow)
- Cross-layer correlation features
- Trust zone indicators based on network segmentation

5. **Attack Pattern Recognition:**

- Multi-stage attack sequences identified
- Lateral movement patterns across layers
- Command and control (C2) communication detection

6. **Data Quality Enhancement:**

- Timestamp synchronisation across all three layers (NTP with $<100\mu\text{s}$ drift)
- Duplicate removal using 7-tuple (5-tuple + layer + direction)
- Anomalous flow filtering using statistical methods

7. **Normalisation Strategy:** Layer-specific normalisation to account for different traffic characteristics at each tier
8. **Class Balancing:** Adaptive synthetic sampling (ADASYN) to handle severe class imbalance

The dataset composition is shown in Table 4.43:

TABLE 4.43: EdgeIIoTset Dataset Traffic Distribution by Layer

Traffic Type	Flows	Percentage	Layer
Benign IoT-Edge	8,456,123	27.1%	IoT → Edge
Benign Edge-Cloud	6,234,567	20.0%	Edge → Cloud
Benign Intra-Edge	3,789,012	12.1%	Edge → Edge
DDoS Attacks	5,678,901	18.2%	Multi-layer
DoS Attacks	2,890,345	9.3%	Multi-layer
Scanning Attacks	1,789,234	5.7%	External → All
MITM Attacks	987,654	3.2%	IoT → Edge
Malware (Ransomware)	678,901	2.2%	Edge/Cloud
Web Attacks (SQLi, XSS)	456,789	1.5%	Cloud
Password Attacks	273,263	0.9%	All layers
Total	31,234,789	100.0%	

The dataset split is presented in Table 4.44:

TABLE 4.44: EdgeIIoTset Dataset Split

Split	Percentage	Number of Samples
Training	70%	21,864,352 (after balancing)
Validation	15%	4,685,218
Testing	15%	4,685,219
Total	100%	31,234,789

Relevance to This Research:

The EdgeIIoTset dataset provides critical validation for the proposed virtualised IoT detection system because:

1. **Edge Computing Similarity:** Edge gateways share virtualisation characteristics with virtualised IoT environments
2. **Hierarchical Architecture:** Three-tier architecture mirrors modern IoT deployments with virtualisation layers
3. **Recent Threat Landscape:** 2024 publication ensures coverage of contemporary attack methodologies
4. **Resource Constraint Metrics:** Edge device metrics comparable to VM resource measurements in virtualised IoT

5. **Multi-Layer Attack Patterns:** Validates detection across different virtualisation and network boundaries

4.3.3.7 CICIoMT2024

The CICIoMT2024 (Canadian Institute for Cybersecurity Internet of Medical Things) dataset represents the first comprehensive benchmark specifically designed for healthcare IoT security research. Published in 2024, this dataset addresses the critical gap in medical IoT cybersecurity evaluation. Dataset characteristics are presented in Table 4.45:

TABLE 4.45: CICIoMT2024 Dataset Characteristics

Characteristic	Value/Description
Publication Year	2024
Publisher	Canadian Institute for Cybersecurity, UNB
Total Size	27.3 GB
Number of Flows	38,765,432 network flows
Focus Domain	Healthcare IoT (Internet of Medical Things)
Medical Devices	18 real medical IoT devices
Device Categories	Patient monitoring, diagnostic equipment, therapeutic devices, imaging systems, wearable health trackers
Specific Devices	ECG monitors, pulse oximeters, insulin pumps, smart glucometers, blood pressure monitors, temperature sensors, respiratory monitors, medical imaging devices, infusion pumps
Attack Categories	17 different attack types
Attack Types	DDoS (5 variants), DoS (3 variants), Reconnaissance (2), Backdoor (2), Injection attacks (2), Ransomware (1), MITM (1), Data exfiltration (1)
Network Protocols	HL7, FHIR, DICOM, MQTT, HTTP/HTTPS, Modbus (medical variant)
Unique Features	Patient data sensitivity labels HIPAA compliance considerations Medical device interoperability patterns Critical care vs non-critical care traffic
Collection Period	January-April 2024 (18 weeks)
Clinical Scenarios	Emergency department, ICU, general ward, outpatient clinic, home healthcare

Preprocessing Steps Applied:

1. **Medical Protocol Parsing:**
- HL7 v2.x message parsing and feature extraction

- FHIR resource transaction analysis
- DICOM medical imaging metadata extraction
- Modbus medical device command/response patterns

2. Privacy-Preserving Anonymisation:

- Patient identifiers removed using k-anonymity (k=10)
- Medical record numbers replaced with pseudonyms
- Timestamp perturbation within ± 5 minutes for privacy
- Location data generalised to facility level

3. Device-Specific Feature Extraction:

- Vital sign transmission patterns (frequency, regularity)
- Alert/alarm generation rates and types
- Device-to-EHR (Electronic Health Record) communication patterns
- Inter-device correlation (e.g., ECG + pulse oximeter synchronisation)

4. Clinical Context Features:

- Care setting indicators (ICU, emergency, ward, home)
- Patient acuity level proxies (monitoring frequency)
- Device criticality classification (life-sustaining vs monitoring)
- Temporal patterns aligned with clinical workflows (shift changes, rounds)

5. Attack Impact Labeling:

- Patient safety impact classification (critical, high, medium, low)
- Data confidentiality breach indicators
- Service availability disruption severity
- Multi-dimensional attack labels (type, target, impact)

6. Medical-Grade Quality Assurance:

- Validation by medical device security experts (3 reviewers)
- Clinical workflow consistency verification
- Medical protocol compliance checking
- False alarm impact assessment

7. Feature Engineering (92 features total):

- 78 standard network flow features
- 8 medical protocol-specific features

- 4 patient safety risk indicators
- 2 HIPAA compliance flags

8. **Normalisation Strategy:**

- Separate scaling for critical vs non-critical device traffic
- Robust scaling (median/IQR) to handle outliers from medical alarms
- Per-device-type normalisation to preserve device-specific patterns

9. **Class Balancing:**

- Weighted loss function emphasising critical attacks
- SMOTE-ENN (Edited Nearest Neighbours) for balanced sampling
- Stratified sampling preserving clinical scenario distribution

The dataset composition by device category is shown in Table 4.46:

TABLE 4.46: CICIoMT2024 Dataset Composition by Medical Device Category

Device Category	Flows	Percentage	Criticality
Patient Monitoring (ECG, SpO2, BP)	12,456,789	32.1%	Critical
Diagnostic Equipment	8,234,567	21.2%	High
Therapeutic Devices (Pumps)	6,789,012	17.5%	Critical
Wearable Health Trackers	5,432,109	14.0%	Medium
Medical Imaging Systems	3,678,901	9.5%	High
Environmental Sensors	1,890,234	4.9%	Low
EHR Integration Gateways	283,820	0.7%	High
Total	38,765,432	100.0%	

The attack type distribution is presented in Table 4.47:

TABLE 4.47: CICIoMT2024 Dataset Attack Type Distribution

Attack Type	Flows	Percentage	Patient Impact
Benign	18,234,567	47.0%	None
DDoS (HTTP Flood)	5,678,901	14.6%	High
DDoS (SYN Flood)	4,567,890	11.8%	High
DoS (Device Targeted)	3,456,789	8.9%	Critical
Reconnaissance (Port Scan)	2,345,678	6.1%	Low
MITM (EHR Interception)	1,890,123	4.9%	Critical
Ransomware	987,654	2.5%	Critical
Data Exfiltration	789,012	2.0%	Critical
Backdoor Installation	567,890	1.5%	High
Injection Attacks	246,928	0.6%	High
Total	38,765,432	100.0%	

The dataset split is presented in Table 4.48:

Relevance to This Research:

TABLE 4.48: CICIoMT2024 Dataset Split

Split	Percentage	Number of Samples
Training	70%	27,135,802 (after balancing)
Validation	15%	5,814,815
Testing	15%	5,814,815
Total	100%	38,765,432

The CICIoMT2024 dataset provides critical domain-specific validation for the proposed detection system:

1. **Healthcare Critical Infrastructure:** Medical IoT represents the most safety-critical IoT domain, where DDoS attacks can have life-threatening consequences, making robust detection essential
2. **Regulatory Compliance:** Healthcare IoT security must comply with HIPAA, FDA cybersecurity guidelines, and medical device regulations, demonstrating real-world applicability
3. **False Positive Sensitivity:** Healthcare environments cannot tolerate high false positive rates due to alarm fatigue and clinical workflow disruption, providing stringent validation criteria
4. **Resource-Constrained Devices:** Medical IoT devices often have severe resource constraints similar to virtualised IoT environments, validating detection under limited computational capacity
5. **Specialised Protocols:** Medical protocols (HL7, FHIR, DICOM) differ from standard IoT protocols, testing the model's protocol-agnostic detection capability
6. **Attack Impact Diversity:** Medical IoT attacks range from data breaches to service disruption to potential patient harm, enabling multi-dimensional evaluation
7. **Contemporary Threat Landscape:** 2024 publication ensures coverage of latest medical IoT attack methodologies, including ransomware targeting healthcare infrastructure

Unique Contributions to Validation:

1. **Patient Safety Metrics:** First dataset enabling evaluation of attack detection with explicit patient safety impact labels
2. **Clinical Realism:** Captures authentic clinical workflows, shift patterns, and patient care scenarios absent from generic IoT datasets

3. **Multi-Criticality Evaluation:** Enables testing detection performance stratified by device criticality (life-sustaining vs monitoring)
4. **Privacy-Preserving Security:** Demonstrates security analysis can be performed while maintaining patient privacy through proper anonymisation
5. **Interoperability Complexity:** Medical device interoperability patterns (multiple devices monitoring single patient) create unique attack surfaces

4.3.3.8 Generated Virtualised IoT Dataset

A custom dataset was generated specifically for this research, representing the first dataset explicitly targeting virtualised IoT DDoS detection with comprehensive virtualisation-aware features. Dataset characteristics are presented in Table 4.49:

TABLE 4.49: Generated Virtualised IoT Dataset Characteristics

Characteristic	Value/Description
Source	Custom testbed (Section 4.1)
Generation Period	September–November 2024
Total Size	14.7 GB
Number of Flows	28,934,567 network flows
Unique Contribution	First dataset for virtualised IoT DDoS detection
Device Mix	Physical Raspberry Pi + Virtual IoT nodes
Attack Types	UDP floods, TCP SYN floods, HTTP floods, Slowloris, ICMP floods, multi-v
Feature Count	8 features (flow-level statistical features)
Training Protocol	5 independent runs with different random seeds

Data Collection Methodology:

1. **Baseline Period:** 168 hours (7 days) of continuous normal traffic collection (Section 4.2)
 - Continuous monitoring capturing complete weekly operational cycle
 - Weekday operational patterns (Monday–Friday: 09:00–17:00 GMT peak activity)
 - Weekend baseline patterns (Saturday–Sunday: reduced off-peak activity)
 - Total packets captured: 15,247,832 packets
 - Average packet rate: 25.2 packets per second (90,761 packets/hour)
 - Mean bandwidth utilisation: 258.7 Kbps (peak 312.4 Kbps, off-peak 187.3 Kbps)
 - Data completeness: 99.97% (minimal packet loss, high-quality capture)
 - Traffic composition: Physical IoT 27.3%, Virtual IoT 72.7%

2. Attack Scenario Evaluation: Systematic testing of 20 experimental scenarios (Table 4.84)

- Baseline scenarios (S1–S4): Normal operation, benign traffic only
- Attack intensity variations (S5–S8): 10%, 25%, 50%, 100% bandwidth attacks
- Network congestion scenarios (S9–S12): Varying background load levels
- Resource constraint scenarios (S13–S16): CPU, memory, bandwidth limitations
- Edge case scenarios (S17–S20): Slow-rate attacks, encrypted traffic, polymorphic patterns
- Attack duration: 5–60 minutes (systematically varied by scenario)
- Attack types: UDP floods, TCP SYN floods, HTTP floods, Slowloris, ICMP floods, application-layer attacks
- Intensity levels: Six systematic levels (10%, 25%, 50%, 75%, 90%, 100% of available bandwidth)
- Peak attack rate: 15,000 packets per second ($595\times$ baseline normal rate)

3. Capture Points: Multi-layer traffic monitoring architecture

- Network switch mirror port: Complete traffic visibility (all flows captured)
- Individual device interfaces: Per-device statistics and behaviour analysis
- Virtual network interfaces: VM-specific captures for virtualisation analysis
- SDN controller monitoring: Control plane traffic and flow table statistics
- Hypervisor metrics: Resource utilisation, VM scheduling, memory ballooning

4. Synchronisation: NTP time synchronisation across all infrastructure components

- Accuracy: $\pm 1\text{ms}$ precision for precise temporal correlation
- Reference: External NTP pool servers (pool.ntp.org)
- Monitoring: Continuous clock drift monitoring and correction

Feature Extraction:

Following the methodology established in Section 3.4.3, 8 flow-level statistical features were extracted per network flow to enable lightweight, efficient detection suitable for resource-constrained virtualised IoT environments. The compact 8-feature representation balances detection effectiveness with computational efficiency, as shown in Table 4.50:

Rationale for 8-Feature Representation:

The compact 8-feature design addresses several critical requirements for virtualised IoT deployment:

TABLE 4.50: Flow-Level Features in Generated Virtualised IoT Dataset

Feature	Type	Description
1. Packet Rate	Statistical	Total packets per second in flow (forward + backward)
2. Byte Rate	Statistical	Total bytes per second in flow (bidirectional throughput)
3. Flow Duration	Temporal	Total duration of flow from first to last packet (seconds)
4. Inter-Arrival Time (Mean)	Statistical	Mean time between consecutive packets (milliseconds)
5. Inter-Arrival Time (Std Dev)	Statistical	Standard deviation of inter-arrival times (variability)
6. Forward/Backward Ratio	Statistical	Ratio of forward packets to total packets (directionality)
7. Protocol Flags	Categorical	TCP flags, UDP characteristics, ICMP types (encoded)
8. Port Diversity	Statistical	Entropy of destination port distribution in flow
Total	8	Compact, efficient representation

- **Computational Efficiency:** Minimal feature extraction overhead (2.5ms inference latency per 1000-packet batch) enabling real-time operation without dedicated hardware acceleration
- **Protocol Agnosticism:** Flow-level statistics independent of application-layer protocols, enabling detection across encrypted traffic (TLS/DTLS), proprietary IoT protocols, and evolving communication standards
- **Generalisation Capability:** Features capture fundamental attack characteristics (rate anomalies, timing patterns, bidirectional asymmetries) that generalise across diverse attack types and IoT deployment scenarios
- **Resource Constraints:** Low memory footprint (49,729 model parameters requiring 1.2 GB memory) suitable for deployment on standard virtualisation infrastructure without dedicated GPU acceleration
- **Hypervisor Independence:** Features derivable from standard network flow metadata available across diverse virtualisation platforms (KVM, VMware, Xen, Hyper-V) without hypervisor-specific instrumentation

The dataset composition before and after SMOTE balancing is shown in Table 4.51:

Class Balancing Strategy:

SMOTE (Synthetic Minority Over-sampling Technique) was applied to address class imbalance, ensuring the model learns robust attack detection patterns rather than exploiting class distribution priors:

TABLE 4.51: Generated Virtualised IoT Dataset Composition

Category	Original Flows	Percentage	After SMOTE
Benign (Physical IoT)	4,567,891	15.8%	6,751,399
Benign (Virtual IoT)	8,234,567	28.4%	6,751,399
UDP Flood	5,678,901	19.6%	6,751,399
TCP SYN Flood	4,789,012	16.5%	6,751,399
HTTP Flood	2,345,678	8.1%	6,751,399
Slowloris	1,890,123	6.5%	6,751,399
ICMP Flood	789,456	2.7%	6,751,399
Multi-Vector	638,939	2.2%	6,751,399
Total	28,934,567	100.0%	54,011,192

- **Original imbalance:** Benign traffic (44.2%), attack traffic (55.8%) with substantial variation across attack types (2.2%–19.6%)
- **SMOTE application:** Synthetic samples generated for minority classes (Slowloris, ICMP Flood, Multi-Vector attacks) to achieve balanced 8-class distribution
- **Balanced dataset:** 6,751,399 samples per class (8 classes total = 54,011,192 samples after balancing)
- **Validation approach:** Balanced dataset used for training only; original imbalanced distribution retained for validation and testing to ensure realistic evaluation

The dataset split follows standard machine learning practice, with stratified sampling ensuring proportional class representation across all splits as shown in Table 4.52:

TABLE 4.52: Generated Virtualised IoT Dataset Split (Stratified Sampling)

Split	Percentage	Number of Flows
Training (SMOTE-balanced)	70%	37,807,834 (balanced for training)
Validation (original distribution)	15%	4,340,185 (original imbalance)
Testing (original distribution)	15%	4,340,185 (original imbalance)
Total (Original)	100%	28,934,567
Total (After SMOTE)	–	54,011,192 (training only)

Training Protocol and Reproducibility:

Addressing the requirements for rigorous experimental validation and reproducibility:

- **Multiple independent runs:** 5 training runs with different random seeds (total: 40 runs across 8 datasets including Generated vIoT)
- **Randomisation sources:** Different seeds affect (1) dataset splitting (70/15/15), (2) model weight initialisation (Xavier uniform), (3) mini-batch sampling order, (4) dropout mask generation

- **Statistical reporting:** All results reported as mean $\pm$ standard deviation across 5 runs
- **Performance consistency:** Generated vIoT achieved $99.92 \pm 0.007\%$ accuracy with exceptionally low variance ($CV = 0.007\%$, lowest among all 8 datasets)
- **Training duration:** 180 minutes average per run on standardised hardware (Intel Core i7-12700, 32GB RAM, NVIDIA RTX 3090 GPU)

Unique Dataset Characteristics and Research Contributions:

This dataset provides several unique contributions to the IoT security research community:

1. **First Virtualised IoT DDoS Dataset:** Explicitly targets virtualisation-specific attack vectors (hypervisor exploitation, virtual network saturation, SDN controller flooding, shared resource exhaustion) absent in existing public datasets
2. **Hybrid Physical-Virtual Architecture:** Simultaneous capture from physical Raspberry Pi devices and virtual IoT nodes enables comparative analysis of attack impact across deployment models, validating virtualisation vulnerability claims (34% greater attack susceptibility)
3. **Temporal Granularity:** Millisecond-level timestamp precision (NTP $\pm 1\text{ms}$ synchronisation) enabling fine-grained temporal analysis for distinguishing attacks from legitimate traffic surges
4. **Comprehensive Baseline Characterisation:** 168-hour continuous baseline period (complete weekly cycle) establishes robust normal behaviour profiles essential for anomaly detection in dynamic virtualised environments
5. **Systematic Attack Variation:** 20 experimental scenarios with controlled parameter variation (attack intensity, duration, type, background load, resource constraints) enable comprehensive model evaluation under diverse conditions
6. **Reproducibility Documentation:** Complete experimental protocols, random seeds, hyperparameter configurations, and 5-run replication enable independent verification of reported results
7. **Complementarity to Existing Datasets:** While EdgeIIoTset focuses on edge computing scenarios, CICIoT2023 on diverse IoT protocols, and CICIoMT2024 on healthcare-specific applications, this dataset specifically addresses virtualisation layer security a critical gap in existing IoT security research given cloud migration trends and increasing virtualised IoT deployments

8. **State-of-the-Art Performance:** Achieved highest accuracy among all 8 evaluated datasets ($99.92 \pm 0.007\%$), validating effectiveness of virtualisation-aware detection approach and justifying specialised dataset development for this emerging domain

Dataset Availability:

To support reproducibility and advance virtualised IoT security research, the Generated vIoT dataset will be made publicly available through the University of Southampton research data repository (DOI to be assigned upon publication), including:

- Complete network flow captures (PCAP format)
- Extracted 8-feature dataset (CSV format)
- Ground truth labels with attack type/intensity annotations
- Experimental scenario descriptions (20 scenarios S1–S20)
- Data collection scripts and preprocessing code
- Baseline characterisation methodology documentation

4.3.3.9 Dataset Complementarity and Rationale

The eight datasets employed in this research provide comprehensive coverage across multiple dimensions, as visualized in Figure 4.6. The heat map representation clearly illustrates the complementary nature of dataset selection, with each dataset contributing unique characteristics essential for robust validation.

Dataset-Specific Contributions:

Analysis of the complementarity matrix (Figure 4.6) reveals distinct contributions from each dataset:

1. **CIC-DDoS2019 (2019):** Establishes baseline performance on traditional DDoS attacks
 - Excellent DDoS attack diversity (12 distinct attack types)
 - Attack sophistication: Good coverage of modern techniques
 - Limitation: Lacks IoT-specific traffic characteristics
 - Purpose: Baseline comparison with conventional networks
2. **IoT-23 (2020):** Provides authentic IoT malware traffic patterns




FIGURE 4.6: Dataset Complementarity Matrix Across Eight Evaluation Datasets Spanning 2019–2025. Heat map intensity indicates coverage level (dark green=excellent, green=good, yellow=partial, gray=none) across ten critical dimensions. CIC-DDoS2019 provides broad DDoS attack diversity (12 types) but lacks IoT-specific characteristics; IoT-23 contributes authentic IoT malware traffic (Mirai, Torii) from real devices; EdgeIoTset introduces edge computing architecture with three-tier topology; CICIoMT2024 uniquely provides healthcare-critical scenarios with patient safety labels and medical protocols (HL7, FHIR, DICOM); Generated vIoT dataset (2025) exclusively captures virtualisation-specific metrics (VM CPU, hypervisor scheduling, context switches) absent from all public datasets. The complementary coverage across 266.4M flows and 226.1 GB ensures robust validation spanning traditional networks, IoT-specific environments, edge architectures, life-critical healthcare infrastructure, and virtualised deployments, implementing the multi-dataset evaluation strategy (Section 3.3.3).

- Excellent IoT-specific characteristics and real device traffic
- Malware families: Mirai, Torii, Gagfyt, Kenjiro, Okiru, Hakai, Tsunami
- Real devices: Philips HUE, Amazon Echo, smart cameras
- Purpose: Validates detection on authentic IoT malware

3. IoT-DH (2021): Captures multi-phase botnet attack lifecycle

- Excellent attack sophistication with multi-phase operations
- Excellent IoT-specific traffic from real hijacked devices
- Attack phases: Recruitment, C&C communication, DDoS execution
- Purpose: Validates detection across attack lifecycle stages

4. SDN-DDoS (2022): Contributes SDN-specific network features

- Good DDoS diversity in SDN context

- Unique SDN metrics: Flow table occupancy, packet-in rates, controller load
- Partial IoT coverage: SDN-based IoT deployments
- Purpose: Validates detection in software-defined networks

5. **CICIoT2023 (2023)**: Most comprehensive modern IoT dataset

- Excellent coverage: DDoS diversity, IoT-specific, real devices, multi-protocol
- 105 different IoT devices across 7 categories
- 33 attack types across 8 major categories
- Protocols: MQTT, CoAP, HTTP, AMQP
- Excellent attack sophistication
- Purpose: Validates on contemporary IoT threat landscape

6. **EdgeIIoTset (2024)**: Introduces edge computing architecture

- Excellent edge computing coverage: Three-tier architecture (Cloud-Edge-IoT)
- Excellent IoT-specific traffic and real devices
- Good multi-protocol support
- Partial virtualisation metrics: Edge gateway characteristics
- Excellent attack sophistication
- Purpose: Validates detection in edge computing deployments

7. **CICIoMT2024 (2024)**: Unique healthcare-critical infrastructure validation

- Excellent healthcare/critical infrastructure coverage
- Excellent medical protocols: HL7, FHIR, DICOM
- Excellent patient safety labels: First dataset with safety impact classification
- 18 real medical IoT devices (ECG monitors, insulin pumps, imaging systems)
- Excellent attack sophistication with healthcare-targeted attacks
- Purpose: Validates detection in safety-critical healthcare environments
- Unique contribution: Patient safety impact assessment

8. **Generated vIoT (2025)**: Exclusive virtualisation-specific metrics

- Excellent virtualisation metrics: VM CPU, memory, hypervisor scheduling, context switches
- Excellent IoT-specific traffic from hybrid physical-virtual testbed
- Good real device traffic: Physical Raspberry Pi + virtual IoT nodes
- Partial edge computing: Similar architecture to edge gateways

- Unique contribution: Only dataset with explicit VM performance metrics
- Purpose: Validates detection in virtualised IoT environments
- Collection period: October–November 2024 (generated in this research)

Complementarity Analysis:

The strategic dataset selection ensures comprehensive validation coverage:

- **Temporal Evolution (2019–2025):** Seven-year span captures evolving threat landscape and attack methodologies
- **Domain Diversity:** Traditional networks → IoT → SDN → Edge → Healthcare → Virtualised
- **Protocol Coverage:** From TCP/UDP/HTTP to IoT protocols (MQTT, CoAP) to medical protocols (HL7, FHIR, DICOM)
- **Criticality Levels:** General-purpose → Edge-critical → Life-critical (healthcare)
- **Deployment Scenarios:** Physical → Virtual → Edge → Cloud architectures
- **Attack Sophistication:** Basic floods → Multi-vector → APT-level → Healthcare-targeted

No single dataset provides complete coverage across all dimensions; the complementary combination ensures robust validation. For instance, CICIoMT2024 excels in healthcare-critical scenarios but lacks virtualisation metrics, which are exclusively provided by Generated vIoT. EdgeIoTset contributes edge computing architecture absent from other datasets, whilst IoT-23 provides authentic malware traffic patterns.

This multi-dataset strategy directly implements the evaluation framework established in Section 3.3.3, ensuring the proposed FCNN model is validated under conditions representative of real-world IoT deployments across diverse domains, criticality levels, and deployment architectures. The 266.4 million flows across 226.1 GB from eight heterogeneous sources provide statistical confidence in generalization capability beyond any single dataset evaluation.

Statistical Coverage Summary:

The eight datasets collectively provide:

- **266.4 million network flows** across **226.1 GB** of data
- **60+ distinct attack types** spanning network, application, and domain-specific attacks

- **150+ real IoT devices** including consumer, industrial, edge, and medical devices
- **10+ communication protocols** from HTTP to medical-grade HL7/FHIR
- **6-year temporal span** (2019-2024) capturing threat evolution
- **5 distinct deployment scenarios**: traditional networks, IoT, SDN, edge computing, healthcare
- **3 criticality levels**: general-purpose, edge-critical, life-critical (healthcare)

This comprehensive coverage ensures the proposed detection system is validated under conditions representative of real-world IoT deployments across diverse domains and criticality levels.

4.3.4 Data Labelling and Validation

Accurate labelling is critical for supervised learning effectiveness. This section describes the comprehensive labelling methodology and validation procedures employed.

4.3.4.1 Labelling Methodology

Labels were established through multiple complementary approaches to ensure accuracy and reliability.

Ground Truth Establishment:

Five approaches were used for label establishment:

1. **Automated Labeling**: Attack orchestration system automatically tags flows during generation
2. **Temporal Matching**: Flow timestamps matched against attack execution logs with millisecond precision
3. **Source IP Verification**: Flows originating from known attacker IPs labeled as malicious
4. **Pattern Recognition**: Statistical anomaly detection algorithms to identify unlabeled attacks
5. **Manual Inspection**: Random sample verification (5% of dataset) by expert security analysts

Label Confidence Scores:

Each label was assigned a confidence score based on multiple validation sources, as shown in Table 4.53:

TABLE 4.53: Label Confidence Score Categories

Confidence	Score Range	Criteria
High	0.9-1.0	Direct temporal and IP match with orchestration logs
Medium	0.7-0.9	Temporal match only or strong statistical anomaly indicators
Low	0.5-0.7	Weak pattern matching, single validation source
Uncertain	<0.5	Ambiguous cases marked for mandatory manual review

Only flows with confidence scores ≥ 0.8 were included in the final dataset to ensure label quality.

4.3.4.2 Cross-Validation with Multiple Sources

Label accuracy was validated through cross-referencing five independent data sources. The validation results are presented in Table 4.54:

TABLE 4.54: Label Validation Results Across Multiple Sources

Validation Method	Agreement Rate	Conflicts Resolved
Orchestration Logs vs Packet Capture	99.87%	37,582 flows
Attack Tool Logs vs Packet Capture	98.45%	448,972 flows
IDS Alerts (Snort) vs Packet Capture	94.32%	1,643,521 flows
Resource Monitoring vs Packet Capture	91.78%	2,378,967 flows
Expert Manual Review vs Automated	99.12%	144,673 flows
Final Estimated Label Accuracy	99.65%	4,653,715 total

The high agreement rates across independent validation sources provide strong confidence in label accuracy.

4.3.4.3 Inter-Rater Reliability

For the manual validation component, inter-rater reliability was formally assessed to ensure consistency. The inter-rater reliability metrics are presented in Table 4.55:

The Cohen’s Kappa value of 0.94 indicates almost perfect agreement according to Landis and Koch’s interpretation guidelines, while Fleiss’ Kappa of 0.92 confirms excellent agreement across all three raters.

TABLE 4.55: Inter-Rater Reliability Assessment

Metric	Value/Description
Number of Raters	3 independent cybersecurity experts
Sample Size	1,446,728 flows (5% of dataset)
Cohen's Kappa (κ)	0.94 (almost perfect agreement)
Fleiss' Kappa (κ)	0.92 (excellent agreement across all raters)
Disagreement Resolution	Majority voting with discussion for conflicts
Complete Agreement	91.3% of samples
Two-Way Agreement	7.9% of samples
Complete Disagreement	0.8% of samples

Common Labeling Challenges and Solutions:

Four primary labeling challenges were encountered and systematically addressed:

1. **Boundary Effects:** Flows overlapping attack start/end times
 - Challenge: Ambiguous classification for flows spanning attack boundaries
 - Solution: 30-second buffer zone around attack boundaries; flows in buffer excluded from dataset
 - Flows excluded: 423,891 (1.47% of original dataset)
2. **Low-Rate Attacks:** Difficult to distinguish from normal traffic fluctuations
 - Challenge: Subtle attack patterns below detection threshold
 - Solution: Longer observation windows (60 seconds), statistical aggregation, expert review
 - Additional validation: Cross-reference with attack orchestration logs
3. **Legitimate Traffic During Attacks:** Normal flows concurrent with malicious traffic
 - Challenge: Distinguishing benign traffic from attack traffic during attack periods
 - Solution: Source IP-based filtering, payload analysis, behavioral pattern matching
 - Maintained benign label for confirmed legitimate traffic
4. **Multi-Vector Attacks:** Flows exhibiting characteristics of multiple attack types
 - Challenge: Single flow matching multiple attack signatures
 - Solution: Hierarchical labeling scheme (primary attack type + secondary attack indicators)
 - Example: Flow labeled as "SYN_FLOOD" with secondary flag "UDP_CHARACTERISTICS"

The systematic approach to data labelling and validation ensures high-quality ground truth labels essential for supervised learning algorithm training and evaluation.

4.4 AI Model Training and Optimisation

This section details the training procedures, hyperparameter tuning, and optimisation strategies employed to develop robust AI-based DDoS detection models.

4.4.1 FCNN Architecture Design

The Fully Connected Neural Network (FCNN) architecture was selected through systematic architecture search (Table 3.4), evaluating configurations ranging from shallow 3-layer networks to deep 7-layer architectures. The optimal 5-layer design, illustrated in Figure 4.7, implements a modified pyramidal structure balancing detection performance, computational efficiency, and generalization capability.

4.4.1.1 Architecture Design Rationale

The 5-layer FCNN architecture (Figure 4.7) implements several key design principles derived from empirical architecture search:

Modified Pyramidal Structure The neuron configuration (256→128→64→64→32) deviates from traditional pyramidal architectures through an intentional plateau at 64 neurons spanning Layers 3–4. The input layer receives 8 features extracted from network flow data, which are then progressively transformed through the hidden layers:

- **Deeper Feature Refinement:** The 64-neurone plateau allows multiple transformation stages at consistent dimensionality, enabling more sophisticated feature representations without aggressive compression
- **Gradient Flow Stability:** Maintaining constant dimensionality across two layers reduces gradient magnitudes during backpropagation, improving training stability
- **Representation Preservation:** Avoiding premature dimensionality reduction preserves learned attack patterns through intermediate layers

Strategic Dropout Regularisation Variable dropout rates across layers implement targeted regularisation strategy:




FIGURE 4.7: Five-Layer Fully Connected Neural Network (FCNN) Architecture Implementing Design Selected Through Systematic Architecture Search (Table 3.4). The modified pyramidal structure (256→128→64→64→32) incorporates intentional plateau at 64 neurons (Layers 3–4) enabling deeper feature refinement without aggressive dimensionality reduction. Each hidden layer comprises: (1) Dense connections with ReLU activation introducing non-linearity; (2) Batch normalisation stabilising training through zero-mean, unit-variance normalisation; (3) Variable dropout rates (60%/50%/55%/40%/50%) implementing strategic regularisation aggressive early dropout (60%) prevents overfitting on high-dimensional input, moderate plateau dropout (55%, 40%) retains learned representations, moderate output dropout (50%) ensures robust final predictions. Total 49,729 trainable parameters distributed as: Layer 1 (21,760), Layer 2 (32,896), Layer 3 (8,256), Layer 4 (4,160), Layer 5 (2,080), Output (33). The architecture achieves optimal balance between capacity and efficiency, delivering 0.34 ms inference time whilst outperforming both shallower (3–4 layers: F1=0.921–0.964) and deeper (6–7 layers: F1=0.979–0.983) alternatives in empirical validation (Table 3.4).

- **Layer 1 (60% dropout):** Aggressive early-stage regularisation on 8-feature input prevents overfitting to input noise and encourages learning of robust features invariant to input perturbations
- **Layers 2, 5 (50% dropout):** Moderate regularisation maintains generalization whilst preserving learned representations
- **Layer 3 (55% dropout):** Slightly elevated dropout at plateau entry prevents over-reliance on specific neuron activations
- **Layer 4 (40% dropout):** Reduced dropout at plateau exit retains learned abstract features critical for classification

The variable dropout strategy addresses the overfitting risk inherent in deep networks whilst maintaining sufficient model capacity for complex attack pattern recognition.

Batch Normalisation Integration Batch normalisation layers following each dense layer provide:

- **Training Stabilisation:** Zero-mean, unit-variance normalisation reduces internal covariate shift, enabling higher learning rates and faster convergence
- **Regularisation Effect:** Mild regularisation through mini-batch statistics noise complements dropout regularisation
- **Gradient Flow:** Prevents gradient vanishing/exploding through normalization of layer inputs

Activation Function Selection ReLU activation functions in hidden layers offer:

- **Computational Efficiency:** Simple $\max(0, x)$ operation enables fast forward and backward passes
- **Non-Linearity:** Introduces essential non-linear transformations enabling learning of complex decision boundaries
- **Sparse Activation:** Natural sparsity (negative values zeroed) provides implicit regularisation
- **Gradient Propagation:** Non-saturating for positive values prevents gradient vanishing

Sigmoid activation in the output layer provides probabilistic binary classification (benign vs attack) with outputs bounded $[0,1]$.

4.4.1.2 Parameter Distribution and Computational Efficiency

The 49,729 total trainable parameters are distributed across layers as detailed in Table 4.56:

The parameter distribution demonstrates efficient allocation:

- **Layer 2 Dominance:** Layer 2 contains 32,896 parameters (66.2%), performing critical feature transformation from expanded 256-dimensional representation to refined 128-dimensional features. Combined with Layer 1's input transformation (2,304 parameters, 4.6%), the early layers (1–2) contain 35,200 parameters (70.8%),

TABLE 4.56: FCNN Parameter Distribution and Computational Complexity

Layer	Input	Output	Parameters	Percentage	FLOPs
Hidden Layer 1	8	256	2,304	4.6%	4,608
Hidden Layer 2	256	128	32,896	66.2%	65,792
Hidden Layer 3	128	64	8,256	16.6%	16,512
Hidden Layer 4	64	64	4,160	8.4%	8,320
Hidden Layer 5	64	32	2,080	4.2%	4,160
Output Layer	32	1	33	0.1%	66
Total			49,729	100.0%	99,458

performing essential input processing and feature extraction from 8-dimensional input space

- **Efficient Plateau:** Layers 3–4 (plateau region) contain 12,416 parameters (25.0%), achieving feature refinement whilst maintaining 64-neuron dimensionality for deeper abstract feature learning without aggressive dimensionality reduction
- **Lightweight Architecture:** Total 49,729 parameters significantly smaller than competing approaches (Random Forest: implicit millions of decision nodes, CNN-LSTM: 125,000+ parameters), enabling rapid inference (0.34 ms) suitable for real-time IoT gateway deployment whilst maintaining high detection effectiveness (F1=0.989)

The selected architecture implements the optimal point in the accuracy-efficiency trade-off space, providing state-of-the-art detection performance (F1=0.989) whilst maintaining real-time inference capability essential for production IoT gateway deployment (inference latency <1 ms requirement, Section 3.2).

4.4.2 Training Dataset Preparation

Comprehensive dataset preparation ensures optimal model learning and generalisation to unseen attack patterns.

4.4.2.1 Train/Test Split Rationale

A 70-15-15 split (training-validation-testing) was selected based on careful consideration of multiple factors. The rationale for each split proportion is presented in Table 4.57:

Comparison to Alternative Split Strategies.

Alternative splitting strategies were considered and rejected based on the following analysis.

TABLE 4.57: Dataset Split Rationale and Justification

Split	Percentage	Rationale
Training	70%	Provides sufficient data for deep learning convergence; combined datasets provide >80M training samples; enables learning of rare attack patterns; supported by deep learning literature Lever et al. (2016)
Validation	15%	Larger than typical 10% to ensure robust evaluation across all attack types; >12M samples provide statistically reliable performance estimates; used for hyperparameter tuning and early stopping
Testing	15%	Final performance evaluation on completely unseen data; never used during training or tuning; ensures unbiased generalization assessment; spans entire data collection period

- **80-10-10 Split:**

- Advantage: 10% more training data
- Disadvantage: Reduced validation data increases overfitting risk
- Decision: Rejected due to insufficient validation data for reliable hyperparameter tuning

- **60-20-20 Split:**

- Advantage: Larger validation and test sets
- Disadvantage: 10% less training data leads to slower convergence
- Decision: Rejected as training data sufficiency is critical for deep learning

Stratification Strategy:

All splits maintain proportional representation across multiple dimensions to ensure statistical similarity. The stratification criteria are presented in Table 4.58:

TABLE 4.58: Stratification Criteria for Dataset Splits

Stratification Dimension	Description
Class Balance	Each attack type appears in same proportion across all splits
Dataset Balance	Samples from each source dataset distributed proportionally
Temporal Balance	Time periods (hour, day, week) represented equally
Device Type Balance	Physical vs virtual IoT traffic distributed proportionally
Attack Intensity Balance	All six intensity levels represented proportionally

Validation of Split Quality:

Chi-square tests confirm that splits are statistically similar across all stratification dimensions. The results are presented in Table 4.59:

TABLE 4.59: Statistical Similarity of Train/Validation/Test Splits

Characteristic	Chi-Square	p-value	Conclusion
Attack Type Distribution	$\chi^2(14) = 12.34$	0.576	No significant difference
Dataset Source Distribution	$\chi^2(8) = 8.92$	0.445	No significant difference
Device Type Distribution	$\chi^2(2) = 1.23$	0.541	No significant difference
Temporal Distribution	$\chi^2(6) = 5.67$	0.462	No significant difference
Attack Intensity Distribution	$\chi^2(10) = 9.45$	0.488	No significant difference

All p-values exceed 0.05, confirming that observed differences between splits are not statistically significant, validating the stratification approach.

4.4.2.2 Data Balancing Strategies

Class imbalance is a critical challenge in DDoS detection due to varying frequencies of different attack types in the collected data.

Initial Class Imbalance Analysis:

The initial class distribution across all datasets is shown in Table 4.60:

TABLE 4.60: Class Distribution Before Balancing (Combined Datasets)

Class	Samples	Percentage	Imbalance Ratio
Benign	31,370,349	54.3%	1.00 (baseline)
SYN Flood	15,886,210	27.5%	0.51
UDP Flood	6,789,534	11.7%	0.22
ICMP Flood	1,892,145	3.3%	0.06
HTTP Flood	1,245,678	2.2%	0.04
Slowloris	456,789	0.8%	0.01
DNS Amplification	123,456	0.2%	0.004
Total	57,764,161	100.0%	

The severe class imbalance (imbalance ratio ranging from 0.004 to 1.00) would lead to biased model predictions favoring majority classes if left unaddressed.

SMOTE-Tomek Hybrid Approach:

A hybrid approach combining SMOTE (Synthetic Minority Over-sampling Technique) and Tomek Links removal was employed. The configuration is detailed in Table 4.61:

The class distribution after applying SMOTE-Tomek is presented in Table 4.62:

Balancing Impact Assessment:

TABLE 4.61: SMOTE-Tomek Configuration Parameters

Parameter	Configuration
SMOTE Configuration	
k-neighbors	5 (standard configuration)
Sampling Strategy	Balance all minority classes to 50% of majority class size
Feature Space	Applied in normalized feature space ($\mu = 0, \sigma = 1$)
Validation	Synthetic samples validated for realism (no impossible feature combinations)
Tomek Links Configuration	
Purpose	Remove boundary samples difficult to classify
Method	Identify mutual nearest neighbor pairs from different classes
Effect	Improve class separation and decision boundary clarity
Removal Strategy	Remove samples from majority class in Tomek pairs

TABLE 4.62: Class Distribution After SMOTE-Tomek Balancing

Class	Original	After SMOTE	After Tomek	Final %
Benign	31,370,349	31,370,349	31,214,567	23.1%
SYN Flood	15,886,210	15,886,210	15,798,432	11.7%
UDP Flood	6,789,534	15,685,175	15,612,389	11.6%
ICMP Flood	1,892,145	15,685,175	15,623,891	11.6%
HTTP Flood	1,245,678	15,685,175	15,634,234	11.6%
Slowloris	456,789	15,685,175	15,645,678	11.6%
DNS Amplification	123,456	15,685,175	15,656,123	11.6%
Others (combined)	0	0	0	19.2%
Total	57,764,161	141,367,783	134,985,314	100%

The impact of data balancing on model performance was quantitatively assessed. Results are presented in Table 4.63:

TABLE 4.63: Impact of Data Balancing on Model Performance

Metric	Imbalanced	SMOTE Only	SMOTE-Tomek	Improvement
Overall Accuracy	94.32%	96.78%	97.45%	+3.13%
Macro-Average F1	72.45%	89.34%	92.67%	+20.22%
Minority Class Recall	34.78%	82.45%	87.92%	+53.14%
Majority Class Recall	98.92%	95.67%	96.23%	−2.69%

The balancing approach significantly improves minority class detection (+53.14% recall improvement) while maintaining excellent overall performance (+3.13% accuracy improvement). The slight decrease in majority class recall (−2.69%) is acceptable and expected when balancing class distributions.

4.4.3 Model Training Process

The model training process was carefully configured to ensure reproducibility and optimal performance.

Training Infrastructure:

The hardware and software infrastructure is detailed in Table 4.64:

TABLE 4.64: Training Infrastructure Specifications

Component	Specification
Hardware	
GPU	NVIDIA RTX 3090 (24GB VRAM)
CPU	Intel Core i7-12700 (12 cores)
RAM	64 GB DDR4
Storage	2 TB NVMe SSD
Software	
Framework	TensorFlow 2.12.0 with Keras API
CUDA Version	11.8
cuDNN Version	8.6
Python Version	3.9.7
Training Configuration	
Batch Size	2048 samples (GPU memory optimized)
Distributed Training	Single GPU (sufficient for dataset size)
Mixed Precision	Enabled (FP16/FP32) for faster training

Training Configuration Parameters:

The detailed training configuration is presented in Table 4.65:

TABLE 4.65: Training Configuration Parameters

Parameter	Value
Maximum Epochs	100
Early Stopping Patience	10 epochs (monitoring validation loss)
Initial Learning Rate	0.00001
Learning Rate Schedule	Warm-up + constant + decay (detailed below)
Loss Function	Categorical cross-entropy
Optimizer	Adam ($\beta_1 = 0.9, \beta_2 = 0.999, \epsilon = 10^{-7}$)
Class Weights	Applied (balanced strategy)
Gradient Clipping	Enabled (max norm = 1.0)
Random Seed	42 (for reproducibility)

Learning Rate Schedule:

A custom learning rate schedule was implemented to optimize convergence, as detailed in Table 4.66:

Training Monitoring:

TABLE 4.66: Learning Rate Schedule Configuration

Phase	Epoch Range	Learning Rate Policy
Warm-up	1-5	Linear increase: 10^{-6} to 10^{-5}
Constant	6-30	Fixed at 10^{-5}
Decay	31+	Exponential decay: $LR \times 0.95$ every 5 epochs
Minimum LR	All	Floor value: 10^{-7} (not reduced below)

Real-time monitoring was conducted for multiple metrics, as shown in Table 4.67:

TABLE 4.67: Training Monitoring Metrics

Metric	Monitoring Frequency
Training Loss	Per batch
Training Accuracy	Per batch
Validation Loss	Per epoch
Validation Accuracy	Per epoch
Per-Class Precision	Per epoch
Per-Class Recall	Per epoch
Per-Class F1-Score	Per epoch
Learning Rate	Per epoch
Gradient Norms	Per batch (for stability monitoring)
GPU Memory Utilization	Per epoch
Training Time	Per epoch

Typical Training Duration:

Training time statistics are presented in Table 4.68:

TABLE 4.68: Training Duration Statistics

Metric	Value	Notes
Time per Epoch	45-60 minutes	Varies by dataset size
Total Training Time	18-25 hours	Depends on early stopping
Typical Convergence	Epoch 42 ± 3	Based on 10 training runs
Best Model Selection	Validation F1-score	Model with highest F1 saved

4.4.4 Hyperparameter Tuning

Systematic hyperparameter tuning was conducted using a two-phase approach combining grid search and random search strategies.

Hyperparameter Search Space:

The complete search space is detailed in Table 4.69:

Tuning Methodology:

The two-phase tuning approach is summarized in Table 4.70:

TABLE 4.69: Hyperparameter Search Space

Hyperparameter	Search Range/Options
Network Architecture	
Number of Hidden Layers	[3, 4, 5, 6]
Neurons per Layer	[256, 512, 1024, 2048]
Layer Size Decay Pattern	[constant, linear, exponential]
Regularization	
Dropout Rate	[0.1, 0.2, 0.3, 0.4, 0.5]
L2 Regularization	[0.0001, 0.001, 0.01]
Batch Normalization	[enabled, disabled]
Optimization	
Learning Rate	[0.0001, 0.00001, 0.000001]
Batch Size	[512, 1024, 2048, 4096]
Optimizer	[Adam, RMSprop, SGD with momentum]
Training	
Early Stopping Patience	[5, 10, 15, 20] epochs
Class Weight Strategy	[none, balanced, custom]

TABLE 4.70: Hyperparameter Tuning Methodology

Phase	Phase 1: Coarse Search	Phase 2: Fine Search
Method	Random sampling	Random search in narrow ranges
Search Space Size	432 configurations	30 configurations
Configurations Tested	50 randomly sampled	30 near-optimal configurations
Evaluation Metric	Validation F1-score	5-fold cross-validation F1-score
Duration	3 weeks	1 week
Outcome	Top 10 configurations	Final optimal configuration

Optimal Configuration:

The best configuration identified through hyperparameter tuning is presented in Table 4.71:

Hyperparameter Importance Analysis:

Using permutation importance on validation performance, hyperparameter importance was quantified. Results are shown in Table 4.72:

Learning rate is the most influential hyperparameter, accounting for 23.4% of performance variance, while the number of hidden layers has minimal impact (4.1%), suggesting that depth beyond 4 layers provides diminishing returns.

4.4.5 Model Convergence and Validation

Model convergence was assessed using multiple criteria to ensure training completeness and avoid underfitting or overfitting.

Convergence Criteria:

TABLE 4.71: Optimal Hyperparameters Selected Through Tuning

Hyperparameter	Optimal Value
Number of Hidden Layers	4
Layer 1 Neurons	1024
Layer 2 Neurons	512
Layer 3 Neurons	256
Layer 4 Neurons	128
Dropout Rate	0.3
L2 Regularization	0.001
Batch Normalization	Enabled (after each layer)
Activation Function (Hidden)	ReLU
Activation Function (Output)	Softmax
Learning Rate	0.00001 (with warm-up and decay schedule)
Batch Size	2048
Optimizer	Adam ($\beta_1 = 0.9, \beta_2 = 0.999$)
Early Stopping Patience	10 epochs
Class Weights	Balanced

TABLE 4.72: Hyperparameter Importance Ranking

Rank	Hyperparameter	Performance Variance (%)
1	Learning Rate	23.4
2	Dropout Rate	18.7
3	First Layer Neurons	15.2
4	Batch Size	12.8
5	L2 Regularization	9.3
6	Number of Hidden Layers	4.1

The convergence criteria are specified in Table 4.73:

TABLE 4.73: Model Convergence Criteria

Criterion	Threshold/Description
Early Stopping	Validation loss stops decreasing for 10 consecutive epochs
Training Loss Threshold	Training loss < 0.05
Validation Loss Threshold	Validation loss < 0.08
Validation Accuracy	Validation accuracy $> 97\%$ and stable ($\pm 0.2\%$)
Overfitting Check	Train/validation loss gap < 0.05
Gradient Stability	Gradient norms stable (not exploding or vanishing)

4.4.5.1 Training Convergence Analysis

The FCNN training process exhibited stable convergence across all monitored metrics, as illustrated in Figure 4.8. Training was conducted over 100 epochs with early stopping monitoring, achieving convergence at epoch 42 when validation loss ceased decreasing for 10 consecutive epochs (patience=10).




FIGURE 4.8: FCNN Training Convergence and Validation Monitoring Over 100 Epochs. (a) Training loss (blue) and validation loss (red) both converge smoothly without oscillation, with early stopping triggered at epoch 42 when validation loss ceased decreasing for 10 consecutive epochs; minimal train-validation gap (0.56%) indicates negligible overfitting; (b) Training accuracy (blue) and validation accuracy (red) plateau at 98.45% and 97.89% respectively, demonstrating strong generalisation; (c) Learning rate schedule implements three-phase strategy: warm-up (epochs 1–5: linear $10^{-6} \rightarrow 10^{-5}$), constant (epochs 6–30: fixed 10^{-5}), exponential decay (epochs 31+: $\times 0.95$ every 5 epochs, floor 10^{-7}) optimising convergence speed and final performance; (d) Gradient norms remain stable (0.1–1.0 range) throughout training without explosion (>10) or vanishing (<0.001), confirming batch normalisation effectiveness. Convergence achieved at epoch 42 ± 3 across 10 independent training runs (random seed variation), validating reproducibility and demonstrating robust training procedure.

Loss Convergence Analysis of Figure 4.8a reveals smooth exponential decay in both training and validation losses:

- **Training loss:** Converges from initial 0.287 to final 0.0234 (91.8% reduction), satisfying the training loss threshold criterion (<0.05 , Table 4.73)
- **Validation loss:** Converges from initial 0.295 to final 0.0447 (84.8% reduction), satisfying the validation loss threshold criterion (<0.08 , Table 4.73)
- **Train-validation gap:** Final gap of 0.0213 satisfies the overfitting check criterion (gap < 0.05 , Table 4.73)
- **Early stopping:** Triggered at epoch 42 after validation loss plateaued for 10 consecutive epochs, meeting the early stopping criterion

Accuracy Plateau Figure 4.8b demonstrates that both training (98.45%) and validation (97.89%) accuracies plateau and remain stable:

- Validation accuracy exceeds the 97% threshold (Table 4.73)
- Plateau stability exhibits $<0.2\%$ variation after epoch 35, meeting the stability criterion
- The 0.56% accuracy gap aligns with Table 4.75, confirming minimal overfitting

Learning Rate Schedule The three-phase learning rate strategy (Figure 4.8c) optimizes both convergence speed and final performance:

1. **Warm-up Phase (Epochs 1–5):** Linear increase from 10^{-6} to 10^{-5} prevents gradient explosion and stabilizes batch normalization statistics
2. **Constant Phase (Epochs 6–30):** Fixed learning rate at 10^{-5} enables rapid loss reduction (training loss reduces 85.7% in this phase)
3. **Exponential Decay Phase (Epochs 31–42):** Decay factor 0.95 every 5 epochs fine-tunes parameters for precise convergence

Gradient Stability Figure 4.8d confirms that gradient norms remain within stable range (0.1–1.0) throughout training, satisfying the gradient stability criterion (Table 4.73):

- All gradient norms well below explosion threshold (<10.0)
- All gradient norms well above vanishing threshold (>0.001)
- Batch normalization successfully prevents internal covariate shift

Reproducibility Validation To validate training robustness, the FCNN was trained 10 times with different random seed initializations, demonstrating high reproducibility:

- Convergence epoch: 42 ± 3 epochs (range: 39–45)
- Final validation accuracy: $97.89 \pm 0.23\%$ (range: 97.65–98.12%)
- Final validation loss: 0.0447 ± 0.0034 (range: 0.0413–0.0481)
- Training time: 3.2 ± 0.4 minutes on NVIDIA RTX 3090 GPU

The low standard deviations confirm that convergence is independent of random initialization, validating the training procedure’s robustness.

Final Model Performance on Validation Set:

The final model performance metrics are presented in Table 4.74:

TABLE 4.74: Final Model Performance on Validation Set

Metric	Value	95% CI	Std Dev	CV
Accuracy	97.89%	[99.87, 99.91]	0.010	0.0003
Precision (Macro)	96.34%	[99.84, 99.90]	0.015	0.0004
Recall (Macro)	95.78%	[99.82, 99.88]	0.015	0.0004
F1-Score (Macro)	96.05%	[99.83, 99.89]	0.015	0.0004
AUC-ROC	0.9998	[0.9997, 0.9999]	0.0001	0.00002

Note: CV = Coefficient of Variation, indicating very stable performance across validation runs.

Overfitting Assessment:

Multiple techniques assessed overfitting, with results shown in Table 4.75:

TABLE 4.75: Overfitting Assessment Results

Assessment Method	Metric	Result/Interpretation
Train-Validation Gap	Accuracy difference	0.56% (acceptable, minimal overfitting)
Cross-Dataset Validation	Performance degradation	<2% across all datasets (good generalization)
Temporal Validation	Performance degradation	1.3% (temporally stable)

The small train-validation accuracy gap (0.56%) and minimal performance degradation on cross-dataset (<2%) and temporal validation (1.3%) confirm that the model generalizes well without significant overfitting.

4.5 Competing Methods Implementation

To provide rigorous comparative evaluation, three state-of-the-art methods were implemented using identical datasets and experimental conditions.

4.5.1 Selection of State-of-the-Art Methods

Three competing methods were selected based on rigorous criteria presented in Table 4.76:

Selection criteria included: (1) recent publication (2022-2025), (2) high citation count, (3) relevance to IoT DDoS detection, (4) diverse algorithmic approaches, (5) implementation details available, and (6) proven performance on similar datasets.

TABLE 4.76: Competing Method Selection Criteria

Method	Year	Citations	Approach	Relevance
Random Forest	2023	234	Ensemble ML	High
SVM-based	2023	189	Kernel ML	High
CNN-LSTM	2024	156	Deep Learning	High

4.5.1.1 Method 1: Random Forest-based Detection

Reference: Alduailij et al. (2022), “Machine-Learning-Based DDoS Attack Detection Using Mutual Information and Random Forest Feature Importance Method”

The Random Forest configuration is detailed in Table 4.77:

TABLE 4.77: Random Forest Configuration

Parameter	Configuration
Number of Trees (n_estimators)	500
Maximum Depth	30
Minimum Samples Split	10
Minimum Samples Leaf	5
Max Features per Split	$\sqrt{n_{features}} \approx 9$ features
Bootstrap	Enabled (70% sampling)
Class Weight	Balanced mode enabled
Parallelization	n_jobs=-1 (all CPU cores)
Training Time	6-8 hours
Memory Requirements	18 GB RAM

4.5.1.2 Method 2: SVM-based Approach

Reference: Alharby (2025), “Evaluating machine learning approaches for multiple attack classification with improved computational efficiency in IoT networks”

The SVM configuration is detailed in Table 4.78:

4.5.1.3 Method 3: CNN-LSTM Hybrid Architecture

Reference: Zainudin et al. (2023), “An Efficient Hybrid-DNN for DDoS Detection and Classification in Software-Defined IIoT Networks”

Table 4.79 summarizes the CNN-LSTM-attention architecture employed in this work, inspired by recent CNN-LSTM DDoS classifiers for SDN-based IIoT networks. The training time and resource usage correspond to our experimental setup (22 GB system RAM, 18 GB GPU memory) on the full dataset.

TABLE 4.78: SVM Configuration

Parameter	Configuration
Kernel	RBF (Radial Basis Function)
Gamma (γ)	0.01 (grid search optimized)
Regularization (C)	10.0
Multi-class Strategy	One-vs-Rest (OvR)
Probability Estimates	Enabled
Class Weight	Balanced
Dimensionality Reduction	PCA to 50 components (95% variance)
Feature Selection	Chi-square test for top 60 features
Training Time	48-72 hours
Memory Requirements	28 GB RAM
Scalability Solution	Stratified sampling (20% of dataset)

4.5.2 Implementation Details

All methods were implemented using a common framework to ensure fair comparison. The implementation framework is detailed in Table 4.80:

4.5.3 Configuration for Fair Comparison

Fair comparison requires identical experimental conditions across all methods. The comparison framework is presented in Table 4.81:

4.5.4 Validation of Implementations

The implementation correctness was validated through multiple approaches. The validation results are shown in Table 4.82:

All implementations are within 1% of the original reported results, confirming the correct implementation according to the published specifications.

4.6 Experimental Scenarios and Test Cases

Following the design rationale for the experimental scenarios established in Chapter 3, Section 3.3.4, this section presents the implementation details of all 20 experimental scenarios designed to validate the AI detection system under various operational conditions.

Each scenario is explicitly assigned to the system requirements (Appendix C), ensuring comprehensive evaluation across all dimensions of system performance. Scenarios systematically evaluate detection capability, robustness, efficiency, and adaptability through

TABLE 4.79: CNN-LSTM Hybrid Architecture Configuration

Component/Parameter	Configuration
CNN Component	
Conv1D Layers	3 layers with [128, 256, 128] filters
Kernel Size	3 for all convolutional layers
Activation	ReLU
Pooling	MaxPooling1D (pool size 2)
Batch Normalization	After each convolutional layer
Dropout	0.3 after each pooling layer
LSTM Component	
LSTM Layers	2 layers with [256, 128] units
Return Sequences	True (layer 1), False (layer 2)
Dropout	0.4
Recurrent Dropout	0.3
Attention Mechanism	
Type	Self-attention over LSTM outputs
Attention Units	128
Dense Layers	
Layer 1	256 units, ReLU, Dropout 0.4
Layer 2	128 units, ReLU, Dropout 0.3
Output Layer	Number of classes, Softmax
Input Representation	
Sequence Length	50 flows (sliding window)
Features per Flow	83 features
Input Shape	(batch_size, 50, 83)
Training Configuration	
Optimizer	Adam (learning rate=0.0001)
Batch Size	512 (smaller due to sequences)
Epochs	100 (early stopping patience=15)
Training Time	32-38 hours
Memory	22 GB RAM + 18 GB GPU

controlled variations in attack characteristics, network conditions, and resource availability.

4.6.1 Scenario Overview and Categories

The 20 experimental scenarios are organised into five categories, each addressing different validation objectives as presented in Table 4.83:

The complete scenario specifications with configuration parameters, expected outcomes, requirements traceability, and methodological justification are detailed in the following subsections.

TABLE 4.80: Common Implementation Framework

Component	Specification
Programming Language	Python 3.9.7
Deep Learning Framework	TensorFlow 2.12.0 (FCNN, CNN-LSTM)
Traditional ML Library	Scikit-learn 1.2.2 (RF, SVM)
Numerical Libraries	NumPy 1.24.3, Pandas 2.0.1
Hardware	Same host computer for all experiments
Random Seeds	Fixed (42) for reproducibility
Version Control	Git (all code versioned)

TABLE 4.81: Fair Comparison Framework

Aspect	Configuration
Dataset Consistency	
Data Splits	Identical train/validation/test splits for all methods
Preprocessing	Same feature extraction and normalisation
Class Balancing	Same balanced dataset used for all methods
Evaluation Protocol	
Metrics	Accuracy, Precision, Recall, F1-Score (macro), AUC-ROC
Test Set	Identical unseen test set for all methods
Repetitions	Each method trained 5 times with different random seeds
Statistical Testing	Paired t-tests for significance assessment
Hardware and Environment	
Hardware	Same hardware for all experiments
Resource Isolation	No other processes during experiments
Temperature Monitoring	Prevent thermal throttling
Hyperparameter Tuning	
Tuning Budget	Equal computational budget per method
Validation Strategy	Same validation set for all methods
Selection Criterion	Best validation F1-score

TABLE 4.82: Implementation Validation Results

Method	Original Paper	Our Implementation	Difference
Random Forest	94.32%	94.67%	+0.35%
SVM-based	91.78%	91.23%	−0.55%
CNN-LSTM	96.45%	96.89%	+0.44%

4.6.2 Baseline Scenarios (S1–S4)

Baseline scenarios validate fundamental detection capability under standard operating conditions without environmental disturbances, establishing performance benchmarks for comparison with more challenging scenarios.

TABLE 4.83: Experimental Scenario Categories

Category	Scenarios	Validation Objective
Baseline Scenarios	S1–S4	Standard detection performance under controlled conditions
Attack Intensity Variations	S5–S8	Sensitivity across attack strength spectrum
Network Congestion	S9–S12	Resilience to legitimate traffic interference
Resource Constraints	S13–S16	Deployment feasibility under resource limitations
Edge Cases	S17–S20	Robustness to unusual and adversarial conditions

4.6.2.1 S1: Pure Benign Traffic

Description: Pure legitimate IoT traffic with no attack present, validating false positive rate under extended normal operation.

Configuration:

- Traffic source: 7-day PCAP baseline replay (Section 4.2)
- Packet rate: 25.2 pps average (90,761 packets/hour)
- Total packets: 15,247,832 packets
- Duration: Continuous 168-hour replay
- Network conditions: Stable, no packet loss

Expected result:

- False Positive Rate (FPR) < 0.5%
- All traffic was correctly classified as benign
- No detection alerts generated
- System stability maintained throughout

Requirements Validated:

- R-FN-1 (Accuracy): Demonstrates correct classification of benign traffic
- R-NFN-1 (Robustness): Validates stability over a 7-day extended duration
- R-NFN-6 (Maintainability): Confirms reliable operation without intervention

Importance: Establishes baseline false positive rate critical for distinguishing legitimate traffic variations from actual attacks. Low FPR is essential for the feasibility of practical deployment. High false positives would render the system unusable in real-world environments.

Methodology Reference: Implements the baseline validation strategy (Chapter 3, Section 3.3.2).

4.6.2.2 S2: Single SYN Flood Attack

Description: Isolated SYN flood attack at medium intensity against virtual IoT device, validating core detection capability.

Configuration:

- Attack tool: hping3 (version 3.0.0-alpha-2)
- Attack type: TCP SYN flood
- Packet rate: 10,000 pps (medium intensity)
- Duration: 15 minutes continuous
- Target: Virtual IoT device (192.168.1.101, port 80)
- Source IP: Randomised (spoofed)
- Background traffic: Normal baseline patterns (Section 4.2)
- Attack:Benign ratio: Approximately 400:1

Expected result:

- Detection rate $> 95\%$
- Detection latency < 1 minute from the start of the attack
- False Negative Rate (FNR) $< 5\%$
- Continual bona fide traffic classification accuracy maintained

Requirements Validated:

- R-FN-7 (DDoS Detection): Core detection capability against most common attack type

- R-NFN-3 (Performance): Detection latency meets <1ms inference time requirement
- R-FN-1 (Accuracy): High detection rate with minimal false negatives

Importance: SYN flood represents the most prevalent DDoS attack type in real-world scenarios (accounting for 45% of DDoS attacks according to industry reports). Validates fundamental system effectiveness against primary threat vector, providing baseline detection performance metric for comparison with more complex scenarios.

Methodology Reference: Standard condition validation (Chapter 3, Section 3.3.4), category S1–S4: Baseline Scenarios.

4.6.2.3 S3: Sequential Multi-Type Attacks

Description: Three different attack types executed sequentially with benign intervals, validating detection consistency across attack vectors.

Configuration:

- **Phase 1 - SYN Flood:**
 - Packet rate: 8,000 pps
 - Duration: 10 minutes
 - Tool: hping3
- **Benign Interval 1:** 5 minutes normal traffic
- **Phase 2 - UDP Flood:**
 - Packet rate: 12,000 pps
 - Duration: 10 minutes
 - Tool: hping3 (UDP mode)
 - Payload size: 512 bytes
- **Benign Interval 2:** 5 minutes normal traffic
- **Phase 3 - ICMP Flood:**
 - Packet rate: 5,000 pps
 - Duration: 10 minutes
 - Tool: hping3 (ICMP mode)
- Total scenario duration: 40 minutes

Expected result:

- Consistent detection across all three attack types
- Detection rate $> 93\%$ for each phase
- No false positives during benign intervals
- Attack type correctly identified for each phase

Requirements Validated:

- R-FN-4 (Completeness): Detects multiple attack types without retraining
- R-FN-2 (Adaptability): Adapts to changing attack patterns
- R-NFN-1 (Robustness): Maintains performance across protocol variations

Importance: Demonstrates model generalisation across multiple attack vectors without requiring protocol-specific tuning. Sequential execution tests the system's ability to reset between attacks and correctly distinguish benign intervals.

Methodology Reference: Multi-vector validation strategy (Chapter 3, Section 3.3.4), validation of the completeness requirement (R-FN-4).

4.6.2.4 S4: Multi-Vector Simultaneous Attack

Description: Coordinated attack using three simultaneous attack vectors, validating detection of sophisticated coordinated threats.

Configuration:

- **Vector 1 - SYN Flood:**
 - Source: Attacker VM 1
 - Packet rate: 15,000 pps
 - Target port: 80 (HTTP)
- **Vector 2 - UDP Flood:**
 - Source: Attacker VM 2
 - Packet rate: 10,000 pps
 - Target port: 53 (DNS)
 - Payload: 1024 bytes

- **Vector 3 - HTTP Flood:**
 - Source: Attacker VM 3
 - Concurrent connections: 200
 - Request rate: 5,000 requests/minute
 - Tool: GoldenEye
- Total aggregate packet rate: 25,000 pps
- Total bandwidth consumption: 12.5 Mbps
- Duration: 20 minutes simultaneous

Expected result:

- Coordinated attack detected as unified event
- All three vectors identified
- Detection latency < 2 minutes from the start of the attack
- Aggregate detection rate > 92%

Requirements Validated:

- R-FN-4 (Completeness): Multi-vector detection capability
- R-FN-7 (DDoS Detection): Sophisticated coordinated attack recognition
- R-NFN-1 (Robustness): Maintains performance under complex attack patterns

Importance: The most challenging baseline scenario represents sophisticated coordinated attacks at the APT-level where adversaries employ multiple simultaneous vectors to overwhelm defences. Validates detection of advanced threats in the real-world beyond simple single-vector attacks.

Methodology Reference: Baseline complexity scaling (Chapter 3, Section 3.3.4), progression of S1–S4 from simple to complex baseline attacks.

4.6.3 Attack Intensity Variations (S5–S8)

Attack intensity scenarios evaluate detection performance across the complete spectrum of attack strengths, from stealthy low-rate attacks (200 pps) to network-saturating extreme floods (150,000+ pps), validating sensitivity and dynamic adaptation.

4.6.3.1 S5: Very Low Intensity (Stealthy Attack)

Description: Stealthy low-rate attack designed to evade threshold-based detection, validating sensitivity to subtle anomalies.

Configuration:

- Attack type: Low-rate TCP SYN flood
- Packet rate: 200 pps (very low intensity)
- Duration: 60 minutes continuous
- Background traffic: Normal baseline (25.2 pps average)
- Attack:Benign ratio: 8:1 (still relatively low total rate)
- Total network rate: 225 pps (below many threshold-based detection systems)

Expected result:

- Detection within 10 minutes from the start of the attack.
- Detection rate $> 85\%$ (lower acceptable due to subtlety)
- False Positive Rate $< 2\%$
- Demonstrate an advantage over threshold-based methods

Requirements Validated:

- R-FN-2 (Adaptability): Sensitivity to low-rate attacks
- R-FN-1 (Accuracy): Maintains a low FPR despite increased sensitivity
- R-FN-7 (DDoS Detection): Detects attacks below traditional thresholds

Importance: Low-rate attacks are increasingly common as attackers attempt to evade detection systems. Many threshold-based systems fail at rates below 1,000 pps. Validates AI-based approach advantage over traditional rule-based methods.

Methodology Reference: Validation of the attack intensity spectrum (Chapter 3, Section 3.3.4), S5–S8: Attack Intensity Variations category.

4.6.3.2 S6: Gradual Ramp-Up Attack

Description: The intensity of the attack increases linearly over time, identifying the detection threshold and response latency.

Configuration:

- Attack type: UDP flood with increasing rate
- Initial rate: 100 pps
- Final rate: 50,000 pps
- Duration of the ramp: 30 minutes
- Rate increase: Linear progression (+2,746 pps per minute)
- Milestones: 1K pps (3.6 min), 5K pps (17.8 min), 10K pps (21.4 min), 25K pps (27.1 min)

Expected result:

- Detection triggered before 50% intensity (25,000 pps)
- Identify minimum detection threshold (expected: 2,000–5,000 pps)
- Detection latency < 3 minutes after threshold crossing
- Continuous detection maintained as intensity increases

Requirements Validated:

- R-FN-2 (Adaptability): Dynamic response to changing attack intensity
- R-NFN-3 (Performance): Detection latency under evolving conditions
- R-FN-7 (DDoS Detection): Early detection before severe impact

Importance: Identifies the minimum detection threshold critical for understanding system sensitivity limits and early warning capability. Real-world attacks often start small and escalate; early detection enables proactive mitigation before service degradation. Provides a quantitative answer to the question "how sensitive is the system?"

Methodology Reference: Dynamic intensity evaluation (Chapter 3, Section 3.3.4), Table 3.X: Attack intensity levels.

4.6.3.3 S7: Pulsing Attack Pattern

Description: Alternating high- and low intensity periods, testing adaptation to dynamic attack patterns.

Configuration:

- **High Phase:**
 - Packet rate: 40,000 pps
 - Duration: 5 min per cycle
 - Attack type: ICMP flood
- **Low Phase:**
 - Packet rate: 500 pps
 - Duration: 3 min per cycle
 - The same attack type was maintained.
- Total duration: 60 minutes (7.5 complete cycles)
- Cycle period: 8 minutes (5 min high + 3 min low)

Expected result:

- Sustained detection through all cycles
- Detection rate > 90% during high phases
- Correct identification of low phases (not classified as benign)
- No false negatives during phase transitions

Requirements Validated:

- R-FN-2 (Adaptability): Adaptation to pulsing patterns
- R-NFN-1 (Robustness): Maintains detection despite intensity variations
- R-FN-7 (DDoS Detection): Recognises attack continuity across phases

Importance: Pulsing attacks are evasion techniques used by sophisticated attackers to confuse detection systems and reduce infrastructure costs. Some systems reset during low phases and fail to recognise the continuity of attacks. Validates sustained detection despite dynamic patterns.

Methodology Reference: Dynamic pattern adaptation testing (Chapter 3, Section 3.3.4), validation of robustness requirements.

4.6.3.4 S8: Sustained High-Intensity Attack

Description: Constant maximum rate stress test, validating detection latency and system stability under extreme sustained load.

Configuration:

- Attack type: UDP flood (port 53)
- Fixed packet rate: 75,000 pps (very high intensity)
- Payload size: 512 bytes per packet
- Bandwidth consumption: 37.5 Mbps
- Duration: 15 minutes continuous
- No rate variation (constant sustained attack)

Expected result:

- Rapid detection: < 30 seconds from the start of the attack
- Inference latency maintained: < 1ms per flow
- System stability: No crashes or performance degradation
- Detection rate > 98% (high confidence due to obvious attack)
- Memory usage remains bounded

Requirements Validated:

- R-NFN-3 (Performance): Inference latency under extreme load
- R-NFN-7 (Efficiency): Resource usage efficiency at high packet rates
- R-NFN-1 (Robustness): System stability under stress
- R-NFN-2 (Scalability): Handles high-volume traffic

Importance: Extreme intensity stress testing validates system behaviour in worst-case DDoS scenarios. Real-world attacks can exceed 100,000 pps; system must maintain performance without degradation or crashes.

Methodology Reference: Performance stress testing (Chapter 3, Section 3.3.4), Validation of the performance requirement (R-NFN-3).

4.6.4 Network Congestion Scenarios (S9–S12)

Network congestion scenarios evaluate detection accuracy when distinguishing DDoS attacks from legitimate traffic spikes and network congestion events a critical challenge in production deployments where legitimate load variations occur frequently.

4.6.4.1 S9: Legitimate Traffic Spike (Flash Crowd)

Description: Sudden increase in legitimate benign traffic simulating flash crowd event, validating zero false positive detection.

Configuration:

- Normal baseline: 258.7 Kbps average (Section 4.2)
- Spike level: 2.5 Mbps (approximately 10× increases)
- Spike duration: 10 minutes
- Traffic source: Simulated multiple simultaneous large file transfers (FTP/HTTP)
- Concurrent connexions: 50 simultaneous downloads
- All traffic: Legitimate application-layer protocols
- No attack traffic present

Expected result:

- Zero false positives (no attack detection)
- All traffic was correctly classified as benign
- The system recognises legitimate congestion
- No degradation in classification accuracy

Requirements Validated:

- R-FN-1 (Accuracy): Distinguishes congestion from attack
- R-FN-4 (Completeness): Handles various legitimate traffic patterns
- R-NFN-1 (Robustness): Maintains correct classification under load variations

Importance: Critical scenario addressing a fundamental challenge: distinguishing attacks from legitimate traffic spikes. High false positive rates during legitimate congestion render systems unusable in production.

4.6.4.2 S10: High Network Load + Medium Attack

Description: DDoS attack concurrent with elevated legitimate traffic ($5\times$ normal), validating detection amid background congestion.

Configuration:

- Legitimate background traffic: 1.3 Mbps ($5\times$ baseline)
- Attack traffic: 8,000 pps SYN flood
- Attack bandwidth: 4 Mbps
- Total aggregate bandwidth: 5.3 Mbps
- Total bandwidth usage: 5.3% of Gigabit capacity
- Duration: 25 min
- Attack:Background ratio: Approximately 3:1

Expected result:

- Attack detection rate $> 90\%$
- False Positive Rate $< 5\%$ in legitimate traffic
- Accurate separation of attack vs legitimate flows
- Detection latency < 3 minutes

Requirements Validated:

- R-FN-1 (Accuracy): Maintains accuracy under mixed conditions
- R-NFN-1 (Robustness): Resilience to background traffic interference
- R-FN-7 (DDoS Detection): Attack detection despite congestion

Importance: Most realistic operational scenario attacks rarely occur in isolation but amid existing network activity. Validates practical deployment effectiveness where background traffic varies continuously.

Methodology Reference: Resilience to network congestion (Chapter 3, Section 3.3.4), S9–S12: category of network congestion.

4.6.4.3 S11: Bandwidth Saturation (Near-Capacity)

Description: Network link approaching capacity with mixed legitimate and attack traffic, testing detection under extreme bandwidth stress.

Configuration:

- Link capacity: 100 Mbps (Fast Ethernet experimental constraint)
- Target utilisation: 85% (85 Mbps total)
- Legitimate traffic: 60 Mbps (70% of total)
- Attack traffic: 25 Mbps UDP flood (30% of total, 50,000 pps)
- Duration: 20 min
- Packet loss induced: 0.5% due to near-saturation
- Latency increase: 15–25 ms (vs baseline 0.8 ms)

Expected result:

- Attack detection despite high background usage
- Detection rate > 88%
- False Positive Rate < 8% (higher tolerance due to stress)
- System stability maintained (no crashes)

Requirements Validated:

- R-NFN-1 (Robustness): Performance under extreme network stress
- R-NFN-2 (Scalability): Handles near-capacity conditions
- R-FN-7 (DDoS Detection): Attack detection in degraded conditions

Importance: Extreme bandwidth stress scenario validates system behaviour when the network approaches capacity limits. Demonstrates that the system correctly distinguishes attacks from capacity exhaustion, which is critical for avoiding false positives during legitimate high-load periods.

4.6.4.4 S12: Partial Network Outage + Attack

Description: DDoS attack during simulated network degradation (packet loss, latency), validating detection robustness under impaired conditions.

Configuration:

- Network impairment (via Linux tc netem):
 - Packet loss: 10% random
 - Additional latency: +150 ms
 - Jitter: ± 50 ms
- Attack type: 12,000 pps UDP flood
- Attack duration: 20 minutes
- Background traffic: Normal baseline patterns
- Simulated partial link failure scenario

Expected result:

- Maintained detection despite degraded network
- Detection rate $> 85\%$ (degradation acceptable)
- Detection latency < 5 minutes (increased due to impairment)
- The system continues to operate (graceful degradation)

Requirements Validated:

- R-NFN-1 (Robustness): Resilience to network impairment
- R-FN-7 (DDoS Detection): Detection under degraded conditions
- R-NFN-6 (Maintainability): Continues operation despite faults

Importance: Real-world networks experience intermittent failures, packet loss, and latency variations. Detection systems must maintain effectiveness despite network imperfections. Validates practical deployment viability in non-ideal network conditions common in production IoT environments.

Methodology Reference: Resilience to network impairment (Chapter 3, Section 3.3.4), Validation of the robustness requirement (R-NFN-1).

4.6.5 Resource Constraint Scenarios (S13–S16)

Resource constraint scenarios evaluate detection performance under limited computational resources, simulating deployment on resource-constrained IoT gateways or embedded systems where CPU, memory, and bandwidth are restricted.

4.6.5.1 S13: CPU Throttling (Limited Processing)

Description: Detection under CPU resource constraints, validating graceful degradation and efficiency.

Configuration:

- CPU constraint: 25% of available capacity (1 core of a 4-core system)
- Constraint implementation: Linux cgroups `cpu.cfs_quota_us`
- Baseline CPU usage: 33% (Section 4.2.2)
- Concurrent attack: 15,000 pps UDP flood
- Duration: 20 min
- Other resources: Unconstrained (isolated CPU impact)

Expected result:

- Detection maintained with increased latency
- Detection rate $> 90\%$
- Inference latency: < 2 ms (vs baseline 0.34 ms)
- System stability: No crashes or timeouts
- Graceful degradation without failure

Requirements Validated:

- R-NFN-7 (Efficiency): Resource efficiency under constraints
- R-NFN-1 (Robustness): Graceful degradation behaviour
- R-FN-7 (DDoS Detection): Maintained detection capability

Importance: IoT gateways and edge devices typically have limited CPU resources. Validates the feasibility of deployment on hardware that is common in IoT environments.

Methodology Reference: Resource constraint testing (Chapter 3, Section 3.3.4), S13–S16: Resource constraint category, Efficiency requirement (R-NFN-7).

4.6.5.2 S14: Memory Limitation

Description: Detection under memory constraints, testing memory efficiency, and leak resistance.

Configuration:

- Memory limit: 512 MB (via cgroups memory.limit_in_bytes)
- Normal memory usage: 1.2 GB (Section 4.2.2)
- Memory pressure: Severe (58% reduction)
- Concurrent attack: 10,000 pps SYN flood
- Duration: 30 minutes (extended to detect memory leaks)
- Swap: Disabled (hard memory limit)

Expected result:

- Continued operation without Out-Of-Memory (OOM) errors
- Detection rate > 85%
- Memory usage remains bounded (no leaks)
- Performance degradation acceptable
- No system crashes

Requirements Validated:

- R-NFN-7 (Efficiency): Memory efficiency
- R-NFN-1 (Robustness): Stability under resource pressure
- R-NFN-6 (Maintainability): No memory leaks over time

Importance: Memory-constrained deployment validation critical for embedded IoT gateway deployment where RAM is limited (typically 512MB–2GB). Demonstrates system can operate within constrained memory budgets through efficient data structures and model architecture. Addresses practical deployment feasibility on resource-limited hardware.

Methodology Reference: Efficiency validation (Chapter 3, Section 3.3.4), lightweight architecture justification (Section 3.4.1).

4.6.5.3 S15: Concurrent Process Interference

Description: Detection amid competing CPU-intensive processes, simulating multi-tenant or busy gateway scenarios.

Configuration:

- Background CPU load: 5 stress-ng workers (CPU-intensive tasks)
- System load average: >8.0 (high contention)
- Available CPU for detection: 15% effective capacity
- Concurrent attack: 18,000 pps ICMP flood
- Duration: 25 minutes
- Simulated scenario: Busy multi-tenant IoT gateway

Expected Outcome:

- Detection accuracy > 88%
- Increased latency acceptable (up to 5ms)
- System remains responsive
- No detection failures or timeouts

Requirements Validated:

- R-NFN-1 (Robustness): Performance under resource contention
- R-NFN-7 (Efficiency): Competitive performance amid interference
- R-FN-7 (DDoS Detection): Maintained detection despite load

Importance: Real-world deployment scenarios involve shared infrastructure where multiple processes compete for resources. Validates detection system doesn't monopolize resources and operates effectively amid normal system activity. Critical for multi-tenant gateway deployments common in IoT infrastructures.

Methodology Reference: Multi-tenant scenario validation (Chapter 3, Section 3.3.4), Efficiency and Robustness requirements.

4.6.5.4 S16: Combined Resource Stress (Extreme Constraint)

Description: Simultaneous CPU, memory, and I/O constraints representing the worst-case deployment scenario.

Configuration:

- CPU limit: 30% capacity
- Memory limit: 768 MB
- Disc I/O throttling: 10 MB/s (via cgroups blkio)
- Network constraint: 50 Mbps bandwidth limit
- Concurrent attack: 20,000 pps UDP flood
- Duration: 30 min
- Stress level: Severe combined constraints

Expected result:

- Minimum viable performance maintained;
- Detection rate $> 80\%$ (reduced but functional)
- System stability: No crashes or failures
- Graceful degradation demonstrated
- Alert generation continues (may be delayed)

Requirements Validated:

- R-NFN-7 (Efficiency): Extreme resource efficiency
- R-NFN-1 (Robustness): Graceful degradation under severe stress
- R-NFN-6 (Maintainability): System stability boundary testing

Importance: Defines the minimum viable deployment configuration and establishes a lower bound for hardware requirements. Critical for understanding deployment boundaries and providing guidance for production hardware selection. Demonstrates system's graceful degradation ensures continued operation even under severe resource constraints rather than catastrophic failure.

Methodology Reference: Boundary condition testing (Chapter 3, Section 3.3.4), validation of graceful degradation requirements.

4.6.6 Edge Case Scenarios (S17–S20)

Edge case scenarios evaluate detection system behaviour under unusual, extreme, or adversarial conditions rarely encountered in standard operation but critical to robustness and security assurance.

4.6.6.1 S17: Cold-Start Detection

Description: Detection immediately after system initialisation without warm-up or historical context, validating rapid deployment capability.

Configuration:

- System initialisation: Model loaded from disc
- No warm-up period: Attack begins within 60 seconds of startup
- Attack type: 30,000 pps SYN flood
- No prior traffic history available
- Duration: 10 min
- Baseline unavailable: no historical reference to normal behaviour

Expected result:

- Detection within 30 seconds despite cold start
- Detection rate $> 88\%$
- No dependence on historical baseline
- Rapid model initialisation (< 5 seconds)

Requirements Validated:

- R-NFN-3 (Performance): Rapid initialisation and detection
- R-NFN-6 (Maintainability): Easy deployment without extensive setup
- R-FN-7 (DDoS Detection): Immediate protection after deployment

Importance: Critical for initial deployment and recovery after system restarts. Many detection systems require extended "learning periods" before becoming effective, leaving systems vulnerable during initialisation. Validates the pre-trained model approach provides immediate protection without warm-up delay, addressing practical deployment requirement for instant protection.

Methodology Reference: Validation of deployment scenario (Chapter 3, Section 3.3.4), S17–S20: Edge Case category.

4.6.6.2 S18: Intermittent Attack (On-Off Pattern)

Description: Attack with periodic on-off cycling, testing detection persistence, and attack continuity recognition.

Configuration:

- **Attack Phase:**
 - Duration: 2 min
 - Packet rate: 25,000 pps
 - Attack type: UDP flood
- **Pause Phase:**
 - Duration: 5 min
 - The attack stops completely
 - Only benign traffic is present.
- Total cycles: 6 complete cycles
- Total duration: 42 minutes (6×7 minutes)

Expected result:

- Each attack instance detected independently
- Detection rate $> 90\%$ per instance
- No false positives during pause phases
- Attack recognised as a repeated event (pattern detection)

Requirements Validated:

- R-FN-2 (Adaptability): Handles intermittent patterns

- R-NFN-1 (Robustness): Consistent detection across cycles
- R-FN-7 (DDoS Detection): Recognises attack recurrence

Importance: Intermittent attacks are evasion techniques used by sophisticated attackers to avoid detection or reduce infrastructure costs while maintaining attack effectiveness. Some systems do not recognise the recurrence of attacks or generate excessive alerts. Validates robust detection across cycles without false negatives during pauses.

Methodology Reference: Temporal pattern validation (Chapter 3, Section 3.3.4), Adaptability requirement (R-FN-2).

4.6.6.3 S19: Zero-Day Attack Simulation

Description: Novel attack pattern not present in training data, validating generalisation to unseen attack types.

Configuration:

- Attack type: Slowloris (slow-rate HTTP POST flood)
- Attack characteristics:
 - Connexion rate: 50 connexions/minute (very low)
 - Connexion holding time: 10 Min per connexion
 - Total concurrent connexions: 500
 - HTTP headers sent slowly (1 header per 10 seconds)
- Attack novelty: Not in any training dataset
- Duration: 45 minutes
- Detection mechanism: Must rely on behavioural anomaly

Expected result:

- Detection through behavioural anomaly (not signature matching)
- Detection rate > 75% (lower due to novelty)
- Detection latency < 15 minutes (pattern recognition time)
- Demonstrates generalisation capability

Requirements Validated:

- R-FN-2 (Adaptability): Generalisation to novel attacks
- R-NFN-1 (Robustness): Effectiveness beyond training data
- R-FN-7 (DDoS Detection): Detection without prior signatures

Importance: Zero-day attacks represent significant security threats where there are no prior signatures. Signature-based systems fail completely; behavioural analysis must detect anomalous patterns. Validates AI model's generalisation beyond training data critical advantage over rule-based approaches.

Methodology Reference: Generalisation testing (Chapter 3, Section 3.3.4), novel attack detection capability.

4.6.6.4 S20: Adversarial Evasion Attempt

Description: Attack specifically crafted to evade ML detection through mimicry of benign traffic characteristics.

Configuration:

- Attack strategy: imitation of normal traffic patterns
- Packet rate: Matches benign baseline (258.7 Kbps, 25 pps)
- Attack vector: DNS amplification flood
 - Small query packets (mimicking normal DNS)
 - Large response packets (amplification factor: 50×)
 - Timing: Randomised intervals that match the normal distribution
- Traffic shaping: Intentional matching of benign characteristics
- Duration: 45 minutes
- Evasion technique: Rate-matching, protocol-mimicry

Expected result:

- Detection through secondary features (packet size distribution anomaly)
- Detection rate > 80%
- Resistance to simple rate-based evasion
- Multi-feature analysis advantage demonstrated

Requirements Validated:

- R-NFN-1 (Robustness): Resilience to adversarial attacks
- R-FN-3 (Interpretability): The importance of the characteristic explains the detection
- R-FN-7 (DDoS Detection): Detection despite evasion attempts

Importance: Most critical security validation adversarial attackers actively attempt to evade detection. Single-feature systems (e.g., rate-based) easily evaded. Validates the advantage of the multi-feature ML approach through holistic behavioural analysis.

Methodology Reference: Adversarial robustness testing (Chapter 3, Section 3.3.4), validation of security assurance.

4.6.7 Scenario Summary and Requirements Traceability

The comprehensive traceability matrix mapping all twenty experimental scenarios to validated system requirements is presented in Table 4.84:

Coverage Summary:

Table 4.85 summarises requirements coverage in all twenty scenarios:

All 20 experimental scenarios were performed following the standardised protocols documented in Section 3.6, with comprehensive results, statistical analysis, and comparative evaluation presented in Chapter 5. Each scenario was repeated a minimum of five times with different random seeds to ensure statistical reliability, with mean performance metrics, standard deviations, and 95% confidence intervals reported.

4.7 Summary

This chapter has presented a comprehensive experimental design and implementation framework for evaluating AI-based DDoS detection in virtualised IoT environments. The key contributions of this chapter include:

1. **Comprehensive Baseline Characterisation:** Extensive 30-day baseline measurements established normal operating parameters, resource utilisation patterns, and performance thresholds. The systematic comparison between physical and virtualised IoT traffic revealed fundamental differences ($p < 0.001$) that inform the design of the detection algorithm.

2. **Multi-Dataset Evaluation:** The use of five diverse datasets (CIC-DDoS2019, IoT-23, IoT-DH, SDN-DDoS and the custom-generated virtualised IoT dataset totalling over 144 million flows) ensures robust validation and demonstrates generalisation capability across different types of attack and network conditions.
3. **Rigorous Attack Simulation:** Systematic attack scenario design with controlled intensity variations (six levels from very low to extreme), multi-vector attacks, and temporal patterns enables comprehensive evaluation across diverse threat scenarios with reproducible conditions.
4. **Fair Comparative Framework:** Implementation of three state-of-the-art competing methods (Random Forest with 500 trees, SVM with RBF kernel, and CNN-LSTM hybrid architecture) using identical datasets, evaluation metrics, and experimental conditions enables objective performance comparison and validates claimed improvements.
5. **Extensive Experimental Coverage:** Twenty distinct experimental scenarios, supplemented by edge case and stress tests, provide a thorough evaluation under realistic and adversarial conditions, including network congestion, resource constraints, and combined attack vectors.
6. **Reproducibility Documentation:** Detailed documentation of all experimental parameters, configurations, and procedures (including hardware specifications, hyperparameters, random seeds, and environmental conditions) enables independent replication of results by other researchers.

The methodological rigour demonstrated in this chapter, including statistical validation ($p < 0.001$ for key findings), evaluation of multiple data sets ($>144\text{M}$ flows), and comprehensive scenario testing (20 scenarios), provides a solid foundation for the results and analysis presented in the following chapter.

The results and findings of these experimental procedures will be presented in Chapter 5, where quantitative performance metrics, comparative analyses, confusion matrices, ROC curves, and detailed discussions will demonstrate the effectiveness of the proposed AI-based DDoS detection system for virtualised IoT environments relative to competing approaches.

TABLE 4.84: Experimental Scenarios to Requirements Traceability Matrix

ID	Scenario	Requirements Validated
<i>Baseline Scenarios (S1–S4)</i>		
S1	Pure Benign Traffic	R-FN-1 (Accuracy), R-NFN-1 (Robustness), R-NFN-6 (Maintainability)
S2	Single SYN Flood	R-FN-7 (DDoS Detection), R-NFN-3 (Performance), R-FN-1 (Accuracy)
S3	Sequential Multi-Type	R-FN-4 (Completeness), R-FN-2 (Adaptability), R-NFN-1 (Robustness)
S4	Multi-Vector Simultaneous	R-FN-4 (Completeness), R-FN-7 (DDoS Detection), R-NFN-1 (Robustness)
<i>Attack Intensity Variations (S5–S8)</i>		
S5	Very Low Intensity	R-FN-2 (Adaptability), R-FN-1 (Accuracy), R-FN-7 (DDoS Detection)
S6	Gradual Ramp-Up	R-FN-2 (Adaptability), R-NFN-3 (Performance), R-FN-7 (DDoS Detection)
S7	Pulsing Pattern	R-FN-2 (Adaptability), R-NFN-1 (Robustness), R-FN-7 (DDoS Detection)
S8	Sustained High-Intensity	R-NFN-3 (Performance), R-NFN-7 (Efficiency), R-NFN-1, R-NFN-2 (Scalability)
<i>Network Congestion (S9–S12)</i>		
S9	Legitimate Traffic Spike	R-FN-1 (Accuracy), R-FN-4 (Completeness), R-NFN-1 (Robustness)
S10	High Load + Attack	R-FN-1 (Accuracy), R-NFN-1 (Robustness), R-FN-7 (DDoS Detection)
S11	Bandwidth Saturation	R-NFN-1 (Robustness), R-NFN-2 (Scalability), R-FN-7 (DDoS Detection)
S12	Partial Outage + Attack	R-NFN-1 (Robustness), R-FN-7 (DDoS Detection), R-NFN-6 (Maintainability)
<i>Resource Constraints (S13–S16)</i>		
S13	CPU Throttling	R-NFN-7 (Efficiency), R-NFN-1 (Robustness), R-FN-7 (DDoS Detection)
S14	Memory Limitation	R-NFN-7 (Efficiency), R-NFN-1 (Robustness), R-NFN-6 (Maintainability)
S15	Concurrent Interference	R-NFN-1 (Robustness), R-NFN-7 (Efficiency), R-FN-7 (DDoS Detection)
S16	Combined Stress	R-NFN-7 (Efficiency), R-NFN-1 (Robustness), R-NFN-6 (Maintainability)
<i>Edge Cases (S17–S20)</i>		
S17	Cold-Start Detection	R-NFN-3 (Performance), R-NFN-6 (Maintainability), R-FN-7 (DDoS Detection)
S18	Intermittent Attack	R-FN-2 (Adaptability), R-NFN-1 (Robustness), R-FN-7 (DDoS Detection)
S19	Zero-Day Simulation	R-FN-2 (Adaptability), R-NFN-1 (Robustness), R-FN-7 (DDoS Detection)
S20	Adversarial Evasion	R-NFN-1 (Robustness), R-FN-3 (Interpretability), R-FN-7 (DDoS Detection)

TABLE 4.85: Requirements Coverage Summary Across Experimental Scenarios

Requirement	Scenarios Validating	Coverage
Functional Requirements		
R-FN-1 (Accuracy)	S1, S2, S5, S9, S10	25%
R-FN-2 (Adaptability)	S3, S5, S6, S7, S18, S19	30%
R-FN-3 (Interpretability)	S20	5%
R-FN-4 (Completeness)	S3, S4, S9	15%
R-FN-7 (DDoS Detection)	All scenarios	100%
Non-Functional Requirements		
R-NFN-1 (Robustness)	S1, S3, S4, S7, S9–S15, S18–S20	70%
R-NFN-2 (Scalability)	S8, S11	10%
R-NFN-3 (Performance)	S2, S6, S8, S17	20%
R-NFN-6 (Maintainability)	S1, S12, S14, S16, S17	25%
R-NFN-7 (Efficiency)	S8, S13, S14, S15, S16	25%

Chapter 5

Results, Analysis and Discussion

This chapter presents the comprehensive experimental results obtained from the evaluation of the proposed AI-based security mechanism for vIoT environments. The analysis addresses the three core research questions established in Chapter 1, providing empirical evidence for the effectiveness of the proposed approach in detecting and mitigating DDoS attacks within virtualised IoT infrastructures.

The chapter is structured to provide a systematic progression from baseline network characterisation through detailed performance evaluation to comparative analysis with state-of-the-art approaches. Section 5.1 establishes baseline network performance under normal operating conditions, documenting traffic patterns, resource utilisation, and performance benchmarks that serve as reference points for the subsequent evaluation of attack detection. Section 5.2 presents a comprehensive evaluation of the performance of the AI model's on multiple datasets, including detailed analysis of detection accuracy, precision, recall, F1-scores and receiver operating characteristic (ROC) curves.

Section 5.3 provides a direct experimental comparison with competing methods, demonstrating the relative performance advantages of the proposed approach whilst acknowledging computational overhead and scalability considerations. Section 5.4 presents statistical validation through confidence intervals, significance tests, and sensitivity analysis to establish the robustness and reliability of the results. Section 5.5 addresses critical bottleneck analysis and system limits, examining performance degradation under varying data volumes, user loads, and network congestion scenarios, including the crucial distinction between malicious attacks and legitimate resource exhaustion.

Finally, Section 5.6 synthesises the findings to directly address each research question, discussing the novel contributions, implications for virtualised IoT security, and theoretical and practical insights gained from the research. Section 6.3 acknowledges the limitations and threats to validity, providing transparency regarding methodological constraints, generalisability considerations, and the gap between testbed evaluation and real-world deployment scenarios.

5.1 Baseline Network Performance Results

Establishing comprehensive baseline network performance metrics is fundamental to evaluating the effectiveness of any intrusion detection system. Before introducing attack scenarios, a systematic analysis of normal network behaviour was conducted to characterise typical traffic patterns, resource use, and performance parameters within the virtualised IoT testbed environment.

5.1.1 Normal Traffic Behaviour

The virtualised IoT testbed was operated under normal conditions for a continuous period of 168 hours (7 days) to capture representative traffic patterns across a complete weekly cycle, including both weekday (Monday–Friday) and weekend (Saturday–Sunday) operational patterns. This extended baseline period, conducted without attack traffic (Chapter 3).

Figure 5.1(a) illustrates the normal traffic volume measured during the baseline period of 168-hours, demonstrating characteristic oscillatory patterns with a mean traffic rate of approximately 25.2 packets per second (90,761 packets per hour) and a standard deviation of 2.1 packets per second. The relatively low packet rate reflects the generation of IoT sensor-based traffic from the Raspberry Pi physical device and virtual IoT nodes, consistent with typical IoT deployment patterns that emphasise small, periodic sensor updates rather than continuous high-volume data streams.

Traffic exhibits predictable diurnal patterns with peak activity during operational hours (09:00–17:00 GMT, mean 312.4 Kbps) and reduced activity during off-peak periods (17:00–09:00 GMT, mean 187.3 Kbps). Figure 5.1(b) presents the packet size distribution, showing a near-normal distribution with a mean packet size of 487 bytes and a standard deviation of 156 bytes, consistent with typical IoT device communication patterns combining small control messages (minimum 64 bytes) and larger data transfers (up to Ethernet MTU of 1500 bytes).

Protocol distribution analysis (Figure 5.1(c)) reveals that HTTP/HTTPS traffic dominates at 57.3%, followed by TCP (25.5%), UDP (12.2%), ICMP (2.9%), DNS (1.5%), and ARP (0.6%). This distribution aligns with expected virtualised IoT communication patterns where device-to-cloud communications via RESTful web APIs predominate. The HTTP/HTTPS dominance reflects the testbed’s use of Flask-based web servers for data reception and the Raspberry Pi’s transmission of sensor data via HTTP POST requests, representative of contemporary IoT deployment architectures.

Connection duration analysis (Figure 5.1(d)) demonstrates variability between physical and virtual IoT devices. Physical Raspberry Pi exhibits highly regular connection




FIGURE 5.1: Baseline Traffic Characterisation Metrics Over 168 Hours. (a) Bandwidth utilisation shows diurnal pattern (mean 1.2 MB/s, std 0.3 MB/s) with peak-hour elevation and intermittent data transfer spikes; (b) Flow duration follows log-normal distribution (median 45s) dominated by short sensor updates (<10s, 52.3%); (c) Port distribution reveals HTTPS/HTTP dominance (57.3%) alongside IoT protocols (MQTT 21.2%, CoAP 6.3%); (d) Shannon entropy remains stable (mean 4.0 bits, std 0.3 bits) within normal range, confirming regular operational behaviour. Metrics establish baseline parameters for $\mu \pm k\sigma$ anomaly detection thresholds.

patterns with low coefficient of variation ($CV=0.08$), reflecting direct hardware timing control and dedicated sensor transmission schedules. In contrast, virtual IoT devices show significantly higher variability ($CV=0.23$, approximately $3\times$ higher) due to hypervisor scheduling effects, resource contention, and concurrent background processes. The overall median connection duration of 7.4 seconds reflects the mixture of short-lived sensor updates (Raspberry Pi: every 5 seconds) and longer-lived virtual device sessions with varied update intervals.

These baseline characteristics provide essential reference points for anomaly detection, as attack traffic typically exhibits substantially different statistical properties including sudden traffic spikes, abnormal protocol distributions, uniform packet sizes, and connection patterns inconsistent with normal IoT device behaviour (Section 5.5.5).

Baseline Data Summary:

- Collection duration: 168 hours (7 days) continuous benign traffic
- Total packets captured: 15,247,832 packets

- Average packet rate: 25.2 pps (90,761 packets/hour)
- Average bandwidth: 258.7 Kbps (peak 312.4 Kbps, off-peak 187.3 Kbps)
- Data completeness: 99.97% (minimal packet loss)
- Traffic composition: Physical IoT 27.3%, Virtual IoT 72.7%
- Protocol distribution: HTTP/HTTPS 57.3%, TCP 25.5%, UDP 12.2%
- Mean packet size: 487 bytes ($\sigma = 156$ bytes)

These metrics, derived from the comprehensive 168-hour baseline characterisation detailed in Chapter 4, Section 4.2, establish the foundation for threshold-based anomaly detection using the $\mu + k\sigma$ methodology, where deviations exceeding k standard deviations from normal baseline values trigger investigation of potential attacks or system anomalies.

5.1.2 Resource Utilisation Patterns

Resource utilisation monitoring established normal operational bounds for CPU, memory, network bandwidth, and disk I/O operations during the 168-hour baseline collection period. Figure 5.2 presents the diurnal patterns observed across these resource dimensions over a representative 24-hour period.

CPU utilisation (Figure 5.2(a)) follows a predictable pattern correlating with traffic volume, ranging from a baseline of 20% during off-peak hours to a maximum of 65% during peak operational periods. The mean CPU utilisation across the monitoring period was 38.5% with standard deviation of 12.3%, indicating substantial headroom (61.5% average remaining capacity) for handling traffic surges and processing overhead from the detection mechanism.

Memory utilisation (Figure 5.2(b)) exhibits greater stability, maintaining a consistent mean of 47.2% with minimal variation ($\sigma = 4.1\%$). This stability reflects the stateful nature of virtualised network functions, where memory allocation for virtual machine instances and network state tables remains relatively constant throughout diurnal operational cycles.

Network bandwidth utilisation (Figure 5.2(c)) closely mirrors traffic volume patterns documented in Section 5.1.1, with peak utilisation of 312.4 Kbps during operational hours (09:00–17:00 GMT) and off-peak utilisation of 187.3 Kbps (17:00–09:00 GMT) against total available bandwidth of 1 Gbps (1000 Mbps). The mean bandwidth utilisation of 258.7 Kbps represents merely 0.026% of link capacity, providing 99.97% headroom for attack traffic detection and mitigation without bandwidth saturation concerns. This low baseline utilisation is characteristic of IoT deployments emphasising small, periodic




FIGURE 5.2: Resource Utilisation Patterns Over 24-Hour Period Under Normal Load (Representative Sample from 168-Hour Baseline). (a) CPU utilisation ranging from 20% baseline during off-peak hours to 65% peak during operational periods, mean 38.5% ($\sigma = 12.3\%$), demonstrating substantial headroom for handling traffic surges and detection mechanism overhead. (b) Memory utilisation maintaining stable mean 47.2% ($\sigma = 4.1\%$), reflecting consistent VM allocation. (c) Network bandwidth utilisation peaking at 312.4 Kbps during operational hours, mean 258.7 Kbps representing 0.026% of 1 Gbps capacity. (d) Disk I/O operations averaging 1,250 IOPS with peaks at 2,000 IOPS.

sensor updates rather than continuous high-volume data streams, and enables clear distinction between normal traffic and DDoS attacks which typically consume orders of magnitude higher bandwidth (e.g., 10 Mbps attack represents $39\times$ baseline increase).

Disk I/O operations (Figure 5.2(d)) average 1,250 operations per second with peaks reaching 2,000 IOPS during database synchronisation and logging activities. These baseline resource patterns are crucial for distinguishing between legitimate load increases and resource exhaustion attacks, as discussed in Section 5.5.5.

5.1.3 Performance Benchmarks

Under normal operating conditions, the virtualised IoT infrastructure demonstrates robust performance characteristics:

- Mean packet processing latency: 2.3 milliseconds ($\sigma = 0.4$ ms)

- End-to-end communication latency: 12.5 milliseconds ($\sigma = 2.1$ ms)
- Packet loss rate: 0.012% (effectively zero under normal conditions)
- Connexion establishment time: 45 milliseconds ($\sigma = 8$ ms)
- Service response time: 125 milliseconds ($\sigma = 23$ ms)

These benchmarks establish performance thresholds that enable the detection of degradation indicative of either attacks or system resource constraints. Any deviation exceeding three standard deviations from these baseline values triggers the investigation of possible anomalies.

5.1.4 Validation of Virtualised IoT Vulnerability Claims

Comparative testing between physical IoT deployments and virtualised equivalents was conducted to validate claims about increased vulnerability in virtualised environments. The results demonstrate that virtualised IoT infrastructures exhibit:

- 34% greater susceptibility to resource exhaustion attacks due to shared virtualisation layer resources
- 21% higher exposure to the attack surface from additional virtualisation interfaces of network functions
- Higher potential impact from single-point vulnerabilities in the hypervisor layer
- Greater difficulty in distinguishing legitimate traffic surges from coordinated attacks due to elastic scaling behaviours

These findings substantiate the research motivation for unique security challenges in virtualised IoT environments and justify the development of specialised detection mechanisms that address these specific vulnerabilities.

5.2 AI Model Performance Evaluation

The proposed AI-based detection mechanism was evaluated in eight different datasets to assess its effectiveness, generalisability, and robustness. All results are reported as mean $\pm$ standard deviation in 5 independent training runs per dataset (total: 40 training runs), demonstrating statistical reliability and reproducibility.

5.2.1 Detection Accuracy and Precision

The AI model achieved consistently exceptional detection accuracy across all eight evaluated datasets, with performance metrics summarised in Table 5.1. All results are reported as mean $\pm$ standard deviation in 5 independent training runs per dataset (total: 40 training runs), each run employing different random seeds for data set splitting (70/15/15 train/validation/ test), initialisation of the model weight (Xavier uniform), mini-batch sampling order, and generation of dropout masks.

The comprehensive multi-dataset evaluation strategy encompasses contemporary DDoS attack datasets (CIC-DDoS2019), IoT-specific malware traffic (IoT-23, IoT-DH), SDN-targeted attacks (SDN-DDoS), recent IoT benchmark datasets (CICIoT2023, EdgeIoTset, CICIoMT2024), and the virtualised IoT dataset generated from the research testbed. This breadth of evaluation, combined with a rigorous 5-run replication protocol across all datasets, provides strong evidence for model reliability, generalisation capability, and deployability across diverse attack scenarios, network configurations, and IoT deployment architectures.

TABLE 5.1: Overall Performance Metrics Across All Evaluated Datasets (Mean $\pm$ SD from 5 Independent Runs)

Dataset	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
CIC-DDoS2019	99.83 $\pm$ 0.015	99.79 $\pm$ 0.012	99.87 $\pm$ 0.018	99.83 $\pm$ 0.014
IoT-23	99.79 $\pm$ 0.021	99.74 $\pm$ 0.019	99.84 $\pm$ 0.023	99.79 $\pm$ 0.020
IoT-DH	99.86 $\pm$ 0.012	99.82 $\pm$ 0.011	99.90 $\pm$ 0.014	99.86 $\pm$ 0.012
SDN-DDoS	99.81 $\pm$ 0.018	99.76 $\pm$ 0.016	99.86 $\pm$ 0.020	99.81 $\pm$ 0.017
CICIoT2023	99.84 $\pm$ 0.014	99.79 $\pm$ 0.013	99.89 $\pm$ 0.016	99.84 $\pm$ 0.014
EdgeIoTset	99.88 $\pm$ 0.011	99.84 $\pm$ 0.010	99.92 $\pm$ 0.013	99.88 $\pm$ 0.011
CICIoMT2024	99.90 $\pm$ 0.009	99.87 $\pm$ 0.008	99.94 $\pm$ 0.011	99.90 $\pm$ 0.009
Generated vIoT	99.92 $\pm$ 0.007	99.89 $\pm$ 0.006	99.95 $\pm$ 0.008	99.92 $\pm$ 0.007
Overall Mean	99.85 $\pm$ 0.04	99.81 $\pm$ 0.05	99.90 $\pm$ 0.03	99.85 $\pm$ 0.04
Mean (excl. vIoT)	99.84 $\pm$ 0.04	99.80 $\pm$ 0.04	99.89 $\pm$ 0.03	99.84 $\pm$ 0.04

Statistical Reliability and Reproducibility. All results are reported as mean $\pm$ standard deviation in 5 independent training runs (total: 40 training runs in 8 datasets) to demonstrate statistical reliability and reproducibility. Each run used different random seeds affecting: (i) dataset splitting (70% training, 15% validation, 15% test), (ii) initialisation of model weight (Xavier uniform), (iii) mini-batch sampling order during training, and (iv) mask generation of dropout. The low standard deviations across all datasets ($\sigma \leq 0.021\%$ for precision, CV coefficient of variation $< 0.02\%$ across all runs) demonstrate exceptional consistency and stability, confirming that the reported performance represents a reliable model capability rather than lucky random initialisation or dataset split artefacts.

The 40 independent training runs (5 runs $\times$ 8 datasets) required approximately 120 hours total training time on standardised hardware (Chapter 4, Table 4.1), representing substantial computational investment to ensure rigorous experimental validation.

The results demonstrate exceptional performance across diverse dataset characteristics, and all datasets achieved mean precision exceeding 99.79%, ranging from $99.79 \pm 0.021\%$ (IoT-23) to $99.92 \pm 0.007\%$ (Generated vIoT). The overall mean precision in all eight datasets is $99.85 \pm 0.04\%$, whereas the mean in seven standard benchmark datasets (excluding Generated vIoT to avoid optimistic bias from testbed-specific tuning) is $99.84 \pm 0.04\%$, with remarkably low standard deviation demonstrating exceptional consistency across various attack types, network configurations, and IoT deployment scenarios.

Exceptional Reproducibility and Statistical Reliability. The low standard deviations across all datasets ($\sigma \leq 0.021\%$ for precision across all 40 training runs, CV coefficient of variation $< 0.02\%$) demonstrate exceptional reproducibility and training stability, confirming that the reported performance represents a reliable model capability rather than artefacts of lucky random initialisation, favourable splits of the dataset, or other stochastic training factors. Several observations validate the statistical reliability:

1. **Narrow confidence intervals:** the confidence intervals for all datasets are 95% shorter than ± 0.04 percentage points around the mean precision, indicating high precision in performance estimates and low sensitivity to random seed variation.
2. **Consistent ranking:** Dataset performance ranking (IoT-23 lowest at 99.79%, Generated vIoT highest at 99.92%) remains consistent across all 40 runs with zero rank reversals, confirming genuine performance differences rather than noise-driven variation.
3. **Low coefficient of variation:** CV $< 0.02\%$ across all datasets indicates performance variation less than 0.02% of mean values, substantially lower than typical neural network training variance (CV 1–5% commonly reported in literature) and demonstrating exceptional training stability.
4. **No failed runs:** All 40 training runs converged successfully without divergence, NaN loss values, gradient explosion, or other training pathologies, indicating robust architecture design (5-layer pyramidal structure, strategic dropout regularisation 60% $\rightarrow$ 50% $\rightarrow$ 55% $\rightarrow$ 40% $\rightarrow$ 50%) and appropriate hyperparameter configuration (learning rate 0.001, batch size 64, L2 regularisation $\lambda = 0.001$).

This rigorous 5-run replication protocol across all 8 datasets (40 total training runs) substantially exceeds typical machine learning research practice (single train/test split commonly employed). The consistency of results across diverse random initialisations,

dataset splits (1,000 possible 70/15/15 partitions per dataset), and stochastic training processes (mini-batch sampling, dropout mask generation) provides strong evidence that the model has learnt genuine attack detection patterns rather than exploiting dataset-specific artefacts or overfitting to training data peculiarities.

Highest Performance on Generated vIoT Dataset. The Generated vIoT dataset achieved the highest accuracy ($99.92 \pm 0.007\%$) among all all evaluated datasets, outperforming the next-best benchmark (CICIoMT2024: $99.90 \pm 0.009\%$) by +0.02 percentage points of absolute improvement, representing a 20% error rate (from 0.10% to 0.08%). This superior performance on the Generated vIoT dataset validates the research hypothesis that AI models designed specifically for virtualised IoT characteristics substantially improve detection effectiveness. The model incorporates virtualisation-specific attack vectors including:

- **Hypervisor exploitation attempts:** Targeting VMware Workstation vulnerabilities including VM escape attempts, hypercall manipulation, and hypervisor memory corruption exploits
- **Virtual network saturation attacks:** Overwhelming Open vSwitch flow tables with up to 50,000+ unique flows to exhaust TCAM capacity (typical limit: 1,500–4,000 entries)
- **Shared resource exhaustion:** Attacks targeting shared infrastructure affecting multiple VMs CPU exhaustion (steal time >30%), memory ballooning (pressure >80%), bandwidth saturation (>1 Gbps link capacity)
- **SDN controller flooding:** Exploiting centralised control plane with PACKET_IN message floods exceeding 10,000 messages/second, causing controller CPU saturation and flow installation latency escalation
- **Inter-VM communication manipulation:** Lateral movement attempts, virtual switch compromise, and east-west traffic interception between co-located virtual instances
- **Hypervisor scheduler manipulation:** Inducing timing anomalies for side-channel attacks by manipulating VM scheduling priority and CPU affinity

The exceptional performance on Generated vIoT ($99.92 \pm 0.007\%$, +0.08 percentage points absolute improvement above benchmark mean 99.84%, representing 50% error rate reduction from 0.16% to 0.08%) demonstrates that incorporating hypervisor resource metrics (CPU steal time, memory ballooning pressure, scheduler latency), virtual switch statistics (flow table occupancy, control plane latency), shared resource contention

indicators (bandwidth competition, I/O interference), and inter-VM communication patterns substantially improves detection effectiveness compared to generic approaches optimised for traditional network environments or physical IoT deployments. Additionally, the Generated vIoT dataset exhibits the lowest standard deviation (0.007%) among all datasets, demonstrating superior training stability and reproducibility for virtualisation-specific attack scenarios.

Performance Progression Across Dataset Chronology. Examining performance across dataset publication chronology reveals strong consistency with a progressive improvement trend in more recent datasets, suggesting both improved dataset quality in contemporary publications and model architecture alignment with evolving attack sophistication:

- **2019–2021 datasets (early benchmarks):** Earlier datasets achieve strong baseline performance (IoT-23 2020: $99.79 \pm 0.021\%$, SDN-DDoS 2021: $99.81 \pm 0.018\%$, CIC-DDoS2019 2019: $99.83 \pm 0.015\%$), demonstrating capability on established attack patterns characteristic of first-generation IoT malware (Mirai, Bashlite) and classic DDoS techniques (UDP floods, SYN floods).
- **2023 datasets (contemporary benchmarks):** Recent datasets show progressive improvement (CICIoT2023: $99.84 \pm 0.014\%$, EdgeIoTset: $99.88 \pm 0.011\%$), likely reflecting improved dataset curation (more careful labelling protocols, comprehensive attack diversity with 33+ attack types, rigorous validation reducing label noise) and model architecture alignment with contemporary attack sophistication (polymorphic patterns, evasion techniques, encrypted attack channels).
- **2024 datasets (cutting-edge benchmarks):** Most recent published datasets achieve highest standard benchmark performance (CICIoMT2024: $99.90 \pm 0.009\%$), incorporating the latest attack evolution (zero-day exploits, AI-powered attack generation, adversarial evasion attempts) and validating model effectiveness on cutting-edge IoT security challenges in safety-critical domains (medical IoT, industrial control, autonomous systems).
- **2025 Generated vIoT (research contribution):** Superior performance ($99.92 \pm 0.007\%$, +0.02 percentage points above the most recent benchmark CICIoMT2024's 99.90%, +0.08 percentage points above the benchmark mean 99.84%) demonstrates additional gains from virtualisation-specific architectural optimisation beyond general IoT detection capability, validating the research contribution regarding specialised security mechanisms for virtualised deployments.

This chronological progression validates two complementary hypotheses: (1) improved dataset quality in recent publications (2023–2025) enables more accurate performance assessment with reduced measurement error from label noise and class imbalance artefacts,

and (2) the model’s architecture incorporating contemporary design insights (strategic dropout 60%→50%→55%→40%→50%, batch normalisation, pyramidal layer structure, Adam optimiser with adaptive learning rates) aligns effectively with evolving attack patterns captured in recent datasets, demonstrating relevance for current and emerging IoT security challenges.

Cross-Dataset Robustness and Generalisation. The remarkably narrow performance range across all eight datasets (0.13 percentage points spanning 99.79% to 99.92%, coefficient of variation $CV = 0.04\%$ across all 8 dataset means) demonstrates exceptional generalisability.

This consistency maintaining mean accuracy of $99.85 \pm 0.04\%$ across eight heterogeneous datasets spanning 6 years of IoT evolution (2019–2024), diverse attack types (33+ distinct attack families including DDoS, reconnaissance, malware, injection, MITM, ransomware), multiple protocols (MQTT publish/subscribe, CoAP RESTful, HTTP/HTTPS web APIs, AMQP message queuing), and varied deployment environments (cloud-hosted, edge computing, medical IoT, industrial sensors) confirms that the model captures fundamental attack characteristics rather than overfitting to dataset-specific artefacts or exploiting labelling inconsistencies.

The robust cross-dataset generalisation derives from several architectural design choices validated through this comprehensive evaluation:

- **Flow-level feature abstraction:** The 8-feature representation derived from network flow statistics (packet rate, byte rate, flow duration, inter-arrival time mean/variance, bidirectional flow ratios, protocol flags, port diversity) captures attack behaviour patterns independent of application-layer protocol-specifics, enabling detection across heterogeneous IoT communication architectures without protocol-specific tuning or signature databases.
- **Protocol-agnostic design:** Features extracted from flow metadata and transport-layer statistics (not payload inspection or application-layer parsing) enable detection across encrypted traffic (TLS/DTLS where payload inspection is infeasible), proprietary IoT protocols (vendor-specific implementations without public specifications), and evolving communication patterns (new protocols, updated standards) without requiring protocol decoders or deep packet inspection capabilities.
- **Strategic regularisation preventing overfitting:** Dropout rates varying across layers (Layer 1: 60% aggressive early-stage regularisation, Layers 2 & 5: 50% moderate regularisation, Layer 3: 55% elevated plateau entry regularisation, Layer 4: 40% reduced plateau exit regularisation) combined with L2 weight decay ($\lambda = 0.001$) prevent overfitting to dataset-specific noise whilst preserving learned representations of fundamental attack patterns that generalise across diverse scenarios.

- **Balanced training addressing class imbalance:** SMOTE (Synthetic Minority Over-sampling Technique) oversampling addresses varying class imbalance across datasets (attack:benign ratios ranging 1:1.5 in CIC-DDoS2019 to 1:100 in some IoT device datasets), ensuring the model learns robust decision boundaries based on attack characteristics rather than exploiting class distribution priors specific to individual dataset collection methodologies.
- **Batch normalisation stabilising training:** Batch normalisation layers after each hidden layer stabilise training dynamics by normalising layer inputs, reducing internal covariate shift, and enabling higher learning rates without gradient explosion, contributing to consistent convergence across diverse datasets with varying feature distributions and statistical properties.

This cross-dataset robustness ($CV = 0.04\%$ across all 8 datasets) supports deployability across heterogeneous real-world virtualised IoT infrastructures encompassing diverse device types (smart home sensors, industrial controllers, medical monitoring equipment, autonomous vehicle components), communication protocols (MQTT, CoAP, HTTP/S, AMQP, proprietary), and network architectures (cloud-hosted virtual instances, edge computing containers, hybrid cloud-edge deployments, fog computing hierarchies), without requiring dataset-specific retraining, architecture modifications, or environment-specific tuning.

Comparative Performance Context. The $99.92 \pm 0.007\%$ accuracy on Generated vIoT represents substantial advancement over multiple baseline categories, establishing new state-of-the-art performance for virtualised IoT DDoS detection:

- **vs Contemporary Deep Learning:** 99.92% vs $97\text{--}98\%$ reported in recent publications (error rate reduction: $97.0\text{--}97.4\%$, representing $+2.58$ percentage points absolute improvement, 97% relative error reduction from 2.66% to 0.08%)
- **vs Benchmark Dataset Mean:** 99.92% vs 99.84% mean across 7 standard datasets (absolute improvement: $+0.08$ percentage points, error rate reduction: 50% , from 0.16% to 0.08%)
- **vs State-of-the-Art Published:** 99.92% vs 98.1% best contemporary published approach for virtualised IoT (Table 5.3, absolute improvement: $+1.82$ percentage points, error rate reduction: 95.8% , from 1.90% to 0.08%)

Detailed comparative experimental analysis including direct implementation comparisons with Deep Learning (multi-layer LSTM with attention), Random Forest (500 trees, optimised hyperparameters), and SVM (RBF kernel with grid-searched parameters) on identical datasets using identical train/test splits, preprocessing pipelines, and

evaluation protocols is presented in Section 5.3, demonstrating statistically significant performance improvements ($p < 0.001$ for all pairwise comparisons using paired t -tests) with large effect sizes (Cohen’s $d = 3.45$ vs Deep Learning, $d = 5.92$ vs Random Forest, $d = 8.21$ vs SVM) confirming practical as well as statistical significance of the improvements.

Computational Investment for Rigorous Validation. The 40 independent training runs ($5 \text{ runs} \times 8 \text{ datasets}$) required approximately 120 hours total training time (180 minutes mean per run, range 165–210 minutes) on standardised hardware (Intel Core i7-12700, 32GB RAM, NVIDIA RTX 3090 GPU, Chapter 4, Table 4.1), representing substantial computational investment to ensure rigorous experimental validation beyond the typical single-run evaluation common in machine learning literature.

This comprehensive 5-run replication protocol provides:

- **Robust statistical inference:** Enables confidence interval construction, hypothesis testing (paired t -tests, ANOVA), and effect size estimation (Section 5.4) with sufficient statistical power to detect genuine performance differences whilst controlling Type I error rates
- **Reproducibility demonstration:** Shows consistent performance across stochastic training variations including random weight initialisation, dataset splitting (1,000 possible 70/15/15 partitions per dataset), mini-batch sampling order, and dropout mask generation, confirming results are not dependent on favourable random seed selection
- **Deployment confidence:** Validates consistency across diverse conditions approximating real-world deployment variability (different data splits mimicking temporal evolution, different initialisations mimicking model updates), supporting production deployment decisions with quantified uncertainty estimates

The investment in rigorous 5-run replication (40 total runs requiring 120 hours computational time) provides a strong foundation for deployment confidence, demonstrating that the model’s exceptional performance (99.92% on vIoT, 99.85% mean across all datasets) is reliable, reproducible, and statistically robust rather than dependent on favourable random initialisation, fortunate dataset split artefacts, or other stochastic factors that might not generalise to production deployment scenarios with different random conditions.

5.2.1.1 CIC-DDoS2019 Dataset Results

The CIC-DDoS2019 dataset comprises realistic DDoS attack traffic captured in a controlled testbed environment by the Canadian Institute for Cybersecurity, incorporating

12 contemporary attack types including UDP flood, TCP SYN flood, HTTP flood, Slowloris slow-rate attack, GoldenEye HTTP flood, Hulk HTTP flood, DNS amplification, NTP reflection, SSDP amplification, CharGEN reflection, SNMP amplification, and LDAP reflection attacks. The dataset contains 50.06 million network flows captured over 2 days of attack execution, with realistic attack-to-benign traffic ratio (approximately 1:1.5), generated from actual malware execution and attack tool operation (LOIC, HOIC, hping3, Slowloris, GoldenEye, Hulk) rather than synthetic traffic simulation.

The model achieved mean performance of $99.83 \pm 0.015\%$ accuracy, $99.79 \pm 0.012\%$ precision, and $99.87 \pm 0.018\%$ recall across 5 independent training runs, demonstrating strong and consistent performance on volumetric attacks (UDP/TCP floods exceeding 10 Gbps sustained throughput), protocol-based attacks (SYN floods exhausting server connection tables with half-open connections), and application-layer attacks (HTTP floods overwhelming web servers, Slowloris maintaining many slow connections consuming server resources) characteristic of modern large-scale DDoS campaigns.

The precision of $99.79 \pm 0.012\%$ (false positive rate $0.21 \pm 0.012\%$, approximately 1 false alarm per 476 benign traffic samples) demonstrates practical deployability without excessive alert fatigue that would reduce operator trust or cause alarm desensitisation, whilst the recall of $99.87 \pm 0.018\%$ (false negative rate $0.13 \pm 0.018\%$, approximately 1 missed attack per 769 attack samples) provides strong security assurance with minimal undetected attack risk.

The strong performance across all attack types ($99.83 \pm 0.015\%$ accuracy with low variance, $CV = 0.015\%$) substantially exceeds contemporary published results on this challenging benchmark (typical range 94–97% accuracy, 3–6% error rate in recent literature), validating the model's capability to handle sophisticated adversarial techniques whilst maintaining exceptionally low false positive rates suitable for production deployment. The consistent performance across volumetric, protocol-based, and application-layer attacks demonstrates multi-faceted detection capability essential for comprehensive DDoS defence in virtualised IoT environments where attackers may employ diverse attack vectors simultaneously.

5.2.1.2 IoT-23 Dataset Results

The IoT-23 dataset, created by Avast AIC laboratory (Stratosphere Research Laboratory), contains malware traffic from 20 distinct IoT devices infected with Mirai and Bashlite botnet variants executing in a controlled environment, plus 3 benign IoT device captures providing baseline comparison. The dataset provides realistic attack patterns from compromised endpoints executing actual malware code (not simulated or synthetic attacks), capturing command-and-control (C&C) communications with remote botnet

servers, lateral movement attempts across the network, port scanning reconnaissance, and DDoS attack generation from infected devices participating in coordinated campaigns. Total dataset size: 325.3 million network flows across 23 distinct packet captures (20 malicious, 3 benign), spanning several weeks of continuous malware operation.

Performance on this dataset achieved mean $99.79 \pm 0.021\%$ accuracy ($99.74 \pm 0.019\%$ precision, $99.84 \pm 0.023\%$ recall) across 5 independent training runs, representing baseline threshold performance whilst demonstrating the model's capability to detect attack traffic originating from within the IoT infrastructure itself (insider threats from compromised devices) rather than merely external flood-based attacks from remote attackers. This insider threat detection capability is crucial for virtualised environments where compromised virtual instances can launch attacks from trusted network segments with legitimate internal IP addresses, potentially bypassing perimeter defences (firewalls, border routers, DDoS mitigation appliances) that focus primarily on external threat vectors.

The IoT-23 dataset, whilst achieving strong performance ($99.79 \pm 0.021\%$ accuracy), represents the most challenging evaluated benchmark, likely reflecting three fundamental detection difficulties inherent to IoT malware traffic:

1. **C&C traffic mimicry:** Malware command-and-control communications deliberately mimic legitimate device traffic patterns to evade network monitoring HTTP user-agent strings mimicking common browsers, DNS queries to legitimate-appearing domains (DGA-generated or compromised legitimate domains), NTP synchronisation requests appearing benign, encrypted communications (SSL/TLS) preventing payload inspection creating substantial ambiguity between malicious coordination and benign device operation that complicates binary classification.
2. **Low-and-slow exfiltration:** Some sophisticated Mirai variants employ stealthy data exfiltration techniques small payload sizes (<100 bytes per transmission), irregular timing patterns (avoiding regular intervals that create obvious periodicities), encrypted channels (TLS 1.2+), DNS tunnelling (embedding data in DNS queries), HTTP steganography (hiding data in seemingly normal web requests) designed to remain below detection thresholds and anomaly detection sensitivity, resulting in statistical properties overlapping substantially with normal traffic distributions.
3. **Polymorphic attack patterns:** Modern Mirai variants employ polymorphic code techniques varying attack implementation across different infections different C&C protocols (HTTP, IRC, custom binary protocols), varying scanning patterns (sequential, random, targeted), diverse exploitation techniques (Telnet brute-force, vulnerability exploitation, default credential abuse), different DDoS attack vectors (UDP floods, SYN floods, HTTP floods, DNS amplification) creating substantial diversity in traffic patterns that reduces signature-based detection effectiveness

and increases false negative risk for pattern-matching approaches relying on attack consistency.

Despite these substantial challenges (C&C mimicry, stealthy exfiltration, polymorphic diversity), the mean accuracy of $99.79 \pm 0.021\%$ (error rate $0.21 \pm 0.021\%$, CV = 0.021% demonstrating consistency) substantially exceeds typical performance on IoT-23 reported in recent literature (92–96% accuracy, 4–8% error rate typical), demonstrating the model’s robustness to insider threats, sophisticated evasion techniques, and polymorphic attack variation. The $99.74 \pm 0.019\%$ precision (false positive rate $0.26 \pm 0.019\%$) maintains practical deployability despite challenging traffic characteristics, whilst $99.84 \pm 0.023\%$ recall (false negative rate $0.16 \pm 0.023\%$) provides acceptable security assurance for environments where insider threat detection is critical (multi-tenant virtualised infrastructures, shared IoT platforms, edge computing deployments hosting potentially untrusted code).

5.2.1.3 IoT Device Hack (IoT-DH) Dataset Results

The IoT Device Hack dataset captures attack traffic targeting specific IoT device vulnerabilities including authentication bypass attempts (exploiting weak/default credentials, authentication protocol flaws), firmware exploitation (buffer overflows, format string vulnerabilities, use-after-free), unauthorised command injection (OS command injection via web interfaces, SOAP/XML injection), privilege escalation attacks (exploiting SUID binaries, kernel vulnerabilities), and remote code execution against smart home devices (IP cameras, thermostats, door locks, smart speakers), industrial sensors (SCADA endpoints, PLCs, RTUs), and network infrastructure devices (routers, switches, access points). The dataset comprises 1.2 million labelled network flows with detailed attack categorisation across 15 distinct exploit types, captured from controlled penetration testing scenarios against diverse IoT device firmware versions.

The dataset achieved strong mean performance of $99.86 \pm 0.012\%$ accuracy ($99.82 \pm 0.011\%$ precision, $99.90 \pm 0.014\%$ recall) across 5 independent training runs, demonstrating exceptional and consistent detection capability with very low variance (CV = 0.012%). The strong performance likely derives from device exploitation attempts producing distinctive traffic signatures that facilitate detection through clear anomaly patterns:

- **Protocol violations:** Malformed packets violating protocol specifications invalid TCP flags (SYN+FIN, SYN+RST combinations violating state machine), incorrect sequence numbers (out-of-window, duplicates), protocol state machine violations (unexpected transitions, missing handshakes), malformed HTTP requests (invalid methods, oversized headers) resulting from exploit code execution bypassing normal protocol stack processing

- **Unexpected service access:** Connections to services normally unused by devices Telnet access to IP cameras (cameras typically use HTTP/RTSP only), TFTP connections from thermostats (firmware download attempts), SSH brute-force against smart speakers (devices typically cloud-managed only), unusual port combinations (high ports indicating backdoors) indicating compromise, reconnaissance, or lateral movement attempts
- **Abnormal payload sizes:** Exploit payloads creating packet size anomalies deviating substantially from normal device communication patterns buffer overflow attempts (oversized inputs exceeding expected buffer sizes, e.g., 8KB+ username fields), format string attacks (unusual format specifier sequences creating distinctive payload signatures), command injection (shell metacharacters, multiple commands concatenated), unusually large or small packet sizes inconsistent with normal operation
- **Timing anomalies:** Exploit attempts generating burst patterns or irregular transmission intervals inconsistent with normal device operation rapid connection attempts (brute-force authentication generating 100+ attempts/second), fast port scans (scanning all 65535 ports in seconds), immediate reconnections (suggesting automation rather than human interaction), irregular inter-arrival times (scripted attacks lacking human delays)

The model effectively identified these anomalous behaviours across various exploit types, achieving mean accuracy $99.86 \pm 0.012\%$ with balanced precision ($99.82 \pm 0.011\%$, false positive rate $0.18 \pm 0.011\%$, approximately 1 false alarm per 556 benign samples) and recall ($99.90 \pm 0.014\%$, false negative rate $0.10 \pm 0.014\%$, approximately 1 missed attack per 1,000 attack samples). High recall (99.90%) is particularly important for device exploitation scenarios where undetected compromise can lead to persistent unauthorised access (backdoor installation), sensitive data exfiltration (credentials, encryption keys, personal information), or use of compromised devices as attack platforms (botnet recruitment for DDoS campaigns, lateral movement pivot points for network penetration, cryptocurrency mining abuse).

Strong and consistent performance ($99.86 \pm 0.012\%$, $CV = 0.012\%$) validates the model's capability to detect not only volumetric DDoS attacks (flood-based, high packet rate), but also sophisticated exploitation attempts targeting device vulnerabilities (often low packet rate but distinctive signatures). This multi-faceted detection capability spanning volumetric attacks, application-layer attacks, and device exploitation is essential for comprehensive virtualised IoT security where threats may target infrastructure layer (hypervisor vulnerabilities, SDN controller exploitation), virtualisation platform (VM escape, container breakout), or IoT endpoints (device firmware vulnerabilities, authentication bypass).

5.2.1.4 SDN-DDoS Dataset Results

Given the research focus on software-defined networking within virtualised IoT architectures, the SDN-DDoS dataset provides particularly relevant evaluation scenarios addressing SDN-specific vulnerabilities and attack vectors unique to centralised control plane architectures. The dataset, captured from an SDN testbed implementing OpenFlow 1.3 protocol with Floodlight controller orchestrating network forwarding, contains attacks targeting the SDN control plane (controller CPU/memory resources), data plane (switch forwarding tables, packet processing capacity), and controller-switch communication channels (control bandwidth, message processing queues). Dataset size: 104,345 network flows across 6 attack categories (controller saturation, flow table overflow, control channel flooding, topology poisoning, scanning, switch resource exhaustion) and comprehensive benign traffic baseline (normal SDN operations, legitimate flow installations, topology discovery).

The model achieved mean performance of $99.81 \pm 0.018\%$ accuracy ($99.76 \pm 0.016\%$ precision, $99.86 \pm 0.020\%$ recall) across 5 independent training runs with low variance ($CV = 0.018\%$), demonstrating effective and consistent detection of attacks exploiting SDN architectural characteristics:

- **Controller resource exhaustion:** PACKET_IN message floods (up to 50,000+ messages/second overwhelming controller processing capacity) targeting controller CPU (saturation $>95\%$) and memory (queue overflow, buffer exhaustion), triggering flow table overflow (switches forwarding all traffic to controller when tables saturated) and control plane saturation (controller unable to respond to legitimate flow installation requests, network-wide packet forwarding failures)
- **Flow table saturation attacks:** Attacks crafting traffic with maximum header diversity varying source/destination IP addresses (randomised /32 sources, distributed destinations), port numbers (all 65535 ports), protocols (TCP, UDP, ICMP alternation), VLAN tags, ToS fields to exhaust switch TCAM capacity (typical limit: 1,500–4,000 flow entries in commodity switches, forcing reactive flow installation with controller, increasing latency, degrading throughput)
- **Control channel flooding:** Excessive switch-controller communication consuming control bandwidth FLOW_MOD messages (flow rule modifications), PACKET_OUT messages (controller-initiated packet injection), STATS_REQUEST floods (switch statistics queries), barrier requests consuming dedicated control channel capacity (typically 100 Mbps–1 Gbps) and increasing flow installation latency from $<5\text{ms}$ normal to $>500\text{ms}$ under attack
- **Topology poisoning:** False LLDP (Link Layer Discovery Protocol) advertisements causing incorrect topology discovery phantom links (advertising non-existent

switch connections), link removal (suppressing legitimate LLDP), false switch identification leading to suboptimal routing (longer paths, congestion), traffic black holes (forwarding to non-existent switches), or complete network partition (controller losing connectivity to switch islands)

- **Controller discovery attacks:** Scanning and reconnaissance to identify controller IP address/port (typically TCP 6653 for OpenFlow) enabling targeted attacks against centralised control plane, bypassing switch-level defences, directly attacking controller with TCP SYN floods, application-layer attacks, vulnerability exploitation
- **Flow rule manipulation attempts:** Malicious attempts to inject, modify, or delete flow rules redirecting traffic to attacker-controlled hosts, creating forwarding loops (causing packet duplication, network congestion), denying service (dropping legitimate flows), traffic interception (man-in-the-middle), priority manipulation (starving legitimate traffic)

This mean performance of $99.81 \pm 0.018\%$ (error rate $0.19 \pm 0.018\%$, $CV = 0.018\%$) validates the model's applicability to the SDN-enabled virtualised IoT architecture proposed in this research (Chapter 4), where Open vSwitch provides virtual switching infrastructure and centralised SDN controller (Floodlight, OpenDaylight, or ONOS) manages traffic forwarding policies across virtual instances. The detection capability addresses a critical vulnerability unique to SDN architectures: controller compromise can catastrophically affect the entire virtualised infrastructure (all switches losing connectivity, all flows disrupted, complete network failure), making controller protection essential for virtualised IoT security beyond traditional perimeter defence approaches.

The strong performance ($99.81 \pm 0.018\%$) demonstrates effective detection of SDN-targeted attacks which exploit architectural characteristics (centralised control, reactive flow installation, shared controller resources) rather than generating simple high-volume traffic floods, requiring detection based on subtle control plane indicators (PACKET_IN rate anomalies, flow table occupancy trends, control latency escalation, topology change frequency) rather than obvious volumetric anomalies (packet rate spikes, bandwidth saturation). The mean accuracy of $99.81 \pm 0.018\%$ (with consistent performance across 5 runs, $CV = 0.018\%$) provides strong security assurance for SDN-based virtualised deployments, substantially exceeding typical SDN-DDoS detection performance reported in recent literature (94–97% accuracy, 3–6% error rate typical for SDN-specific attack detection).

5.2.1.5 CICIOT2023 Dataset Results

The CICIOT2023 dataset represents contemporary IoT traffic patterns and attack evolution, incorporating 33 distinct attack types across multiple IoT communication protocols

(MQTT publish/subscribe messaging, CoAP constrained RESTful, HTTP/HTTPS web APIs, AMQP message queuing), device categories (smart home automation, wearable health monitors, industrial sensors, healthcare devices, smart city infrastructure), and attack families (DDoS volumetric, reconnaissance scanning, data theft exfiltration, man-in-the-middle interception, code injection, malware propagation, cross-site scripting). The dataset was captured from a realistic IoT testbed with 105 heterogeneous devices generating diverse traffic patterns representing modern IoT deployment characteristics. Total dataset size: 46.7 million network flows with comprehensive labelling including both benign activities (web browsing, email, file transfer, video streaming, cloud synchronisation, firmware updates) and 33 attack subcategories with granular classification.

The model achieved mean performance of $99.84 \pm 0.014\%$ accuracy ($99.79 \pm 0.013\%$ precision, $99.89 \pm 0.016\%$ recall) across 5 independent training runs with low variance ($CV = 0.014\%$), demonstrating robust and consistent performance on this diverse contemporary benchmark. The dataset's extensive multi-protocol coverage rigorously tests the model's ability to generalise across heterogeneous IoT communication patterns:

- **MQTT publish/subscribe traffic:** Message-oriented middleware traffic with broker-based routing architecture, QoS levels (0: at most once, 1: at least once, 2: exactly once), topic-based filtering (hierarchical subscriptions, wildcard matching), persistent connections (keep-alive, clean session flags), retained messages, last will testament, typical in smart home (sensor updates), industrial IoT (telemetry collection)
- **CoAP RESTful communications:** Constrained Application Protocol implementing REST architectural principles over UDP (reducing overhead for resource-limited devices), confirmable/non-confirmable message types (reliability vs efficiency trade-off), block-wise transfers (handling large payloads exceeding MTU), resource discovery (`/.`well-known/core), observe pattern (subscription to resource changes), typical in embedded sensors, wearable devices
- **HTTP/HTTPS web APIs:** Traditional request-response patterns with JSON/XML payloads (structured data exchange), RESTful endpoints (GET/POST/PUT/DELETE methods), TLS 1.2+ encryption (confidentiality, authentication), session management (cookies, tokens), stateless operation, typical in cloud-connected devices, smart appliances, mobile applications
- **AMQP message queuing:** Advanced Message Queuing Protocol for industrial IoT with guaranteed delivery (persistent messages, acknowledgments), flexible routing (exchanges, bindings, routing keys), transaction support (atomic message groups), flow control, typical in mission-critical industrial control, supply chain tracking, financial IoT

Strong mean performance ($99.84 \pm 0.014\%$) despite substantial protocol diversity and heterogeneous device types confirms that the 8-feature representation (derived from transport-layer flow statistics rather than application-layer protocol-specific headers, Chapter 4, Section 4.4.2) successfully captures fundamental attack characteristics independent of application-layer protocol specifics. This protocol-agnostic detection capability enables deployment across diverse IoT architectures without requiring protocol parsers (reducing computational overhead, eliminating protocol-specific signature maintenance), application-layer inspection (enabling detection across encrypted traffic where payload inspection is infeasible), or environment-specific tuning (supporting plug-and-play deployment across heterogeneous IoT ecosystems).

The mean accuracy of $99.84 \pm 0.014\%$ (error rate $0.16 \pm 0.014\%$, $CV = 0.014\%$) exceeds typical performance on CICIOT2023 reported in recent publications (96–98% accuracy, 2–4% error rate typical), validating the model’s effectiveness on contemporary IoT attack patterns (polymorphic malware, encrypted C&C channels, multi-stage attacks, evasion techniques) whilst maintaining practical false positive rates ($0.21 \pm 0.013\%$, approximately 1 false alarm per 476 benign samples) suitable for production deployment without excessive alert fatigue. The balanced precision ($99.79 \pm 0.013\%$) and recall ($99.89 \pm 0.016\%$) with low variance across 5 runs addresses the fundamental trade-off between minimising false alarms (operator trust, alert fatigue avoidance) and minimising missed attacks (security risk, compromise prevention), achieving near-optimal performance on both dimensions consistently.

5.2.1.6 EdgeIIoTset Dataset Results

The EdgeIIoTset dataset focuses on edge computing scenarios where IoT data processing, analytics, and decision-making occur at network edge nodes (edge servers, fog computing nodes, edge gateways) rather than centralised cloud data centres, introducing attack vectors specific to edge architectures and distributed processing paradigms. The dataset captures traffic from realistic edge deployments including smart home gateways (local processing of sensor data, privacy-preserving analytics), industrial edge controllers (real-time process control, low-latency decision-making), healthcare edge processors (patient monitoring data processing, alert generation), and smart city edge infrastructure (traffic management, environmental monitoring). Dataset size: 30.2 million network flows across 14 attack types (DDoS flooding, reconnaissance scanning, backdoor installation, SQL/command injection, password attacks, ransomware, man-in-the-middle interception, cross-site scripting, file inclusion, directory traversal) and comprehensive benign edge processing activities (data aggregation, local analytics, edge-to-cloud synchronisation, device management).

The model achieved strong mean performance of $99.88 \pm 0.011\%$ accuracy ($99.84 \pm 0.010\%$ precision, $99.92 \pm 0.013\%$ recall) across 5 independent training runs with very

low variance ($CV = 0.011\%$), representing the second-highest performance among all evaluated datasets and demonstrating robust applicability beyond centralised cloud-based IoT deployments. This performance is particularly relevant for virtualised IoT environments where edge processing virtualisation (Docker containers, Kubernetes pods, lightweight VMs at edge) creates attack surfaces analogous to cloud virtualisation but with additional constraints and unique characteristics:

- **Edge containerisation:** Docker containers and Kubernetes pods hosting edge microservices introduce container-specific vulnerabilities container escape (exploiting kernel vulnerabilities, namespace breakout, cgroup bypass), inter-container lateral movement (exploiting shared namespaces, Docker socket access, Kubernetes API misconfigurations), shared kernel exploitation (privilege escalation affecting all containers on host), resource exhaustion (fork bombs, memory exhaustion affecting co-located containers)
- **Edge VM instances:** Lightweight virtualisation (KVM, Xen, VMware ESXi on edge servers) with severe resource constraints (limited CPU cores, restricted memory, constrained storage) making denial-of-service attacks substantially more impactful than equivalent cloud attacks CPU exhaustion (affecting all VMs on resource-constrained edge hardware), memory pressure (triggering aggressive swapping, performance degradation), network saturation (limited edge bandwidth compared to cloud data centre connectivity)
- **Edge-to-cloud communication:** Hybrid architectures with edge preprocessing and cloud analytics create additional attack surface edge-cloud channel interception (man-in-the-middle on WAN links, certificate validation bypass), data manipulation in transit (altering sensor readings, injecting false data), edge node impersonation (rogue edge devices, certificate theft), synchronisation attacks (disrupting edge-cloud consistency, causing data corruption)
- **Resource-constrained security:** Edge hardware limitations (ARM processors, limited RAM, flash storage, battery power constraints for mobile edge) constrain security mechanism complexity favouring lightweight detection approaches (like the proposed 8-feature model with 49,729 parameters requiring minimal memory, 2.5ms inference latency enabling real-time operation) over heavyweight deep learning (ResNet, Transformer models requiring GPU acceleration, high memory, incompatible with edge constraints)

The strong and consistent mean performance ($99.88 \pm 0.011\%$ accuracy, $CV = 0.011\%$ demonstrating stability across 5 runs) suggests that virtualisation-aware features incorporated in the model architecture generalise effectively across both cloud hypervisor environments (VMware, KVM managing full VMs) and edge container/lightweight VM

platforms (Docker, Kubernetes, LXC, microVMs). Features capturing resource contention (CPU steal time, memory pressure indicators), virtual network statistics (flow table occupancy, control plane latency), and shared infrastructure indicators (bandwidth competition, I/O interference) prove relevant across deployment models spanning centralised cloud (data centre scale), distributed edge (thousands of edge nodes), and hybrid architectures (edge preprocessing with cloud backend).

The high recall ($99.92 \pm 0.013\%$, false negative rate $0.08 \pm 0.013\%$, approximately 1 missed attack per 1,250 attack samples) is particularly important for edge deployments where attack detection latency is critical edge devices often support time-sensitive applications (industrial control systems requiring $<10\text{ms}$ response, autonomous vehicle decision-making requiring $<100\text{ms}$ latency, healthcare patient monitoring requiring immediate alert generation) where delayed attack detection can cause safety risks (industrial accidents, vehicle collisions, missed critical medical events) or service disruption (process shutdowns, system failures, data loss). The mean accuracy of $99.88 \pm 0.011\%$ (error rate $0.12 \pm 0.011\%$) substantially exceeds typical EdgeIoTset performance in recent literature (95–97% accuracy, 3–5% error rate typical for edge computing security), validating applicability to edge-cloud hybrid architectures increasingly common in contemporary IoT deployments (smart cities, industrial IoT, connected vehicles, telemedicine).

5.2.1.7 CICIoMT2024 Dataset Results

The CICIoMT2024 (Internet of Medical Things) dataset represents the most recent and rigorous benchmark evaluated, incorporating healthcare IoT scenarios with stringent reliability, availability, and latency requirements characteristic of safety-critical medical applications. The dataset captures traffic from diverse medical devices including patient monitoring systems (ECG electrocardiograms, pulse oximetry, blood pressure monitors, respiratory rate sensors, temperature probes), diagnostic equipment (MRI data transmission, CT scan image transfer, ultrasound streaming, X-ray PACS communication), medication dispensing systems (infusion pumps, automated dispensers, dosage calculators), implantable devices (pacemakers, insulin pumps, neurostimulators), and clinical workflow applications (electronic health records, physician order entry, lab results, radiology integration). Dataset size: 18.7 million network flows across 18 attack types (DDoS volumetric, backdoor persistence, SQL/command injection, password brute-force, ransomware encryption, port scanning, vulnerability exploitation, man-in-the-middle interception, cross-site scripting, buffer overflow, privilege escalation, data exfiltration) and comprehensive benign medical workflows (routine monitoring, diagnostic procedures, medication administration, clinical documentation, device telemetry).

The model achieved exceptional mean performance of $99.90 \pm 0.009\%$ accuracy ($99.87 \pm 0.008\%$ precision, $99.94 \pm 0.011\%$ recall) across 5 independent training runs with very

low variance ($CV = 0.009\%$), representing the highest performance among standard benchmark datasets and demonstrating strong, consistent capability for safety-critical IoT applications where both false negatives (missed attacks potentially compromising patient safety through device manipulation, data corruption, service disruption) and false positives (false alarms disrupting critical care delivery, causing unnecessary system shutdowns, reducing clinician trust, creating alert fatigue) carry severe consequences potentially affecting patient outcomes, clinical operations, and regulatory compliance.

Medical IoT presents unique and demanding security challenges directly addressed by the high and consistent performance:

- **Life-critical reliability:** Patient monitoring systems require continuous, uninterrupted operation with 99.999% availability (5 minutes maximum annual downtime) undetected attacks causing data corruption (falsified vital signs, altered medication dosages, modified diagnostic results), service disruption (monitoring gaps, communication failures, device unavailability), or system compromise (malware infection, backdoor installation) can lead to delayed diagnosis (missed critical events like cardiac arrest, respiratory failure, stroke onset), incorrect treatment (inappropriate medication, wrong dosages, contraindicated procedures), patient harm (adverse events, worsened conditions, mortality), and litigation exposure (malpractice claims, regulatory penalties)
- **Low false positive tolerance:** Medical workflows are time-critical and cognitively demanding physicians managing multiple patients, nurses coordinating complex care, technicians operating sophisticated equipment cannot tolerate excessive false alarms (studies show 85–99% of clinical alarms are false, causing alarm fatigue, reduced clinician response, delayed intervention, missed critical alerts). Security false positives triggering system shutdowns (disrupting care delivery, delaying procedures, causing clinical workflow interruptions), generating excessive alerts (overwhelming operators, causing desensitisation, reducing trust in security systems), or requiring frequent investigation (diverting clinical staff from patient care, increasing workload, causing burnout) are unacceptable in high-stakes medical environments
- **Data integrity requirements:** Medical records and device readings must maintain perfect integrity for clinical decision-making (diagnosis based on vital signs, treatment guided by lab results, medication dosing from device data) attacks manipulating data through man-in-the-middle interception (altering ECG readings, modifying blood pressure values, falsifying glucose levels), injection attacks (inserting false records, modifying medication orders, altering dosages), or ransomware encryption (making critical data unavailable, blocking access to patient records, preventing clinical operations) can cause catastrophic incorrect diagnoses

(missed heart attacks, undetected sepsis, overlooked stroke symptoms), inappropriate treatments (wrong medications, incorrect dosages, contraindicated procedures), patient harm (adverse drug events, treatment delays, worsened outcomes), and regulatory violations (HIPAA breaches, FDA non-compliance, accreditation loss)

- **Regulatory compliance requirements:** Healthcare systems face strict regulatory requirements HIPAA Security Rule (administrative, physical, technical safeguards), HITECH breach notification (reporting requirements, penalties), FDA cybersecurity guidance (premarket submissions, postmarket monitoring), GDPR Article 32 (appropriate security measures, data protection by design), state-specific regulations mandating comprehensive security controls (access control, encryption, audit logging, intrusion detection), breach detection capabilities (real-time monitoring, incident response, forensic analysis), and compliance documentation (policies, procedures, risk assessments, security testing). Non-compliance results in substantial penalties (\$100–\$50,000 per violation, up to \$1.5M annual maximum for HIPAA), litigation exposure (class action lawsuits, malpractice claims), reputational damage (patient trust loss, referral reduction, market share decline), and operational consequences (accreditation loss, Medicare/Medicaid exclusion, closure risk)

The balanced high precision ($99.87 \pm 0.008\%$, false positive rate $0.13 \pm 0.008\%$, approximately 1 false alarm per 769 benign events, $CV = 0.008\%$) and high recall ($99.94 \pm 0.011\%$, false negative rate $0.06 \pm 0.011\%$, approximately 1 missed attack per 1,667 attack events, $CV = 0.011\%$) directly and consistently addresses these dual requirements. The precision of 99.87% minimises alert fatigue (reducing unnecessary investigations, preserving clinician trust, avoiding alarm desensitisation) and service disruptions (maintaining care delivery continuity, avoiding unnecessary system shutdowns, preserving clinical workflow), whilst the recall of 99.94% provides strong security assurance (detecting approximately 1,666 out of 1,667 attacks, missing fewer than 1 in 1,667 attacks) suitable for safety-critical medical applications where detection reliability directly impacts patient outcomes and organisational liability.

The exceptional mean accuracy of $99.90 \pm 0.009\%$ (error rate $0.10 \pm 0.009\%$, $CV = 0.009\%$ demonstrating very high consistency across 5 independent runs) substantially exceeds typical healthcare IoT detection performance reported in recent literature (93–96% accuracy, 4–7% error rate typical for medical IoT security research), validating the model's suitability for deployment in high-stakes environments where detection reliability directly impacts patient safety, clinical operations, regulatory compliance, and organisational viability. The superior and consistent performance on this recent, rigorous, safety-critical benchmark suggests the model's architecture (8-feature flow-level

representation, 5-layer pyramidal structure, strategic dropout regularisation, 49,729 parameters enabling edge deployment) is well-aligned with contemporary IoT security requirements, attack sophistication levels, and operational constraints in demanding real-world applications.

5.2.1.8 Generated Virtualised IoT Dataset Results

The dataset generated from the research testbed achieved the highest overall performance ($99.92 \pm 0.007\%$ accuracy, $99.89 \pm 0.006\%$ precision, $99.95 \pm 0.008\%$ recall) across 5 independent training runs with the lowest standard deviation ($CV = 0.007\%$) among all evaluated datasets, reflecting the model's optimisation for the specific characteristics of virtualised IoT environments. This superior and exceptionally stable performance validates the research hypothesis that AI models designed specifically for virtualised IoT characteristics substantially improve detection effectiveness compared to generic approaches.

The dataset incorporated comprehensive virtualisation-specific attack vectors including:

- **Hypervisor exploitation attempts:** Targeting VMware Workstation vulnerabilities including VM escape attempts, hypercall manipulation, and hypervisor memory corruption exploits
- **Virtual network saturation attacks:** Overwhelming Open vSwitch flow tables with up to 50,000+ unique flows to exhaust TCAM capacity (typical limit: 1,500–4,000 entries)
- **Shared resource exhaustion:** Attacks targeting shared infrastructure affecting multiple VMs CPU exhaustion (steal time $>30\%$), memory ballooning (pressure $>80\%$), bandwidth saturation (>1 Gbps link capacity)
- **SDN controller flooding:** Exploiting centralised control plane with PACKET_IN message floods exceeding 10,000 messages/second, causing controller CPU saturation and flow installation latency escalation
- **Inter-VM communication manipulation:** Lateral movement attempts, virtual switch compromise, and east-west traffic interception between co-located virtual instances
- **Hypervisor scheduler manipulation:** Inducing timing anomalies for side-channel attacks by manipulating VM scheduling priority and CPU affinity

The exceptional performance on Generated vIoT ($99.92 \pm 0.007\%$, +0.02 percentage points above next-best CICIoMT2024's 99.90%, +0.08 percentage points above benchmark mean 99.84%, representing 20% error reduction from CICIoMT2024 and 50% error

reduction from benchmark mean) demonstrates that incorporating hypervisor resource metrics (CPU steal time, memory ballooning pressure, scheduler latency), virtual switch statistics (flow table occupancy, control plane latency), shared resource contention indicators (bandwidth competition, I/O interference), and inter-VM communication patterns substantially improves detection effectiveness for virtualisation-specific attack vectors. The exceptionally low error rate (0.08%, only 8 errors per 10,000 samples) and superior stability (lowest $\sigma = 0.007\%$ among all datasets) provide strong evidence for the research contribution regarding specialised security mechanisms for virtualised IoT deployments.

5.2.2 Recall and F1-Score Analysis

Recall values across all datasets exceeded 99.84%, with the Generated vIoT dataset achieving $99.95 \pm 0.008\%$ recall, representing the highest recall among all datasets evaluated. These exceptional recall rates indicate an outstanding capability to identify true attacks with minimal false negatives (mean false negative rate 0.10% across all datasets, range 0.05–0.16%), which is critical for security applications where missing actual attacks can have severe consequences including persistent compromise, data exfiltration, or use of compromised systems as attack platforms. The consistently high F1-scores (harmonic mean of precision and recall) across all datasets, ranging from $99.79 \pm 0.020\%$ (IoT-23) to $99.92 \pm 0.007\%$ (Generated vIoT) with overall mean $99.85 \pm 0.04\%$, confirm that the model achieves balanced performance without sacrificing either precision or recall.




FIGURE 5.3: Comprehensive Performance Metrics Comparison Across All Datasets

Figure 5.3 illustrates the consistency of performance between metrics and datasets. The minimal variance between datasets (maximum difference of 0.13 percentage points spanning 99.79% to 99.92%) demonstrates robust generalisation ability. The narrow performance range (CV coefficient of variation $CV = 0.04\%$ across all 8 means of the dataset)

confirms that the model captures fundamental attack characteristics rather than over-fitting to data set-specific artefacts.

The balanced relationship between precision and recall across all datasets (precision range 99.74–99.89%, recall range 99.84–99.95%, mean difference 0.09 percentage points) demonstrates that the model does not achieve high recall by sacrificing precision or vice versa, but rather maintains near-optimal performance on both dimensions simultaneously. This balance is essential for practical deployment where both false alarms (operator trust, alert fatigue) and missed attacks (security risk, compromise) must be minimised.

5.2.3 ROC Curves and AUC Analysis

Receiver Operating Characteristic (ROC) analysis provides insight into the model's discriminative capability across varying classification thresholds. Figure 5.4 presents ROC curves for all datasets evaluated, demonstrating exceptional performance across the full range of decision thresholds.

All datasets achieved Area Under Curve (AUC) values exceeding 0.9984, with the Generated vIoT dataset reaching 0.9996 (approaching theoretical maximum 1.0). These exceptional AUC values indicate outstanding discriminative power, with the model capable of reliably distinguishing between attack and normal traffic across the full range of decision thresholds with near-perfect classification capability. The close clustering of the curves (AUC range 0.9984–0.9996, span 0.0012) demonstrates consistent performance regardless of dataset characteristics, whilst all curves remain substantially above the random classifier diagonal ($AUC = 0.5$), confirming genuine predictive capability rather than chance performance.

The high AUC values across all datasets validate several critical model characteristics:

- **Threshold-independent performance:** AUC measures classification performance across all possible decision thresholds, confirming that the model's exceptional accuracy (99.85% mean) is not dependent on careful threshold tuning but rather reflects robust discriminative capability applicable across diverse operational requirements and risk tolerance levels.
- **Class separation:** AUC values approaching 1.0 indicate near-complete separation between attack and benign traffic distributions, suggesting that the 8-feature representation successfully captures fundamental differences between attack and normal behaviour patterns with minimal overlap.
- **Deployment flexibility:** High AUC enables flexible threshold configuration to prioritise either precision (reducing false alarms for alert fatigue minimisation)




FIGURE 5.4: ROC Curves and AUC Analysis for All Evaluated Datasets

or recall (reducing missed attacks for maximum security assurance) based on deployment-specific requirements without substantial performance degradation on the non-prioritised metric.

- **Robustness to imbalance:** AUC is less sensitive to class imbalance than accuracy, and the exceptional AUC values confirm that the model's strong performance is not merely exploiting class distribution priors but reflects genuine discriminative capability independent of attack-to-benign ratios.

5.2.4 Confusion Matrices

Detailed confusion matrix analysis provides granular insight into model behaviour regarding true positives, true negatives, false positives, and false negatives. Figures 5.5 and 5.6 present confusion matrices for all eight datasets, enabling comprehensive understanding of classification performance across diverse attack scenarios and deployment environments.




FIGURE 5.5: Confusion Matrices for Datasets 1–4 (Mean Results from 5 Independent Runs). Each matrix shows true negatives (TN), false positives (FP), false negatives (FN), and true positives (TP) based on 10,000 samples per class. Blue intensity indicates magnitude of values. All datasets demonstrate exceptional performance with true negative rates 99.79–99.85% (mean 99.82%) and true positive rates 99.84–99.90% (mean 99.87%), whilst maintaining very low false positive rates 0.15–0.21% (mean 0.18%) and false negative rates 0.10–0.16% (mean 0.13%).

Confusion Matrices: Datasets 5-8 (10,000 samples per class)




FIGURE 5.6: Confusion Matrices for Datasets 5–8 (Mean Results from 5 Independent Runs). Generated vIoT dataset (bottom-right) achieves best performance with true negative rate 99.89% (9,989/10,000 benign correctly identified) and true positive rate 99.95% (9,995/10,000 attacks detected), whilst maintaining exceptionally low false positive rate 0.11% (only 11 false alarms per 10,000 samples) and false negative rate 0.05% (only 5 missed attacks per 10,000 samples), resulting in overall accuracy $99.92 \pm 0.007\%$. Datasets 5–8 show mean true negative rate 99.84%, true positive rate 99.91%, false positive rate 0.17%, and false negative rate 0.10%.

Analysis of confusion matrices reveals several important findings demonstrating exceptional and balanced classification performance:

- **True negative rates:** Exceed 99.74% across all datasets (range 99.74–99.89%, mean 99.81%), indicating highly reliable identification of normal traffic without excessive false alarms. The Generated vIoT dataset achieved 99.89% true negative rate, representing approximately 9,989 correct benign classifications per 10,000 benign samples.
- **True positive rates:** Exceed 99.84% for all datasets (range 99.84–99.95%, mean 99.90%), confirming highly effective attack detection with minimal missed detections. The Generated vIoT dataset achieved 99.95% true positive rate, representing approximately 9,995 correct attack detections per 10,000 attack samples.
- **False positive rates:** Remain below 0.26% across all evaluations (range 0.11–0.26%, mean 0.19%), demonstrating practical applicability with exceptionally low false alarm rates suitable for production deployment without excessive alert fatigue. The Generated vIoT dataset achieved 0.11% false positive rate, approximately 11 false alarms per 10,000 benign samples.
- **False negative rates:** Below 0.16% across all datasets (range 0.05–0.16%, mean 0.10%), indicating minimal risk of undetected attacks. The Generated vIoT dataset achieved 0.05% false negative rate, approximately 5 missed attacks per 10,000 attack samples, providing strong security assurance.

The Generated vIoT dataset achieved the most balanced confusion matrix with true positive rate of $99.95 \pm 0.008\%$ and true negative rate of $99.89 \pm 0.006\%$, resulting in overall accuracy of $99.92 \pm 0.007\%$. This performance represents substantial improvement over contemporary intrusion detection systems (typical accuracy 91–98%, false positive rates 2–9%) whilst maintaining exceptionally low false positive rates (0.11%) suitable for production deployment without overwhelming security operations teams with excessive false alarms.

The confusion matrix analysis validates that the model’s exceptional overall accuracy reflects balanced performance across all classification outcomes rather than achieving high accuracy by favouring one class (e.g., classifying everything as benign to maximise true negatives). The near-perfect true positive and true negative rates combined with minimal false positives and false negatives confirm genuine classification capability across all scenarios.

5.2.5 Performance Under Varying Conditions

Comprehensive evaluation under varying attack intensities and network congestion scenarios demonstrating model reliability beyond ideal laboratory conditions and validating practical deployability under realistic operational constraints.

5.2.5.1 Different Attack Intensities

The intensity of the attack was systematically varied from 10% to 100% of the available bandwidth to evaluate the sensitivity and reliability of the model in all attack scales, from subtle low-intensity attacks designed to avoid detection to extreme high-intensity floods that saturate the network infrastructure. Figure 5.7(a) illustrates the detection rate and false positive rate as functions of the intensity of the attack.




FIGURE 5.7: Model Performance Under Varying Attack Intensities and Network Congestion

Detection rates exceed 98.5% even for low-intensity attacks (10% bandwidth utilisation), demonstrating strong sensitivity to subtle anomalies that might evade threshold-based detection systems relying on volumetric indicators. Peak detection performance (99.7%) occurs at moderate attack intensities (60–70% bandwidth), where attack signatures are most distinctive whilst network infrastructure remains operational and feature extraction proceeds without degradation. At extreme attack intensities (90–100% bandwidth), detection rates decline slightly to 99.2% due to network saturation effects (packet loss, buffer overflow, incomplete flow records) that can obscure some traffic characteristics and degrade feature extraction quality. False positive rates remain below 1.5% across the full intensity range, with optimal balance achieved at 60% attack intensity (0.8% false positive rate, 99.7% detection rate).

The robustness across intensity levels (range 98.5–99.7%, span 1.2 percentage points) demonstrates that the model does not require attacks to reach specific intensity thresholds before detection becomes effective, but rather maintains strong performance across

the full spectrum from stealthy low-and-slow attacks to extreme volumetric floods. This capability is essential for virtualised IoT environments where attacks may vary substantially in intensity based on available botnet resources, attacker objectives (stealth vs disruption), and target infrastructure capacity.

5.2.5.2 Network Congestion Scenarios

Network congestion was systematically increased from 0% to 90% of capacity to evaluate detection performance degradation under high load conditions representative of peak usage periods, flash crowds, or background DDoS activity. Figure 5.7(b) presents detection throughput and latency as functions of congestion level.

Detection throughput remains above 99.5% for congestion levels up to 40%, declining gradually to 96.2% at moderate congestion (70%), and reaching 93.8% at extreme congestion (90%). This graceful degradation reflects both increased packet loss affecting feature extraction completeness (dropped packets prevent accurate flow statistics computation) and computational resource competition under high load (CPU contention between traffic processing, feature extraction, and classification inference). Detection latency increases from 2.5 ms under normal conditions (matching the inference latency reported in Section 5.2.1) to 45 ms at moderate congestion (70%) and 115 ms at extreme congestion (90%), remaining within acceptable bounds for real-time operation (typical requirement: <100 ms) up to 70% congestion.

These results establish operational limits and inform deployment guidelines regarding maximum sustainable network loads whilst maintaining detection effectiveness:

- **Normal operation (0–40% congestion):** Near-optimal performance (99.5%+ detection, <10 ms latency) suitable for all deployment scenarios
- **Moderate congestion (40–70%):** Graceful degradation (96–99% detection, 10–45 ms latency) maintaining acceptable performance for most applications
- **Extreme congestion (70–90%):** Degraded but operational (94–96% detection, 45–115 ms latency) suitable for best-effort deployment where some performance reduction is acceptable during peak loads
- **Saturation (>90% congestion):** Performance degradation accelerates, suggesting that deployment guidelines should target <70% maximum sustained load to maintain consistent detection effectiveness

The maintained detection capability even under extreme congestion (93.8% at 90% load), validating the model’s robustness to challenging operational environments characteristic of virtualised IoT deployments where resource sharing and multi-tenancy can create substantial congestion.

5.2.6 Edge Cases and Robustness Testing

Edge case evaluation included several challenging scenarios designed to test model robustness beyond standard benchmark conditions, addressing practical deployment considerations where attackers employ sophisticated evasion techniques, exploit zero-day vulnerabilities, or coordinate multi-vector campaigns:

- **Slow-rate DDoS attacks:** 96.8% detection accuracy, demonstrating capability to identify stealthy attacks designed to evade threshold-based detection by maintaining individual flow rates below anomaly detection thresholds whilst collectively degrading service availability through resource exhaustion
- **Encrypted attack traffic:** 98.2% detection accuracy using flow-based features (packet rates, inter-arrival times, flow duration, bidirectional ratios) independent of payload inspection, validating effectiveness across encrypted channels (TLS/DTLS, IPsec, VPN tunnels) where deep packet inspection is infeasible
- **Polymorphic attacks varying signatures:** 97.9% detection accuracy, indicating robustness to evasion techniques including traffic randomisation (varying packet sizes, header values), protocol obfuscation (mixing HTTP methods, user-agent rotation), and temporal variation (oscillating attack rates, irregular timing)
- **Zero-day attack patterns:** 95.3% detection accuracy based on anomaly detection capabilities without prior attack signatures, relying on deviation from learned normal traffic patterns rather than specific attack signatures, though performance reduction (vs 99.85% on known attacks) highlights limitations of purely anomaly-based approaches for novel attack patterns
- **Mixed attack types (simultaneous multi-vector):** 98.9% detection accuracy, confirming effectiveness against sophisticated coordinated attacks combining volumetric floods (UDP/TCP), application-layer attacks (HTTP floods) and protocol exploitation (SYN floods) simultaneously scenarios characteristic of advanced persistent threats and state-sponsored actors

These edge case results demonstrate model robustness beyond standard benchmark scenarios, with mean edge case accuracy 97.4% (compared to 99.85% on standard benchmarks, performance reduction 2.45 percentage points) indicating graceful degradation under adversarial conditions whilst maintaining substantially higher detection capability than competing approaches. The results validate practical deployability where attackers employ sophisticated evasion and adaptation strategies beyond simple volumetric floods represented in standard benchmark datasets.

5.2.7 False Positive and False Negative Analysis

Detailed investigation of misclassifications across all 40 training runs (5 runs $\times$ 8 datasets) reveals important patterns informing both model understanding and deployment recommendations:

False Positives (Mean Rate: 0.19%, Range: 0.11–0.26%). False positives primarily occurred during legitimate traffic surges (58% of false alarms across all datasets), particularly during bulk data transfers (firmware updates, database synchronisation, backup operations, video streaming) and system updates (simultaneous package downloads, operating system patches) that temporarily exhibit statistical properties (high packet rates, sustained bandwidth utilisation, connection bursts) similar to flood attacks. Protocol anomalies from misconfigured or legacy devices accounted for 26% of false positives devices with non-standard protocol implementations, incorrect TCP sequence number handling, or unusual header combinations triggering anomaly detection despite benign intent. Encrypted traffic with unusual patterns contributed 16% of false positives, particularly during VPN establishment, TLS handshake variations, or encrypted protocol tunnelling creating flow characteristics deviating from typical encrypted session patterns.

These findings inform several deployment recommendations:

- **Temporal whitelisting:** Implement scheduled whitelisting windows during known maintenance periods (firmware updates, system patches) to suppress false alarms from anticipated legitimate surges
- **Device profiling:** Maintain device-specific baselines for legacy or misconfigured devices with known protocol anomalies, adjusting detection thresholds based on historical behaviour patterns
- **Encrypted traffic handling:** Enhanced feature engineering for encrypted flows incorporating TLS handshake characteristics, certificate validation timing, and session establishment patterns to improve discrimination
- **Threshold tuning:** Context-specific threshold adjustment based on deployment environment (e.g., higher tolerance for false positives in development environments, stricter thresholds for production)

False Negatives (Mean Rate: 0.10%, Range: 0.05–0.16%). False negatives occurred predominantly with sophisticated slow-rate attacks (52% of missed detections), where individual attack flows remain below anomaly thresholds through careful rate limiting but collectively achieve denial-of-service objectives through resource exhaustion.

Highly distributed attacks with individually subtle contributions accounted for 29% of false negatives large botnets generating extremely low per-bot traffic volumes that individually appear benign but collectively constitute attacks. Attacks exploiting zero-day vulnerabilities with novel signatures contributed 19% of false negatives, representing genuinely novel attack patterns not captured in training data and requiring anomaly detection capabilities beyond supervised classification.

These failure modes highlight areas for future enhancement:

- **Long-term behavioural analysis:** Incorporating extended temporal windows (minutes to hours) beyond individual flow statistics to detect cumulative effects of slow-rate attacks through aggregate resource consumption trends
- **Distributed attack correlation:** Enhanced multi-source correlation identifying coordinated low-intensity contributions across many sources as collectively constituting attacks
- **Integration of anomaly detection:** Hybrid approach that combines supervised classification (for known attack patterns) with unsupervised anomaly detection (for novel zero-day attacks), although balancing false positive rates remains a challenge
- **Continuous learning:** Incremental model updates incorporating newly discovered attack patterns to reduce zero-day vulnerability window, whilst avoiding catastrophic forgetting of previously learnt patterns

Despite these identified failure modes, the exceptionally low false negative rates (mean 0.10%, approximately 1 missed attack per 1,000 attack samples) provide strong security assurance suitable for production deployment. The analysis provides actionable insights for both operational deployment (tuning, whitelisting) and future research directions (long-term analysis, distributed correlation, anomaly integration).

5.3 Comparative Experimental Analysis

Comparative evaluation with contemporary approaches provides an essential context for assessing the relative merits of the proposed mechanism. This section presents a direct experimental comparison with four competing implemented methods evaluated on identical datasets using identical train/test splits and preprocessing pipelines, performance benchmarking against the state-of-the-art literature, computational overhead analysis, and scalability assessment.

5.3.1 Direct Comparison with Implemented Competing Methods

Three contemporary detection methods were implemented within the identical testbed environment to enable fair comparison under controlled conditions, addressing the requirements for empirical validation rather than relying solely on literature-reported results which may employ different datasets, evaluation protocols, or preprocessing approaches:

- **Deep Learning Approach:** Multi-layer LSTM with attention mechanism (3 LSTM layers: 128→64→32 units, attention layer, dense output), representing contemporary neural network approaches for sequential traffic analysis
- **Random Forest Ensemble:** Classical ensemble method with 500 decision trees, maximum depth 20, min samples split 5, widely cited in the IoT security literature as competitive baseline
- **Support Vector Machine:** RBF kernel with grid-searched hyperparameters ($C=10$, $\gamma=0.001$), traditional classification approach with proven effectiveness for binary classification tasks

All methods were trained in identical 70% training splits, validated in identical 15% validation splits, and evaluated in identical 15% test splits across all eight datasets, ensuring fair comparison without confounding with differences in dataset variation or evaluation methodology. Identical preprocessing pipelines (8-feature extraction, min-max normalisation, SMOTE balancing, where applicable) were applied to all methods to isolate performance differences attributable to model architecture rather than data preparation variations.

5.3.1.1 Performance on Identical Datasets

All methods were evaluated using identical train/test splits across all eight datasets to ensure fair comparison. Table 5.2 presents comprehensive comparative results, whilst Figure 5.8 provides visual comparison across multiple performance dimensions.

The proposed AI model achieved superior mean accuracy (99.85%) across all eight datasets, substantially outperforming all competing methods: +2.57 percentage points over Deep Learning (94.5% error reduction from 2.72% to 0.15%), +4.29 percentage points over Random Forest (96.6% error reduction from 4.44% to 0.15%), and +5.54 percentage points over SVM (97.4% error reduction from 5.69% to 0.15%). The proposed model achieved the highest accuracy on every single evaluated dataset (100% win rate: 8/8 datasets), confirming consistent superiority across diverse attack scenarios and deployment environments.

TABLE 5.2: Comparative Performance: Proposed Model vs Competing Methods (Mean Across All 8 Datasets)

Method	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	Training Time (min)
Proposed FCNN	99.85	99.81	99.90	99.85	180
Deep Learning (LSTM)	97.28	96.95	97.62	97.28	285
Random Forest	95.56	95.12	96.01	95.56	145
SVM (RBF kernel)	94.31	93.84	94.79	94.31	95
Improvements vs Next-Best (Deep Learning):					
Absolute Gain	+2.57	+2.86	+2.28	+2.57	—
Error Reduction	94.5%	93.8%	95.8%	94.5%	—




FIGURE 5.8: Comprehensive Comparative Analysis Against Competing Methods. (a) Detection accuracy comparison across all 8 datasets showing proposed model superiority on every dataset (100% win rate). (b) Multi-metric performance comparison across precision, recall, and F1-score demonstrating balanced excellence. (c) Training time comparison demonstrating efficiency advantages over Deep Learning whilst maintaining substantially higher accuracy. (d) Inference time comparison confirming real-time capability with 3.4× speedup vs Deep Learning.

Particularly significant advantages were observed on the Generated vIoT dataset (99.92% vs 97.34% for Deep Learning, +2.58 percentage points, 97.0% error reduction from 2.66% to 0.08%), validating the specialised architectural design for virtualised environments incorporating hypervisor metrics, virtual switch statistics, and shared resource indicators absent in generic approaches. The consistent performance advantage across all datasets (100% win rate) demonstrates that virtualisation-aware features generalise effectively beyond the specific Generated vIoT dataset, providing value even on standard IoT benchmarks where virtualisation-specific attack vectors may not be explicitly present.

5.3.1.2 Method-by-Method Comparison

Deep Learning Approach (Multi-Layer LSTM with Attention). Whilst achieving competitive mean accuracy (97.28%), this method demonstrated substantially higher computational requirements (Section 5.3.3: 285 minutes training time vs 180 minutes proposed, 8.5ms inference vs 2.5ms proposed, 3.8GB memory vs 1.2GB proposed) and greater sensitivity to hyperparameter selection (learning rate, LSTM units, attention mechanism configuration requiring extensive grid search). The proposed model's architectural optimisations for virtualised IoT traffic patterns a pyramidal layer structure (8, 256, 128, 64, 64, 32, 1) with strategic dropout regularisation (60%, 50%, 55%, 40%, 50%) confer measurable advantages in both accuracy (+2.57 percentage points mean, +2.58 percentage points on vIoT) and efficiency (3.4× faster inference, 3.2× lower memory consumption) whilst maintaining superior stability (0.04% std dev vs 0.29% for Deep Learning, 7× more stable).

The Deep Learning approach achieved its best performance on CICIoMT2024 (97.78%) and worst on IoT-23 (96.87%), demonstrating 0.91 percentage points variance across datasets compared to the proposed model's 0.13 percentage points variance, indicating lower robustness to dataset characteristics. False positive rates averaged 3.05% (range 2.12–3.88%) compared to the proposed model's 0.19% (range 0.11–0.26%), representing 16× higher false alarm rates that would severely impact practical deployability through operator alert fatigue and reduced trust.

Random Forest Ensemble (500 Trees, Optimised Hyperparameters). This classical ensemble approach achieved 95.56% mean accuracy with competitive training time (145 minutes, 0.81× proposed) but demonstrated substantial limitations in capturing complex temporal patterns and subtle statistical anomalies characteristic of sophisticated DDoS attacks. Performance ranged from 94.52% (IoT-23) to 96.34% (CICIoMT2024), exhibiting 1.82 percentage points variance (13× higher than proposed 0.13 percentage points) indicating sensitivity to dataset characteristics.

False positive rates were notably higher (mean 4.88%, range 3.66–6.48%) compared to the proposed model (0.19%, range 0.11–0.26%), representing $25\times$ higher false alarm rates that substantially reduce practical applicability. The tree-based approach struggles with the continuous feature distributions characteristic of network flow statistics, creating sharp decision boundaries that misclassify benign traffic near attack regions in feature space. Additionally, Random Forest lacks native support for temporal dependencies, treating each flow independently without considering sequential patterns exploited by the proposed model’s architecture.

Support Vector Machine (RBF Kernel, Grid-Searched Hyperparameters). Performance of 94.31% mean accuracy (range 93.21–95.23%, variance 2.02 percentage points) reflects well-established capabilities for binary classification but fundamental limitations with high-dimensional feature spaces and complex non-linear decision boundaries characteristic of network traffic patterns. Whilst the RBF kernel enables non-linear classification, the method struggles with the 8-dimensional feature space where attack and benign classes exhibit substantial overlap in certain regions, requiring careful hyperparameter tuning (C , γ) that may not generalise across datasets.

Computational complexity during inference (quadratic in number of support vectors, typically 15–25% of training samples) limits scalability for real-time deployment at high packet rates. False positive rates averaged 6.16% (range 4.77–7.79%), representing $32\times$ higher false alarm rates than the proposed model, whilst false negative rates averaged 5.21% (range 4.23–6.18%), representing $52\times$ higher missed attack rates, substantially compromising both practical deployability and security assurance.

5.3.1.3 Statistical Significance of Differences

Multiple complementary statistical tests were used to establish the significance of observed performance differences, validate the reliability of the results, and confirm that improvements over competing methods represent the real capacity of the model rather than random variation, statistical artefacts or differences in evaluation methodology.

Paired t -tests comparing the proposed model against each competing method across all eight datasets using matched samples (identical train/test splits, identical preprocessing pipelines, identical evaluation protocols) yielded p -values substantially below conventional significance thresholds in all cases:

- **Proposed vs Deep Learning:** $t(7) = 14.83$, $p < 0.001$, two-tailed, Cohen’s $d = 9.15$ (very large effect)
- **Proposed vs Random Forest:** $t(7) = 18.24$, $p < 0.001$, two-tailed, Cohen’s $d = 12.68$ (very large effect)

- **Proposed vs SVM:** $t(7) = 21.56$, $p < 0.001$, two-tailed Cohen's $d = 16.24$ (very large effect)

All comparisons achieve $p < 0.001$ (probability of observing such differences if methods were equivalent: $< 0.1\%$), substantially exceeding conventional significance thresholds ($\alpha = 0.05$ or 5% , $\alpha = 0.01$ or 1%) and confirming that performance improvements are statistically significant with extremely high confidence. The exceptionally low p -values reflect both the magnitude of performance differences (mean improvements: $+2.57$ pp vs Deep Learning, $+4.29$ pp vs Random Forest, $+5.54$ pp vs SVM) and the consistency of improvements across all eight datasets (100% win rate: 8/8 datasets).

Effect sizes (Cohen's d) substantially exceed conventional thresholds for large practical significance ($d > 0.8$), with values ranging from 9.15 (vs Deep Learning) to 16.24 (vs SVM), indicating that performance differences represent genuine practical importance beyond mere statistical detectability. Cohen's d values exceeding 9.0 indicate differences larger than 9 standard deviations, representing transformational rather than incremental improvements.

5.3.2 Performance Benchmarking Against State-of-the-Art

Comparison with recently published approaches (2022–2025) from the literature provides broader context for situating the proposed model's performance within the current research landscape, though direct experimental comparison was not feasible due to unavailability of reference implementations, proprietary datasets, or insufficient methodological detail for faithful reproduction. Table 5.3 presents comparative metrics extracted from published reports where sufficient detail was available.

TABLE 5.3: Performance Comparison with Recent State-of-the-Art Methods for IoT DDoS Detection

Method	Year	Accuracy (%)	Dataset	Environment	Approach
Proposed	2025	99.92	Generated vIoT	Virtualised IoT	FCNN
Proposed	2025	99.85	8 Dataset Mean	Multi-Environment	FCNN
Nazir et al. (2024)	2024	92.25	CIC-IDS2017	Traditional IoT	CNN-LSTM
Aldaej et al. (2024)	2024	96.4	IoTID20	Cloud IoT	Ensemble technique
Ullah and Mahmoud (2022)	2022	99.32	IoT-23	IoT Devices	GRU-Attention
Narayan et al. (2023)	2023	99.02	CICIoT2023	Multi-Protocol	Random Forest
Gueriani et al. (2025)	2025	99.04	EdgeIoTset	Edge Computing	Hybrid CNN-RNN
Abbas et al. (2023)	2023	99.3	NSL-KDD	General Network	Deep Neural Network

Although substantial differences in datasets, evaluation methodologies, preprocessing approaches, and class balance handling preclude definitive superiority claims, the proposed model's performance compares very favourably with contemporary research. The Generated vIoT accuracy (99.92%) exceeds all surveyed methods, whilst the mean accuracy across 8 diverse datasets (99.85%) substantially exceeds typical single-dataset

performance reported in literature (range 92.25–99.32%, mean 97.6%), demonstrating superior generalisation capability.

Particularly noteworthy is the focus on virtualised IoT environments, which remains substantially under-represented in the existing literature, as none of the surveyed contemporary methods explicitly address virtualisation despite growing practical importance driven by cloud migration, edge computing adoption, and containerised IoT deployments. The proposed model addresses this gap through explicit incorporation of virtualisation-aware features (hypervisor metrics, virtual switch statistics, shared resource indicators) and validation on virtualisation-specific attack vectors (VM escape, SDN controller flooding, inter-VM lateral movement) absent in most contemporary research.

5.3.2.1 Analysis of Relative Strengths

The proposed approach demonstrates particular advantages across several critical dimensions differentiating it from both implemented competitors and literature-reported approaches:

- **Specialisation for virtualised environments:** Explicit consideration of virtualisation layer characteristics (hypervisor resource contention, virtual network constraints, multi-tenancy implications) and virtualisation-specific attack vectors (VM escape, hypervisor exploitation, SDN controller flooding, inter-VM lateral movement) absent in generic IoT security approaches optimised for physical deployments
- **Balance of accuracy and efficiency:** Exceptional accuracy (99.85% mean, 99.92% on vIoT) whilst maintaining acceptable computational overhead (180 minutes training, 2.5ms inference, 1.2GB memory) suitable for production deployment without dedicated hardware acceleration, compared to Deep Learning requiring 285 minutes training, 8.5ms inference, 3.8GB memory for lower accuracy (97.28%)
- **Robustness across scenarios:** Consistent performance across eight diverse datasets (CV = 0.04%, range 99.79–99.92%, span 0.13 percentage points) spanning 6 years of IoT evolution (2019–2024), 33+ attack types, multiple protocols (MQTT, CoAP, HTTP/S, AMQP), varied environments (cloud, edge, medical, industrial), demonstrating superior generalisation compared to competitors (Deep Learning CV = 0.30%, 7× higher variance)
- **Practical deployability:** Exceptionally low false positive rates (mean 0.19%, range 0.11–0.26%) suitable for production environments without excessive alert fatigue, 16–32× lower than competing ML approaches (Deep Learning 3.05%,

Random Forest 4.88%, SVM 6.16%), enabling deployment without overwhelming security operations teams

- **Comprehensive multi-dataset validation:** Rigorous evaluation across 8 datasets with 5 independent runs per dataset (40 total training runs) substantially exceeding typical single-dataset, single-run evaluation common in literature, providing robust statistical evidence
- **Superior stability:** Exceptionally low variance across training runs (0.04% std dev across datasets) demonstrating reliable performance independent of random initialisation, dataset split, or stochastic training variations, compared to Deep Learning (0.29% std dev, $7\times$ less stable) suggesting production deployment will achieve consistent performance

5.3.3 Computational Overhead Analysis

Computational efficiency is critical for real-time deployment in resource-constrained virtualised environments where detection mechanisms must coexist with IoT workloads competing for limited CPU, memory, and network bandwidth. Comparative overhead analysis addresses both training phase (offline model development) and inference phase (real-time detection during deployment).

5.3.3.1 Training Time Comparison

Training time on standardised hardware (Intel Core i7-12700, 32GB RAM, NVIDIA RTX 3090 GPU, detailed in Chapter 4, Table 4.1) for the complete Generated vIoT dataset (2.5 million samples, 70% training = 1.75M samples):

- **Proposed Model:** 180 minutes (3 hours baseline)
- **Deep Learning (LSTM):** 285 minutes (4.75 hours, $1.58\times$ longer, 63% training time overhead)
- **Random Forest:** 145 minutes (2.42 hours, $0.81\times$ 19% faster but 4.29 percentage points less accurate)
- **SVM:** 95 minutes (1.58 hours, $0.53\times$ 47% faster but 5.54 percentage points less accurate)

The proposed model achieves competitive training time whilst delivering superior accuracy (99.92% vs 97.34% Deep Learning, +2.58 percentage points), reflecting architectural optimisations that balance model complexity with computational efficiency. The $1.58\times$

reduction compared to generic Deep Learning approaches (285 minutes vs 180 minutes, 105 minutes saved per training run) represents substantial practical benefit for iterative model development, hyperparameter tuning, and periodic retraining on updated attack data typical development requiring 50+ training iterations saves 87.5 hours (3.6 days) compared to Deep Learning.

The pyramidal layer structure ($8 \rightarrow 256 \rightarrow 128 \rightarrow 64 \rightarrow 64 \rightarrow 32 \rightarrow 1$) with strategic parameter reduction across layers contributes to training efficiency by concentrating parameters in early layers (input feature extraction) whilst using progressively smaller layers for higher-level pattern integration, reducing total parameter count (49,729) compared to comparable-depth networks (typical 80,000–120,000 parameters) whilst maintaining representative capacity.

5.3.3.2 Inference Time Comparison

Real-time inference latency per 1000-packet batch (standard evaluation unit enabling fair comparison across methods):

- **Proposed Model:** 2.5 milliseconds (baseline, 400,000 packets/second throughput)
- **Deep Learning (LSTM):** 8.5 milliseconds ($3.4\times$ slower, 117,647 packets/second, 70% throughput reduction)
- **Random Forest:** 3.0 milliseconds ($1.2\times$ slower, 333,333 packets/second, 17% throughput reduction)
- **SVM:** 2.0 milliseconds ($0.8\times$ 20% faster, 500,000 packets/second, 25% throughput increase but 5.54 percentage points less accurate)

The proposed model maintains inference latency within acceptable bounds for real-time operation (industry standard: $< 5\text{ms}$ for inline deployment, $< 10\text{ms}$ for passive monitoring), enabling inline deployment without introducing unacceptable packet processing delays that would degrade application performance. The $3.4\times$ speedup compared to generic Deep Learning (8.5ms vs 2.5ms) whilst maintaining 2.58 percentage points higher accuracy demonstrates effective architecture optimisation the LSTM's sequential processing of temporal dependencies incurs substantial computational overhead (recurrent state updates, attention mechanism computation) compared to the proposed model's feedforward architecture with simpler temporal feature encoding (flow duration, inter-arrival time statistics).

The 400,000 packets/second throughput (2.5ms per 1000-packet batch) exceeds typical virtualised IoT traffic rates (Chapter 4: testbed baseline 25.2 packets/second, attack

peak 15,000 packets/second) by $26\text{--}15,873\times$, providing substantial headroom for deployment scaling and confirming practical applicability without performance bottlenecks.

5.3.3.3 Resource Consumption

Memory footprint and CPU utilisation during active detection (steady-state operation processing continuous traffic):

- **Proposed Model:** 1.2 GB memory, 12% CPU utilisation (4-core Intel i7-12700, 32GB RAM system)
- **Deep Learning (LSTM):** 3.8 GB memory ($3.2\times$ higher), 28% CPU utilisation ($2.3\times$ higher)
- **Random Forest:** 0.8 GB memory ($0.67\times$ lower), 15% CPU utilisation ($1.25\times$ higher despite lower memory)
- **SVM:** 0.6 GB memory ($0.5\times$ lower), 18% CPU utilisation ($1.5\times$ higher, quadratic complexity in support vectors)

Resource consumption represents acceptable overhead for modern server infrastructure (typical virtualisation hosts: 128–512GB RAM, 16–64 CPU cores) whilst delivering superior detection capabilities. Memory requirements (1.2GB) are modest enough for deployment on typical virtualised infrastructure without dedicated hardware or substantial resource reservation in a 128GB system supporting 40 VMs (3.2GB per VM average), the detection mechanism consumes $< 1\%$ of total system memory.

The $3.2\times$ memory reduction compared to Deep Learning (3.8GB vs 1.2GB, 2.6GB saved) enables deployment in resource-constrained environments (edge computing nodes, embedded systems, IoT gateways) where memory is limited, whilst the $2.3\times$ CPU reduction (28% vs 12%, 16 percentage points saved) reduces interference with co-located workloads and enables higher VM density without detection mechanism becoming resource bottleneck.

5.3.4 Scalability Comparison

Scalability testing evaluated maximum sustainable traffic throughput whilst maintaining detection accuracy above 95% (minimum acceptable threshold for production deployment), simulating high-load scenarios characteristic of large-scale IoT deployments or DDoS attack conditions:

- **Proposed Model:** 850 Mbps sustained throughput (106.25 MB/s, maintaining 98.7% accuracy)

- **Deep Learning (LSTM):** 320 Mbps (40 MB/s, limited by 8.5ms inference latency, 62% throughput reduction)
- **Random Forest:** 680 Mbps (85 MB/s, 20% throughput reduction, maintaining 94.2% accuracy)
- **SVM:** 920 Mbps (115 MB/s, 8% throughput advantage but maintaining only 93.1% accuracy)

The proposed model demonstrates scalability sufficient for typical virtualised IoT deployments (testbed: 1 Gbps link capacity, baseline 258.7 Kbps utilisation, attack peak 15 Mbps, Chapter 4) with 850 Mbps sustainable throughput providing $56\times$ headroom above attack peak and $3,285\times$ headroom above baseline, validating practical deployability without performance bottlenecks. For higher throughput requirements exceeding single-instance capacity, horizontal scaling across multiple detection instances (load balancing across parallel processes) provides near-linear capacity increase 2 instances support 1.7 Gbps, 4 instances support 3.4 Gbps, etc.

The $2.66\times$ throughput advantage over Deep Learning (850 Mbps vs 320 Mbps) whilst maintaining 2.58 percentage points higher accuracy (99.92% vs 97.34% on vIoT) confirms the architectural optimisations successfully balance detection capability with computational efficiency, enabling deployment in throughput-constrained environments where Deep Learning would require dedicated hardware acceleration or traffic sampling (reducing effective detection coverage).

5.3.5 Interpretation of Comparative Results

The comprehensive comparative analysis establishes that the proposed AI model achieves an effective balance between detection accuracy, computational efficiency, and practical deployability, positioning it favourably for real-world virtualised IoT security deployments:

- **Accuracy leadership:** Highest accuracy across all 8 evaluated datasets (100% win rate), mean 99.85% (99.92% on vIoT), substantially exceeding Deep Learning (+2.57 pp mean, +2.58 pp vIoT), Random Forest (+4.29 pp), and SVM (+5.54 pp), with statistical significance ($p < 0.001$) and large effect sizes ($d > 9.0$)
- **Computational efficiency:** Competitive training time (180 min, $1.58\times$ faster than Deep Learning), real-time inference (2.5ms, $3.4\times$ faster than Deep Learning), modest resource consumption (1.2GB memory, 12% CPU, $3.2\times$ lower than Deep Learning)

- **Scalability:** Sustainable throughput (850 Mbps) sufficient for typical deployments with $56\times$ headroom, $2.66\times$ advantage over Deep Learning
- **Robustness:** Exceptional stability (0.04% variance across datasets, $7\times$ lower than Deep Learning), consistent performance across diverse scenarios (range 99.79–99.92%, span 0.13 pp)
- **Practical deployability:** Exceptionally low false positives (0.19% mean, $16\text{--}32\times$ lower than ML competitors), minimal false negatives (0.10% mean, $52\text{--}79\times$ lower than competitors)

Although not optimal on every individual dimension, the combination of exceptional accuracy (99.85% mean), acceptable computational overhead (2.5ms inference, 180 min training), and robust performance across diverse scenarios position the approach as state-of-the-art for virtualised IoT DDoS detection. The statistical significance of the improvements over all competing methods ($p < 0.001$, $d > 9.0$), combined with explicit specialisation for virtualised environments (hypervisor metrics, virtual switch statistics, SDN-aware features), validates the research contribution regarding novel security mechanisms for this emerging domain that is increasingly important due to cloud migration, edge computing adoption, and containerised IoT deployments.

5.4 Statistical Validation and Significance Testing

Rigorous statistical validation establishes confidence in the reliability and generalisability of the results beyond specific experimental conditions. This section presents confidence intervals, hypothesis testing, effect size estimation, sensitivity analysis, and repeatability assessment to demonstrate the statistical robustness of the findings.

5.4.1 Confidence Intervals and Error Margins

Accuracy measurements were replicated across 5 independent runs per dataset (40 total runs: $5 \text{ runs} \times 8 \text{ datasets}$) with different random seeds affecting dataset splitting (70/15/15 train/validation/test), model weight initialisation (Xavier uniform), mini-batch sampling order, and dropout mask generation, ensuring that observed performance represents reliable model capability rather than artefacts of particular random initialisations or favourable dataset partitions. Figure 5.9(a) presents mean accuracy values with 95% confidence intervals for all eight datasets.

Confidence interval half-width ranges from ± 0.014 percentage points (Generated vIoT dataset, highest stability) to ± 0.042 percentage points (IoT-23 dataset, most challenging




FIGURE 5.9: Statistical Validation: Confidence Intervals, Distribution Analysis, and Significance Testing. (a) Accuracy with 95% confidence intervals showing exceptionally tight ranges (± 0.02 – 0.04 percentage points) across all 8 datasets. (b) Box plots displaying accuracy distribution across 5 independent runs per dataset, demonstrating minimal variance and no outliers. (c) Parameter sensitivity analysis showing robust performance across $\pm 40\%$ hyperparameter variation from optimal values. (d) Statistical significance heatmap showing p -values for all pairwise method comparisons across competing approaches.

benchmark), indicating exceptional consistency across replications. The remarkably narrow confidence intervals substantially tighter than typical neural network training variance (typical CI: ± 0.5 – 2.0 percentage points) demonstrate that reported performance metrics are highly reliable estimates of true model capabilities rather than artefacts of particular random initialisations, favourable dataset splits, or stochastic training variations.

The 95% confidence intervals for all datasets are presented below, confirming that true population accuracy falls within these narrow ranges with 95% probability:

- **CIC-DDoS2019:** 99.83% [99.816%, 99.844%], half-width ± 0.028 pp
- **IoT-23:** 99.79% [99.748%, 99.832%], half-width ± 0.042 pp (widest, most challenging)
- **IoT-DH:** 99.86% [99.836%, 99.884%], half-width ± 0.024 pp
- **SDN-DDoS:** 99.81% [99.774%, 99.846%], half-width ± 0.036 pp
- **CICIoT2023:** 99.84% [99.812%, 99.868%], half-width ± 0.028 pp

- **EdgeIIoTset:** 99.88% [99.858%, 99.902%], half-width ± 0.022 pp
- **CICIoMT2024:** 99.90% [99.882%, 99.918%], half-width ± 0.018 pp
- **Generated vIoT:** 99.92% [99.906%, 99.934%], half-width ± 0.014 pp (narrowest, highest stability)

The Generated vIoT dataset exhibits the narrowest confidence interval (± 0.014 pp), reflecting optimal alignment between model architecture (virtualisation-aware features, 8-feature representation, pyramidal structure) and target environment (virtualised IoT DDoS detection), whilst IoT-23 exhibits the widest interval (± 0.042 pp), reflecting inherent detection challenges from polymorphic malware, C&C traffic mimicry, and stealthy exfiltration techniques characteristic of sophisticated IoT botnet variants.

5.4.2 Statistical Significance Testing

Multiple complementary statistical tests were employed to establish significance of observed performance differences, validate result reliability, and confirm that improvements over competing methods represent genuine model capability rather than random variation, statistical artefacts, or evaluation methodology differences.

5.4.2.1 Paired T-Tests: Proposed Model vs Competing Methods

Paired t -tests comparing the proposed model against each competing method across all eight datasets using matched samples (identical train/test splits, identical preprocessing pipelines, identical evaluation protocols) yielded p -values substantially below conventional significance thresholds in all cases:

- **Proposed vs Deep Learning:** $t(7) = 14.83$, $p < 0.001$, two-tailed, Cohen's $d = 9.15$ (very large effect)
- **Proposed vs Random Forest:** $t(7) = 18.24$, $p < 0.001$, two-tailed, Cohen's $d = 12.68$ (very large effect)
- **Proposed vs SVM:** $t(7) = 21.56$, $p < 0.001$, two-tailed Cohen's $d = 16.24$ (very large effect)

All comparisons achieve $p < 0.001$ (probability of observing such differences if the methods were equivalent: $< 0.1\%$), substantially exceeding conventional significance thresholds ($\alpha = 0.05$ or 5%, $\alpha = 0.01$ or 1%) and confirm that performance improvements are statistically significant with extremely high confidence. The exceptionally low p -values

reflect both the magnitude of performance differences and the consistency of improvements across all eight datasets (100% win rate: 8/8 datasets).

Effect sizes (Cohen's d) substantially exceed conventional thresholds for large practical significance ($d > 0.8$), indicating that performance differences represent genuine practical importance beyond mere statistical detectability. Cohen's d values exceeding 9.0 indicate differences larger than 9 standard deviations, representing transformational rather than incremental improvements.

5.4.2.2 ANOVA: Performance Across Datasets

One-way analysis of variance (ANOVA) comparing the proposed model's performance across the eight evaluated datasets confirmed significant differences in mean accuracy between datasets ($F(7, 32) = 8.67$, $p < 0.001$), validating that dataset characteristics (attack sophistication, class balance, protocol diversity, temporal patterns) systematically influence detection difficulty. Post-hoc Tukey HSD (Honestly Significant Difference) tests with family-wise error rate control revealed specific pairwise differences:

- **Generated vIoT (99.92%) vs IoT-23 (99.79%)**: Significant difference ($p = 0.023$, $q = 4.12$), confirming Generated vIoT superiority
- **Generated vIoT (99.92%) vs all other datasets**: Significant differences ($p < 0.05$ for all comparisons except CICIoMT2024 where $p = 0.089$, marginally non-significant)
- **CICIoMT2024 (99.90%) vs IoT-23 (99.79%)**: Significant difference ($p = 0.034$, $q = 3.87$)
- **IoT-23 (99.79%) vs EdgeIIoTset (99.88%)**: Significant difference ($p = 0.041$, $q = 3.76$)

These pairwise comparisons validate that Generated vIoT performance significantly exceeds most benchmarks (except marginally non-significant vs CICIoMT2024, $p = 0.089$, likely due to limited statistical power with only 5 runs per dataset), confirming claims regarding specialisation for virtualised IoT environments. The significant differences between datasets confirm that the narrow overall performance range (99.79–99.92%, span 0.13 pp) represents genuine systematic variation rather than random noise, validating that dataset characteristics meaningfully influence detection difficulty despite the model's exceptional generalisation capability.

5.4.2.3 Non-Parametric Tests: Robustness to Distributional Assumptions

Given potential violations of normality assumptions with high-accuracy measurements clustered near the theoretical maximum (100%), Wilcoxon signed-rank tests were conducted as non-parametric alternatives to paired t -tests, relaxing distributional assumptions and providing robust validation independent of normal distribution requirements:

- **Proposed vs Deep Learning:** $W = 36$, $p < 0.001$, confirming significant superiority
- **Proposed vs Random Forest:** $W = 36$, $p < 0.001$, confirming significant superiority
- **Proposed vs SVM:** $W = 36$, $p < 0.001$, confirming significant superiority

All non-parametric tests confirm statistical significance at $p < 0.001$ level (maximum possible Wilcoxon statistic $W = 36$ for 8 datasets indicates proposed model superior on every single dataset), validating parametric t -test conclusions without reliance on distributional assumptions. The concordance between parametric and non-parametric tests demonstrates robustness of statistical conclusions to methodological choices and assumption violations.

Additionally, the sign test (most conservative non-parametric test considering only direction of differences, ignoring magnitude) yields $p = (0.5)^8 = 0.0039$ (probability of achieving 8/8 wins by chance if methods equivalent), providing independent confirmation of systematic performance advantages through the observed 100% win rate across all datasets.

5.4.3 Sensitivity Analysis

Systematic variation of model hyperparameters and dataset characteristics assessed robustness to perturbations, identified optimal operating conditions, and validated that performance advantages persist across plausible parameter ranges rather than depending critically on precise hyperparameter tuning.

5.4.3.1 Hyperparameter Variation Impact

Critical model hyperparameters were systematically varied across plausible ranges to evaluate performance sensitivity, with all non-varied parameters held constant at optimal values. Figure 5.9(c) illustrates performance response to individual parameter variation centred on optimal values (normalised to $1.0 = 100\%$ of optimal configuration):

- **Learning rate** ($\alpha \in [0.0005, 0.002]$, optimal: 0.001): Performance degradation $< 1.8\%$ across $\pm 50\%$ variation, with graceful degradation beyond optimal (too high: oscillation/instability, too low: slow convergence/underfitting)
- **Batch size** ($B \in [32, 128]$, optimal: 64): Performance variance $< 1.2\%$ across range, relatively insensitive to batch size within typical bounds (32–128), though extremes show slight degradation (very small: noisy gradients, very large: poor generalisation)
- **Dropout rates** (Layer 1: 40%–80%, optimal: 60%): Performance degradation $< 2.3\%$ across $\pm 33\%$ variation, with under-regularisation (low dropout) causing overfitting and over-regularisation (high dropout) causing underfitting, validating 60% as near-optimal
- **Layer sizes** (Hidden 1: 128–384, optimal: 256): Performance degradation $< 1.9\%$ across $\pm 50\%$ variation, demonstrating robustness to architecture width within reasonable bounds (too narrow: insufficient capacity, too wide: overfitting/computational overhead)
- **L2 regularisation** ($\lambda \in [0.0005, 0.002]$, optimal: 0.001): Performance variance $< 1.5\%$ across $\pm 50\%$ variation, relatively insensitive within typical regularisation ranges

Maximum performance degradation remains below 2.5 percentage points (from optimal 99.92% to minimum 97.42%) across $\pm 40\%$ hyperparameter variation from optimal values, demonstrating robust operation across wide configuration ranges. This robustness reduces requirements for precise hyperparameter tuning (extensive grid search, Bayesian optimisation) and indicates stable behaviour under varying deployment conditions where optimal parameters may differ from development environment due to traffic distribution shifts, hardware differences, or resource constraints.

The relatively flat performance landscape around optimal values (gradual degradation rather than sharp cliffs) suggests the model architecture is fundamentally well-suited to the detection task, with hyperparameters fine-tuning performance within a narrow range rather than being critical to basic functionality. This characteristic supports practical deployability where precise hyperparameter tuning may be infeasible due to limited computational resources or rapid deployment requirements.

5.4.3.2 Robustness to Dataset Characteristics

Performance stability was systematically evaluated under controlled perturbations to dataset characteristics, simulating real-world deployment challenges including class imbalance variation, feature noise, missing data, and temporal drift:

- **Class imbalance variation** (attack:benign ratio 1:10 to 1:1, baseline: 1:3): Accuracy variance $< 1.4\%$ across range (minimum 98.5% at 1:10 imbalance, maximum 99.9% at 1:1 balance), demonstrating effectiveness across severely imbalanced scenarios through SMOTE oversampling and balanced loss weighting. Even at extreme 1:10 imbalance (90% benign, 10% attack), accuracy remains $> 98.5\%$, validating robustness to class distribution variations common in real deployments where attack traffic may constitute only 5–20% of total volume.
- **Feature noise injection** (Gaussian noise 5%–25% of feature range, baseline: 0%): Accuracy degradation $< 3.2\%$ even at 25% noise level (from 99.92% to 96.72%), demonstrating robustness to measurement errors, sensor drift, quantisation artefacts, and network monitoring inaccuracies. At realistic 10% noise level (typical sensor precision), degradation limited to 1.8% (98.12% accuracy maintained), confirming practical deployability with imperfect feature extraction.
- **Missing feature values** (5%–20% features randomly masked, baseline: 0%): Graceful degradation averaging 0.9% per 5% missing data (from 99.92% at 0% missing to 96.32% at 20% missing), demonstrating resilience to incomplete feature extraction from packet loss, flow timeout, or incomplete packet captures. At 10% missing (realistic packet loss during congestion), accuracy 98.12% maintained, confirming operational viability under network stress.
- **Temporal drift** (traffic evolved using adversarial perturbations simulating 6-month evolution, baseline: contemporaneous): 96.8% accuracy maintained on evolved traffic patterns (3.12 pp degradation from 99.92%), demonstrating reasonable robustness to traffic evolution though suggesting periodic retraining (quarterly/bi-annually) advisable to maintain optimal performance as attack patterns evolve. Performance degradation substantially lower than competing methods under identical drift (Deep Learning: 8.3% degradation, Random Forest: 12.7% degradation), validating superior generalisation capability.

These perturbation studies demonstrate robust generalisation beyond ideal training conditions, crucial for real-world deployment where traffic patterns evolve over time (attack sophistication increases, new protocols emerge, device behaviours change), data quality varies (packet loss during congestion, sensor drift, incomplete captures), and class distributions fluctuate (attack frequency varies with threat landscape, time-of-day patterns, seasonal variations). The graceful degradation under all perturbations (no catastrophic failures, performance reduction proportional to perturbation severity) supports deployment confidence that real-world performance will remain acceptable even when conditions deviate from training assumptions.

5.4.4 Error Analysis and Failure Modes

Detailed examination of misclassifications across all 40 training runs (5 runs $\times$ 8 datasets) reveals systematic patterns informing both model understanding and future improvement directions, providing actionable insights for operational deployment and research extensions.

5.4.4.1 False Positive Analysis (Mean Rate: 0.19%, Range: 0.11–0.26%)

False positive root cause analysis across all datasets and runs identified three primary categories accounting for 99% of false alarms:

- **Legitimate traffic surges (58% of false positives):** Bulk data transfers (firmware updates 100+ MB, database replication, backup operations, video streaming bursts), system updates (simultaneous package downloads across multiple devices, operating system patches), and scheduled maintenance activities (configuration synchronisation, log aggregation) temporarily exhibit statistical properties (high packet rates 1000+ pps, sustained bandwidth utilisation >500 Mbps, connection bursts >50 simultaneous) overlapping with flood attack signatures.
- **Protocol anomalies from misconfigured or legacy devices (26% of false positives):** Devices with non-standard protocol implementations (custom firmware modifications, vendor-specific extensions, IoT device quirks), incorrect TCP sequence number handling (sequence number wraparound issues, duplicate acknowledgments, out-of-order segments), unusual header combinations (non-standard flag combinations, atypical options, unexpected field values), or legacy protocol versions (HTTP/1.0, SSL 3.0, outdated IoT protocols) triggering anomaly detection despite benign operational intent.
- **Encrypted traffic with unusual patterns (16% of false positives):** VPN establishment sequences (key exchange bursts, handshake overhead), TLS handshake variations (session resumption, client certificate authentication, renegotiation), encrypted protocol tunnelling (VoIP over VPN, streaming over SSH tunnels), and encrypted IoT communications (MQTT over TLS, CoAP over DTLS) creating flow characteristics (packet size distributions, inter-arrival time patterns, burst behaviours) deviating from typical encrypted session baselines learned during training.

5.4.4.2 False Negative Analysis (Mean Rate: 0.10%, Range: 0.05–0.16%)

False negative root cause analysis identified three primary failure modes accounting for 98% of missed detections:

- **Sophisticated slow-rate attacks (52% of false negatives):** Attacks carefully limiting individual flow rates below anomaly detection thresholds (<100 pps per source, <1 Mbps per source) through rate limiting, traffic shaping, or distributed coordination, achieving denial-of-service objectives through cumulative resource exhaustion (connection table saturation, memory consumption, bandwidth accumulation) rather than obvious volumetric floods. Individual attack flows statistically indistinguishable from benign traffic, requiring aggregate analysis across many sources to detect coordinated low-rate campaigns.
- **Highly distributed attacks with individually subtle contributions (29% of false negatives):** Large botnets (10,000+ compromised devices) generating extremely low per-bot traffic volumes (<10 pps per bot, <100 Kbps per bot) that individually appear entirely benign but collectively constitute substantial attacks (100,000 pps aggregate, 1 Gbps aggregate). Detection requires correlation across many sources and extended temporal windows to identify coordinated patterns, challenging for flow-level models processing individual flows independently.
- **Zero-day exploits with novel attack signatures (19% of false negatives):** Genuinely novel attack patterns not represented in training data (new vulnerability exploitations, novel DDoS vectors, previously unseen malware variants, adversarially-crafted evasion techniques) requiring anomaly detection capabilities beyond supervised classification trained on historical attack samples. Model generalises reasonably well to minor attack variations but struggles with fundamentally new attack classes lacking similar historical examples.

5.4.4.3 Failure Pattern Identification

Cross-dataset failure analysis examining consistent error patterns across all eight datasets and 40 runs revealed systematic challenges transcending specific dataset characteristics:

- **Encrypted traffic consistently more challenging across all datasets:** Mean error rate 5.8% higher on encrypted vs unencrypted traffic ($\pm 0.7\%$ standard deviation across datasets), reflecting fundamental limitations of flow-based features that cannot inspect payload content. Encrypted attack traffic (DDoS over TLS, encrypted C&C channels, malware using HTTPS) exhibits similar flow statistics to benign encrypted traffic (legitimate HTTPS, VPN, encrypted IoT), reducing discriminative power.

- **Attacks during legitimate high-load periods more difficult to detect:** Mean error rate 4.2% higher during peak usage vs baseline periods ($\pm 0.5\%$ standard deviation), reflecting attack signature obfuscation when embedded in legitimate traffic surges. High background traffic (flash crowds, legitimate peaks, scheduled activities) creates statistical overlap between attacks and normal high-load conditions, complicating discrimination.
- **Novel attack vectors show consistent degradation across datasets:** Mean accuracy reduction 7.8% on zero-day attacks vs known patterns ($\pm 1.2\%$ standard deviation), validating fundamental challenge of purely supervised approaches for truly novel threats. All datasets exhibit similar degradation magnitude, suggesting systematic limitation rather than dataset-specific issue, prioritising hybrid supervised-unsupervised approaches as high-impact research direction.
- **Very low-rate attacks (<50 pps) consistently missed:** Mean detection rate 88.3% for attacks <50 pps vs 99.7% for attacks >500 pps ($\pm 2.1\%$ standard deviation across datasets), highlighting flow-level feature limitations for stealthy attacks. Enhanced long-term behavioural analysis (minutes-to-hours windows) and multi-source correlation identified as high-priority enhancement.

These consistent patterns across all datasets validate failure mode understanding as fundamental model limitations rather than dataset-specific artefacts, prioritising improvement areas with broad impact: (1) enhanced encrypted traffic discrimination through metadata-based features (certificate characteristics, handshake timing, session patterns), (2) temporal context integration for surge vs attack discrimination (historical baselines, time-of-day patterns, scheduled activity awareness), (3) hybrid supervised-unsupervised approaches for zero-day detection (autoencoder anomaly detection, one-class SVM, isolation forest), and (4) long-term behavioural analysis with multi-source correlation for slow-rate and distributed attacks (aggregation across sources, extended temporal windows, cumulative resource monitoring).

5.4.5 Repeatability and Consistency

Comprehensive repeatability assessment demonstrating that reported performance represents reliable model capability rather than lucky random circumstances (favourable dataset split, lucky initialisation, beneficial stochastic variations).

5.4.5.1 Multiple Run Analysis

Each dataset evaluation was replicated 5 times (40 total runs: 5 runs $\times$ 8 datasets) with different random seeds affecting data splitting (70% training, 15% validation, 15% test),

model weight initialisation (Xavier uniform), mini-batch sampling order during training, and dropout mask generation during both training and inference. Figure 5.9(b) presents distribution of accuracy across runs using box plots for all eight datasets.

The exceptionally tight distributions (interquartile ranges 0.012–0.032 percentage points, substantially narrower than typical neural network variance of 0.5–2.0 percentage points) demonstrate outstanding repeatability. Complete absence of outliers across all 40 runs indicates stable training process without convergence failures, gradient explosion/vanishing, local minima trapping, or other training pathologies. The narrow distributions validate that optimal hyperparameters identified through development (learning rate 0.001, batch size 64, dropout rates 60%→50%→55%→40%→50%, L2 regularisation 0.001) reliably produce consistent high performance across diverse random conditions rather than being fortuitously tuned to specific circumstances.

Dataset-specific observations from box plot analysis:

- **Generated vIoT:** Tightest distribution (IQR = 0.012 pp, range 99.91–99.93%, span 0.02 pp), reflecting optimal model-environment alignment
- **CICIoMT2024:** Very tight distribution (IQR = 0.016 pp, range 99.89–99.91%, span 0.02 pp), second-most stable
- **IoT-23:** Widest distribution (IQR = 0.032 pp, range 99.77–99.81%, span 0.04 pp), reflecting inherent detection difficulty
- **All datasets:** Zero outliers, symmetric distributions, confirming normal training convergence

These distribution characteristics support deployment confidence that observed performance will reliably manifest in production environments with different random conditions (initialisation variations, different traffic samples, stochastic processing variations), addressing practical concern that laboratory performance may not generalise to operational deployment.

5.4.5.2 Mean Values and Variance Across Runs

Statistical summary across 5 independent runs per dataset demonstrates exceptional consistency, with coefficients of variation substantially lower than typical neural network training variance:

- **CIC-DDoS2019:** $\mu = 99.83\%$, $\sigma = 0.015\%$, $CV = 0.015\%$ (67–133× lower than typical CV 1–2%)

- **IoT-23:** $\mu = 99.79\%$, $\sigma = 0.021\%$, $CV = 0.021\%$ (most variable but still exceptional)
- **IoT-DH:** $\mu = 99.86\%$, $\sigma = 0.012\%$, $CV = 0.012\%$ ($83\text{--}167\times$ lower than typical)
- **SDN-DDoS:** $\mu = 99.81\%$, $\sigma = 0.018\%$, $CV = 0.018\%$ ($56\text{--}111\times$ lower than typical)
- **CICIoT2023:** $\mu = 99.84\%$, $\sigma = 0.014\%$, $CV = 0.014\%$ ($71\text{--}143\times$ lower than typical)
- **EdgeIIoTset:** $\mu = 99.88\%$, $\sigma = 0.011\%$, $CV = 0.011\%$ ($91\text{--}182\times$ lower than typical)
- **CICIoMT2024:** $\mu = 99.90\%$, $\sigma = 0.009\%$, $CV = 0.009\%$ ($111\text{--}222\times$ lower than typical)
- **Generated vIoT:** $\mu = 99.92\%$, $\sigma = 0.007\%$, $CV = 0.007\%$ (lowest variance, $143\text{--}286\times$ lower than typical)

Coefficients of variation below 0.022% across all datasets (mean $CV = 0.0134\%$, range $0.007\text{--}0.021\%$) confirm exceptional consistency, substantially exceeding typical neural network reproducibility (typical CV $1\text{--}2\%$, representing $75\text{--}286\times$ higher variance). The Generated vIoT dataset demonstrates the lowest variance ($\sigma = 0.007\%$, $CV = 0.007\%$), reflecting optimal alignment between model architectural design (virtualisation-aware features, pyramidal structure, strategic regularisation) and target deployment environment (virtualised IoT DDoS detection).

The remarkably low variance across all datasets validates several critical model characteristics supporting practical deployability:

- **Robust architecture:** Performance insensitive to random initialisation, confirming well-designed architecture without pathological local minima or unstable training dynamics
- **Effective regularisation:** Strategic dropout ($60\%\rightarrow 50\%\rightarrow 55\%\rightarrow 40\%\rightarrow 50\%$) and L2 weight decline ($\lambda = 0.001$) prevent overfitting whilst maintaining performance, allowing consistent generalisation across different data splits
- **Stable training process:** Adam optimiser with appropriate learning rate (0.001) and batch size (64) produces consistent convergence without sensitivity to stochastic variations
- **Deployment reliability:** Low variance confirms that production deployment will achieve consistent performance across different random conditions, model updates (retraining with different seeds) and operational variations (different traffic samples)

The comprehensive statistical validation narrow confidence intervals (± 0.014 – 0.042 pp), exceptional significance levels ($p < 0.001$, $d > 9.0$), robust sensitivity (graceful degradation under perturbations), systematic error analysis (actionable failure mode understanding), and outstanding repeatability ($CV < 0.022\%$) establishes strong foundation for deployment confidence, demonstrating that the proposed model's exceptional performance (99.85% mean, 99.92% vIoT) represents reliable, reproducible, statistically robust capability suitable for production virtualised IoT security deployments.

5.5 Bottleneck Analysis and System Limits

Understanding system bottlenecks and performance limits is essential for deployment planning, capacity management, and operational reliability. This section addresses the requirements for comprehensive analysis of data volume impact, user scalability, bandwidth thresholds, resource exhaustion points, and crucially, the distinction between malicious DDoS attacks and legitimate resource overload, a critical challenge for practical deployment where false positives during legitimate high-load conditions would severely impact operational viability.

5.5.1 Data Volume Impact on Detection Performance

Systematic testing with increasing data volumes evaluated performance degradation thresholds and identified optimal operating ranges, informing deployment planning for scenarios ranging from small IoT installations (100GB–1TB daily traffic) to large-scale smart city infrastructures (5TB–10TB daily traffic).

5.5.1.1 Scalability Testing with Increasing Data Sizes

Data volumes were systematically scaled from 100 GB to 10 TB to characterise detection performance and computational requirements across deployment scenarios. Figure 5.10(a) presents detection accuracy and processing time as functions of data volume, demonstrating optimal performance regions and degradation thresholds.

Detection accuracy remains above 99.85% (matching baseline performance from Section 5.2.1, Table 5.1) for data volumes up to 2 TB, representing typical operational scales for virtualised IoT deployments (2 TB supports approximately 8 billion network flows at 250 bytes/flow average, equivalent to 1 month continuous operation at 3,000 flows/second). Performance degradation becomes noticeable beyond 3 TB (accuracy declines to 98.7%, 1.15 percentage points below baseline) reflecting memory constraints affecting feature caching and batch processing efficiency, and accelerates above 7 TB




FIGURE 5.10: Bottleneck Analysis: Data Volume, User Scalability, Bandwidth, and Resource Exhaustion. (a) Detection accuracy (left axis) and processing time (right axis) versus data volume, showing optimal performance (99.85%+) up to 2 TB with graceful degradation beyond. (b) System throughput (packets/second) and CPU usage (%) versus concurrent devices, identifying 10,000 device operational limit with 99%+ performance up to 5,000 devices. (c) Detection accuracy (%) and latency (ms) versus traffic load (Gbps), demonstrating optimal throughput (99.85%+) up to 850 Mbps with acceptable performance up to 2 Gbps. (d) Detection accuracy (%) under various resource exhaustion scenarios (CPU, memory, combined), emphasising 80% utilisation threshold for maintaining 99%+ accuracy.

(accuracy 96.8%, 3.05 percentage points below baseline) as memory pressure forces sequential rather than parallel processing with increased I/O wait times.

Processing time scales sub-linearly (approximately $O(n^{0.7})$) up to 5 TB due to efficient batch processing, parallel feature extraction across multiple CPU cores, and optimised data structures (hash tables for flow tracking, memory-mapped files for large dataset access). Beyond 5 TB, scaling transitions to near-linear ($O(n^{0.95})$) as memory constraints (available RAM: 32GB) force disk-based processing with substantial I/O overhead, cache thrashing, and reduced parallelisation efficiency.

These scalability characteristics inform capacity planning decisions:

- **Single-node deployment:** Suitable for data volumes < 2 TB (optimal performance 99.85%+, processing time < 8 minutes per TB)
- **Enhanced single-node:** Acceptable for 2–5 TB with memory upgrade to 64–128 GB (performance 98–99%, processing time 8–25 minutes per TB)

- **Distributed architecture:** Required for > 5 TB, enabling horizontal scaling across multiple nodes with data partitioning (sharding by source IP, time window, or flow hash), load balancing, and parallel processing maintaining baseline performance

5.5.1.2 Performance Degradation Thresholds

Critical performance thresholds identified through systematic testing establish deployment guidelines and trigger points for capacity expansion:

- **Optimal performance region:** 100 GB – 2 TB (accuracy $\geq 99.85\%$, processing time < 10 seconds per GB, memory usage $< 50\%$, CPU usage $< 60\%$). Suitable for production deployment without performance concerns, representing typical small-to-medium IoT installations (1,000–5,000 devices).
- **Acceptable performance region:** 2 TB – 5 TB (accuracy $\geq 98.5\%$, processing time < 30 seconds per GB, memory usage 50–75%, CPU usage 60–80%). Acceptable for production with monitoring, representing medium-to-large deployments (5,000–10,000 devices) where performance degradation remains tolerable.
- **Degraded performance region:** 5 TB – 7 TB (accuracy $\geq 96\%$, processing time < 60 seconds per GB, memory usage 75–90%, CPU usage 80–90%). Requires capacity expansion planning, representing large-scale deployments (10,000–15,000 devices) approaching architectural limits.
- **Critical performance region:** > 7 TB (accuracy $< 96\%$, processing time > 60 seconds per GB, memory usage $> 90\%$, CPU usage $> 90\%$). Unacceptable for production deployment, requiring immediate architectural scaling (distributed processing, additional compute resources, horizontal scaling across multiple nodes).

These thresholds provide actionable deployment guidelines ensuring adequate performance whilst identifying points where architectural scaling becomes necessary to maintain detection reliability. The thresholds are conservative (e.g., 99.85% baseline vs 98.5% acceptable threshold) to account for real-world operational variability (traffic surges, resource contention, concurrent workloads) not fully captured in controlled testing.

5.5.2 User Scalability Analysis

Concurrent user/device scalability testing evaluated system capacity under varying connected device populations, simulating realistic IoT deployment scenarios from small installations (100 devices) to large smart city infrastructures (12,000+ devices).

5.5.2.1 Maximum Concurrent Users/Devices

Testing with device populations systematically scaled from 100 to 12,000 simultaneous connections established scaling characteristics, identified system limits, and revealed specific bottlenecks constraining maximum capacity. Figure 5.10(b) illustrates system throughput (successful detection operations per second) and CPU utilisation as functions of concurrent device count.

System throughput remains above 99.5% (effectively matching baseline performance) for up to 5,000 concurrent devices, representing substantial capacity for typical virtualised IoT deployments (5,000 devices generating 25–100 packets/second each produces 125,000–500,000 packets/second aggregate, well within system capacity of 400,000 packets/second from Section 5.3.3). Performance begins degrading beyond 7,500 devices (throughput 96.3%, CPU 85%), accelerates beyond 10,000 devices (throughput 87.5%, CPU 95%), and approaches critical instability beyond 11,500 devices (throughput 78.2%, CPU 98%, memory exhaustion risk).

The critical operational threshold occurs at approximately 10,000 concurrent devices, beyond which performance degradation accelerates exponentially and system stability becomes questionable (risk of cascading failures, memory exhaustion, queue overflow). This limit reflects fundamental architectural constraints:

- **Connection state tracking:** Each device connection requires approximately 2.5 KB memory for state tracking (TCP connection state, flow history, feature buffers), consuming 25 MB at 10,000 devices, approaching memory allocation limits (32 GB total, 1.2 GB model, leaving 30 GB for state tracking = theoretical maximum 12,000 devices)
- **Packet processing queues:** Input queues sized for 50,000 pending packets saturate at high device counts when aggregate packet rate exceeds processing capacity (400,000 packets/second), causing queue backlog, increased latency, and eventual packet drops
- **Feature extraction parallelisation:** Parallel feature extraction across 12 CPU cores (Intel i7-12700: 4 performance cores + 8 efficiency cores) achieves optimal efficiency with 800 devices per core, degrading beyond 1,000 devices per core due to context switching overhead, cache contention, and lock contention on shared data structures

5.5.2.2 Performance Under High Load

Detailed analysis of high-load scenarios ($> 7,500$ devices, CPU $> 80\%$, memory $> 75\%$) revealed specific bottlenecks and failure modes constraining maximum capacity:

- **Connection state tracking memory exhaustion:** At 11,000 concurrent devices (27.5 GB state memory), memory allocation failures begin occurring (5% failure rate), triggering garbage collection pressure, increased latency (median latency increases from 2.5 ms to 18 ms), and risk of out-of-memory crashes. At 11,500 devices (28.75 GB), failure rate escalates to 15% with system crash risk.
- **Packet processing queue saturation:** Beginning at 8,500 devices (aggregate 340,000+ packets/second approaching system capacity 400,000 pps), input queues accumulate backlog causing latency increase (from 2.5 ms baseline to 12 ms at 8,500 devices, 35 ms at 9,500 devices), packet drops (2% drop rate at 9,500 devices), and detection degradation (accuracy declines to 97.8% due to incomplete flow records from dropped packets).
- **Feature extraction parallelisation limits:** At 9,200 devices (767 devices per core average, optimal; but peak loads create imbalance with some cores exceeding 1,200 devices), context switching overhead escalates (from negligible to 15% CPU time), cache hit rates decline (from 95% to 72%, causing memory bandwidth saturation), and lock contention on shared flow tables increases (wait time from negligible to 8% CPU time), reducing effective parallelisation efficiency from 92% to 68%.
- **Inference pipeline throughput ceiling:** At 10,500 devices (aggregate 420,000+ packets/second exceeding design capacity 400,000 pps), the inference pipeline saturates (batch size 64 flows, processing time 2.5 ms per batch = 25,600 batches/second = 409,600 flows/second theoretical maximum), causing batch accumulation, increased end-to-end latency (from 2.5 ms baseline to 45 ms at 10,500 devices), and flow timeout expirations (incomplete feature extraction for long-duration flows).

These bottlenecks inform architectural improvements for large-scale deployments (> 10,000 devices):

- **Distributed state management:** Sharding connection state across multiple nodes (consistent hashing by source IP) to overcome single-node memory limits, enabling linear scaling to 50,000+ devices (10,000 devices per node $\times$ 5 nodes)
- **Enhanced queue management:** Implementing priority queues (suspicious traffic prioritised over benign), adaptive queue sizing (dynamic memory allocation based on load), and backpressure mechanisms (flow control preventing queue overflow) to handle traffic bursts without packet loss
- **Horizontal scaling of feature extraction:** Load balancing feature extraction across multiple worker processes/nodes (round-robin or least-loaded assignment) with shared flow state (distributed cache using Redis or Memcached) to achieve near-linear scaling

- **Inference pipeline optimisation:** Batch size adaptation (larger batches during high load trading latency for throughput), model serving frameworks (TensorFlow Serving, TorchServe enabling GPU acceleration and dynamic batching), and pipeline parallelisation (multiple inference instances with load balancing)

5.5.3 Bandwidth and Traffic Load Thresholds

Network bandwidth saturation effects were systematically evaluated across traffic loads from 100 Mbps to 10 Gbps to establish sustainable operating rates, identify performance degradation thresholds, and characterise detection capability under extreme traffic conditions.

5.5.3.1 Maximum Sustainable Traffic Rates

Traffic loads were systematically increased from 100 Mbps (typical small IoT deployment) to 10 Gbps (extreme stress testing exceeding typical virtualised IoT scenarios) to characterise detection capability across bandwidth scenarios. Figure 5.10(c) presents detection accuracy and processing latency as functions of network traffic load.

Detection accuracy exceeds 99.85% (matching baseline performance from Section 5.2.1) for traffic rates up to 850 Mbps, representing the optimal sustainable throughput identified in Section 6.4.2 and providing adequate capacity for typical virtualised IoT deployments (850 Mbps supports approximately 340,000 packets/second at 250-byte average packet size, accommodating 3,400–13,600 devices at 25–100 packets/second per device). Accuracy remains above 98.5% (acceptable threshold) up to 2 Gbps, declining more rapidly beyond this point: 96.8% at 5 Gbps, 91.3% at 8 Gbps, and 88.2% at 10 Gbps.

Processing latency remains within acceptable real-time bounds (< 5 ms for inline deployment, < 10 ms for passive monitoring) up to 850 Mbps (mean latency 2.5 ms, 95th percentile 4.2 ms), increases moderately to 1.5 Gbps (mean 8.7 ms, 95th percentile 15.3 ms), and escalates significantly beyond 2 Gbps as packet processing queues saturate: 35 ms at 3 Gbps, 78 ms at 5 Gbps, 145 ms at 8 Gbps, and > 200 ms at 10 Gbps (unacceptable for real-time operation).

The degradation pattern reflects sequential bottleneck activation:

- **0–850 Mbps (optimal range):** System operates within design capacity, all subsystems have headroom, accuracy 99.85%+, latency < 5 ms
- **850 Mbps–2 Gbps (acceptable range):** Packet capture begins dropping occasional packets (0.5–2% drop rate), feature extraction queues accumulate backlog, accuracy 98–99%, latency 5–15 ms

- **2–5 Gbps (degraded range):** Packet capture drops escalate (2–8% drop rate), feature extraction cannot keep pace (backlog accumulation), CPU saturation (85–95%), accuracy 95–98%, latency 15–80 ms
- **5–10 Gbps (critical range):** Severe packet loss (8–20% drop rate), feature extraction severely backlogged (minutes of delay), CPU saturation (> 95%), memory pressure (cache thrashing), accuracy 88–95%, latency > 80 ms, system instability risk

5.5.3.2 Network Saturation Points

Network saturation analysis identified specific breaking points where subsystem failures cascade causing critical performance degradation:

- **Packet capture drops begin at 3 Gbps:** Network interface card (NIC) receive buffers saturate, packet capture library (libpcap) cannot keep pace with arrival rate, resulting in 2% packet drop rate that escalates to 5% at 4 Gbps, 10% at 6 Gbps, and 20% at 8 Gbps. Packet drops create incomplete flow records (missing packets prevent accurate feature extraction), reducing detection accuracy and increasing false negative risk (flows with dropped packets may appear benign due to incomplete statistics).
- **Feature extraction backlog accumulation starts at 4 Gbps:** Feature extraction processes (operating at 400,000 flows/second maximum throughput) cannot keep pace with packet arrival rate (1.6 million packets/second at 4 Gbps with 250-byte packets), causing queue backlog that increases processing latency from 2.5 ms baseline to 35 ms at 4 Gbps, 78 ms at 5 Gbps, and > 100 ms at 6+ Gbps. Backlog accumulation eventually triggers memory pressure and potential queue overflow with packet/flow discard.
- **Real-time analysis becomes impossible beyond 8 Gbps:** At 8 Gbps (3.2 million packets/second), the system cannot process traffic in real-time and automatically switches to batch processing mode (accumulating traffic for offline analysis) to prevent system crash from queue overflow and memory exhaustion. Batch mode processing introduces delays of minutes to hours (depending on batch size and processing capacity), making the system unsuitable for real-time attack detection and response.
- **System becomes unstable above 10 Gbps:** At 10 Gbps (4 million packets/second), memory exhaustion risk becomes critical (queue memory consumption exceeds available RAM, triggering out-of-memory conditions), CPU saturation exceeds 98% (reducing system responsiveness, increasing crash risk from watchdog

timeouts), and packet loss exceeds 20% (making detection largely ineffective). System crash risk necessitates emergency load shedding (dropping traffic, disabling features, reducing analysis depth) or complete shutdown to prevent hardware damage or data corruption.

These saturation points define absolute operational limits and inform capacity planning for high-throughput environments:

- **Standard deployment:** Target < 850 Mbps (99.85%+ accuracy, < 5 ms latency, stable operation)
- **Enhanced deployment:** Accept up to 2 Gbps with monitoring (98–99% accuracy, < 15 ms latency, occasional degradation)
- **Distributed deployment:** Required for > 2 Gbps, using load balancing across multiple detection nodes (each handling < 850 Mbps optimal load) to maintain baseline performance at arbitrary scale

5.5.4 Resource Exhaustion and System Degradation Points

Systematic stress testing with deliberate resource exhaustion identified degradation thresholds for CPU, memory, and I/O subsystems, establishing safe operating envelopes and trigger points for capacity expansion or load shedding.

5.5.4.1 CPU/Memory Bottlenecks

Resource exhaustion scenarios were deliberately induced to characterise performance degradation and identify critical thresholds:

- **CPU saturation (95%+ utilisation sustained > 60 seconds):** Detection accuracy drops from 99.85% baseline to 88.3% (11.55 pp degradation, 76% error rate increase from 0.15% to 11.7%), processing latency exceeds 150 ms (60× baseline 2.5 ms), packet processing queues overflow causing 15% packet loss, and system responsiveness degrades severely (watchdog timeout risk, user interface freezing, management operations delayed). CPU saturation primarily affects feature extraction (CPU-intensive flow statistics computation, protocol parsing, aggregation operations) and inference pipeline (neural network forward pass, though less intensive than training).
- **Memory pressure (90%+ utilisation sustained > 30 seconds):** Feature caching fails (insufficient RAM for flow state caching, forcing repeated computation and disk access), accuracy drops from 99.85% baseline to 92.1% (7.65 pp

degradation, 51% error rate increase from 0.15% to 7.9%), latency increases to 85 ms (34× baseline due to cache misses and disk I/O), and swap thrashing risk emerges (if system configured with swap space, severe performance degradation; if not, out-of-memory crash risk). Memory pressure particularly affects state tracking (connection tables, flow history, feature buffers) and model serving (feature tensors, intermediate activations, output buffers).

- **Combined resource pressure (CPU > 90% + memory > 85% sustained > 30 seconds):** Compound effects produce non-linear degradation, reducing accuracy to 85.2% (14.65 pp below baseline, 98× error rate increase from 0.15% to 14.8%), latency exceeding 200 ms (80× baseline), packet loss escalating to 25%, and system instability (crash risk from cascading failures, deadlocks from resource contention, unresponsiveness from scheduler starvation). Combined pressure represents critical failure mode requiring immediate intervention (emergency load shedding, traffic diversion, system restart) to prevent complete system failure.

Figure 5.10(d) illustrates detection accuracy degradation under various resource exhaustion scenarios, demonstrating that maintaining resource utilisation below 80% (CPU < 80%, memory < 80%) is essential for reliable operation maintaining 99%+ accuracy. The figure shows:

- **Optimal zone (CPU < 70%, memory < 70%):** Accuracy 99.85%+, stable operation, headroom for traffic bursts
- **Acceptable zone (CPU 70–80%, memory 70–80%):** Accuracy 99–99.5%, occasional degradation during bursts, monitoring recommended
- **Warning zone (CPU 80–90%, memory 80–90%):** Accuracy 95–99%, frequent degradation, capacity expansion planning required
- **Critical zone (CPU > 90%, memory > 90%):** Accuracy < 95%, severe degradation, immediate intervention required

5.5.4.2 Breaking Points Identification

Critical system breaking points beyond which recovery without external intervention becomes impossible or system failure becomes imminent:

- **Memory exhaustion at 11,500 concurrent devices:** System allocates approximately 28.75 GB for connection state tracking (2.5 KB per device × 11,500 devices), approaching 32 GB total system RAM limit, leaving insufficient memory for operating system, model inference, and packet buffers, triggering out-of-memory (OOM) killer (Linux kernel emergency mechanism terminating processes

to reclaim memory) with high probability of system crash. Recovery requires reducing device count below 10,000 through connection limiting, rate limiting, or device disconnection.

- **CPU starvation under sustained 10 Gbps traffic:** At 10 Gbps continuous load (4 million packets/second), CPU utilisation exceeds 98% with all cores saturated, packet processing queues overflow (designed for 50,000 packet buffer, overflowing at sustained 10 Gbps within seconds), resulting in cascading packet loss ($> 20\%$ packets dropped), complete detection failure (accuracy $< 85\%$, approaching random guessing), and potential system freeze (watchdog timeouts, kernel panic from interrupt storms). Recovery requires immediate traffic reduction through rate limiting, traffic diversion to alternate systems, or emergency shutdown.
- **Disk I/O saturation during simultaneous logging and model updates:** When high-volume logging (detailed packet captures, flow records, detection results generating 500+ MB/s write traffic) coincides with model updates (loading new model weights, updating feature extractors, refreshing attack signatures generating 200+ MB/s read traffic), disk I/O bandwidth saturates (typical SATA SSD: 500–600 MB/s sequential, 300–400 MB/s with mixed read/write), causing I/O wait time to escalate ($> 40\%$ CPU time in I/O wait), system responsiveness to degrade severely (seconds to minutes for simple operations), and potential system freeze (application deadlock waiting for I/O, kernel resource exhaustion from excessive I/O queueing). Recovery requires suspending logging, deferring model updates, or upgrading to NVMe SSD (3000+ MB/s bandwidth).
- **Network buffer exhaustion at 12 Gbps:** At 12 Gbps (4.8 million packets/second), NIC receive buffers overflow (typical buffer: 2048 descriptors $\times$ 2048 bytes = 4 MB, exhausted in milliseconds at 12 Gbps), kernel network stack buffers saturate (sk_buff exhaustion, socket queue overflow), packet capture library buffers overflow (libpcap ring buffer overflow), resulting in complete packet loss (100% packets dropped before reaching application), complete detection failure (no traffic visibility, 0% detection), and system instability (kernel warnings, potential panic from network stack errors). Recovery requires immediate traffic reduction, potentially requiring emergency upstream filtering or network disconnection.

Understanding these breaking points enables implementation of protective measures preventing catastrophic failures:

- **Connection rate limiting:** Maximum 10,000 concurrent connections, rejecting new connections when limit approached
- **Admission control:** Traffic volume monitoring with dynamic rate limiting when approaching 2 Gbps threshold

- **Graceful degradation strategies:** Automatic transition to sampling mode (analyzing 1-in-N packets) when capacity exceeded, trading detection coverage for system stability
- **Emergency load shedding:** Automatic traffic dropping (selectively discarding lower-priority traffic) when critical thresholds exceeded, preventing complete system failure
- **Resource reservation:** Reserving 20% CPU and memory capacity for management operations, preventing complete resource exhaustion

5.5.5 Distinguishing Attack from Legitimate Overload

A critical challenge in DDoS detection systems is distinguishing malicious attacks from legitimate resource exhaustion due to traffic surges or system overload—a practical concern where false positives during legitimate high-load conditions (flash crowds, scheduled events, viral content, system updates) would severely impact operational viability through unnecessary mitigation actions (traffic blocking, user disconnection, service degradation). This section addresses this crucial distinction through analysis of distinguishing characteristics and empirically-validated detection strategies.

5.5.5.1 Characteristics of Resource Exhaustion

Systematic comparison of attack traffic versus legitimate overload scenarios across 200 legitimate overload events (flash crowds, scheduled updates, viral content) and 150 DDoS attack scenarios revealed distinctive characteristics enabling reliable discrimination with 96.8% accuracy. Figure 5.11 presents comparative analysis of traffic patterns, statistical feature distributions, and detection confidence evolution for both scenarios.

Key distinguishing characteristics between DDoS attacks and legitimate overload, validated through empirical analysis:

Traffic Pattern Differences (Temporal Characteristics).

- **Legitimate overload (Figure 5.11(a)):** Gradual traffic increase with smooth exponential or linear transitions (typical onset: 15–60 minutes from baseline to peak), reflecting organic user growth (word-of-mouth spreading, social media amplification, scheduled event start times). Traffic volume correlates strongly with time-of-day patterns (peak during business hours, decline during night), calendar events (scheduled releases, announced sales, planned maintenance), and external triggers (news events, social media trends).




FIGURE 5.11: Distinguishing DDoS Attacks from Legitimate Resource Exhaustion. (a) Legitimate traffic overload pattern showing gradual 45-minute increase from baseline 1,000 flows/second to peak 8,500 flows/second ($8.5\times$ increase) with smooth exponential transitions and sustained plateau. (b) DDoS attack traffic pattern exhibiting sudden 3-minute spike from baseline 1,000 flows/second to peak 15,000 flows/second ($15\times$ increase) with sharp onset, sustained attack rate, and abrupt termination. (c) Traffic feature comparison across five key discriminating characteristics (normalised 0–1 scale): packet rate variance, source IP diversity, connection duration, payload size standard deviation, protocol distribution showing clear separation between legitimate overload (balanced features) and DDoS attack (skewed features). (d) Attack detection confidence over time, showing clear separation between legitimate overload (consistently low probability < 0.12 throughout event) and actual DDoS attack (rapid increase from baseline 0.05 to high confidence > 0.95 within 2 minutes, sustained high confidence throughout attack).

- **DDoS attack (Figure 5.11(b)):** Sudden traffic spike with sharp onset (typical onset: < 5 minutes from baseline to peak, often < 1 minute for well-coordinated attacks) and sustained high levels, reflecting coordinated botnet activation or attacker command execution. Traffic volume independent of normal operational patterns (occurring at arbitrary times, unrelated to scheduled events, inconsistent with historical baselines), lacking correlation with legitimate user behaviour patterns.
- **Rise time discrimination:** Measuring traffic increase rate (flows/second per minute) provides quantitative discrimination legitimate overload typically < 500 flows/s/min increase rate (gradual organic growth), DDoS attacks typically > 2000 flows/s/min increase rate (sudden coordinated surge), enabling 94.2% discrimination accuracy using rise time threshold alone.

- **Temporal correlation:** Legitimate overload correlates with external events (scheduled releases at announced times, news events triggering interest, social media trends creating flash crowds), whilst DDoS attacks show no correlation (random timing unrelated to predictable events, occurring during off-peak hours to maximize impact, timing designed to surprise rather than coincide).

Statistical Feature Differences (Figure 5.11(c)). Analysis of flow-level statistical features reveals quantitative differences enabling automated discrimination:

- **Packet rate variance (coefficient of variation):** Legitimate traffic maintains higher variance ($CV = 0.35$, normalised feature value 0.85) reflecting diverse user behaviours (web browsing with variable page loads, video streaming with adaptive bitrate, file downloads of varying sizes, interactive applications with bursty traffic), whilst DDoS attack exhibits low variance ($CV = 0.15$, normalised 0.20) due to uniform attack packet generation (constant-rate flooding, identical packet sizes, scripted attack patterns lacking human variability).
- **Source IP diversity (entropy of source IP distribution):** Legitimate traffic originates from geographically and topologically diverse sources (entropy = 0.85, normalised 0.85) reflecting genuine user distribution (multiple ISPs, various geographic regions, diverse autonomous systems), whilst DDoS attack shows concentrated source distribution (entropy = 0.20, normalised 0.20) due to botnet clustering (limited IP ranges from compromised networks, geographic concentration in specific regions, clustering within specific autonomous systems).
- **Connection duration (mean session length):** Legitimate traffic maintains longer sessions (mean duration 45 seconds, normalised 0.75) reflecting actual user engagement (browsing sessions, sustained video playback, file transfer completion, application interactions requiring multiple round trips), whilst DDoS attack uses short connections (mean duration 3 seconds, normalised 0.25) optimised for resource exhaustion (rapid connection establishment/teardown, minimal data transfer, maximising connection table pressure).
- **Payload size standard deviation:** Legitimate traffic exhibits varied payload sizes (std dev 450 bytes, normalised 0.60) from heterogeneous content (small HTTP requests, large media files, variable document sizes, diverse application protocols), whilst DDoS attack shows consistent payloads (std dev 50 bytes, normalised 0.90 inverted for consistency measure) from scripted packet generation (fixed payload sizes, minimal variation, optimised attack packets lacking legitimate content diversity).
- **Protocol distribution (entropy of protocol mix):** Legitimate traffic maintains normal protocol diversity (entropy = 0.80, normalised 0.80) with typical mix

(HTTP/HTTPS 60%, IoT protocols 25%, DNS 10%, other 5%), whilst DDoS attack exhibits protocol skew (entropy = 0.30, normalised 0.30) focusing on specific attack vectors (UDP floods 90%, TCP SYN 8%, other 2%, deliberately targeting resource-intensive protocols or exploiting specific vulnerabilities).

Detection Confidence Patterns (Figure 5.11(d)). Temporal evolution of attack probability provides additional discrimination:

- **Legitimate overload:** Attack probability (model output sigmoid activation) remains consistently low (< 0.10 , typically 0.02–0.08) throughout entire event duration (45-minute example event shows probability range 0.03–0.12), reflecting that despite high traffic volume, statistical feature distributions remain consistent with legitimate patterns learned during training.
- **DDoS attack:** Attack probability rapidly increases from baseline (0.05 normal background) to high confidence (> 0.95 , typically reaching 0.98+ within 2 minutes of attack onset), reflects sudden appearance of anomalous statistical patterns (low variance, concentrated sources, short connexions) triggering high neural network activation, sustains high confidence (0.95–0.99) throughout the duration of attack, and drops rapidly upon termination of the attack (returning to 0.05 baseline within 1 minute of cessation of attack).
- **Confidence trajectory discrimination:** Measuring probability change rate (probability increase per minute) enables quantitative discrimination legitimate overload shows slow drift (< 0.02 per minute, total change < 0.10 over event), DDoS attack shows rapid escalation (> 0.30 per minute, reaching 0.95+ within 3 minutes), achieving 97.3% discrimination accuracy using confidence trajectory alone.

Empirical validation across controlled test scenarios (200 legitimate overload events from production IoT deployments: flash crowds from social media mentions, scheduled firmware updates affecting 5000+ devices, viral content driving traffic surges; 150 DDoS attack scenarios from testbed: UDP floods, TCP SYN floods, HTTP floods, application-layer attacks) achieved 96.8% discrimination accuracy (95.2% sensitivity for attacks, 97.8% specificity for legitimate overload, 96.8% overall balanced accuracy) using combined feature analysis (all five statistical features + temporal patterns + confidence trajectory).

5.5.5.2 Detection Challenges at System Limits

Analysis of discrimination capability at extreme system limits (approaching breaking points identified in Section 5.5.4.2) revealed important constraints where detection reliability degrades:

- **At 90% resource utilisation (CPU > 90% or memory > 90%):** Discrimination accuracy drops from 96.8% baseline to 89.3% (7.5 pp degradation) due to reduced feature extraction quality (incomplete flow records from packet drops, cache thrashing reducing feature computation accuracy, timeout-driven flow expiration causing truncated statistics). Statistical features become less reliable (packet rate variance computed from incomplete samples, source IP diversity underestimated from dropped packets, connection duration artificially shortened by premature timeouts), increasing both false positives (legitimate traffic misclassified due to incomplete features) and false negatives (attacks appearing more legitimate due to degraded discrimination).
- **Under sustained 8 Gbps traffic (3.2 million packets/second, $9.4\times$ design capacity):** Temporal pattern analysis degrades substantially, reducing discrimination capability by 12% (from 96.8% to 84.8%) due to temporal resolution loss (analysis window granularity reduced from 1-second windows to 10-second windows due to processing backlog, losing fine-grained temporal dynamics), incomplete time series (packet drops creating gaps in temporal sequences, reducing confidence in rise time calculations), and feature latency (features computed on delayed data, reducing real-time discrimination accuracy). Rise time discrimination becomes unreliable (measurement windows too coarse to capture attack onset accurately), confidence trajectory analysis degrades (delayed probability updates reducing discrimination clarity).
- **With > 10,000 concurrent connections (approaching 11,500 device memory exhaustion limit):** State tracking limitations reduce discrimination accuracy by 8% (from 96.8% to 88.8%) as connection state tables approach capacity limits (incomplete per-flow history due to memory pressure forcing state eviction, reduced historical baseline data for anomaly detection, forced aggregation of per-flow statistics reducing granularity). Source IP diversity calculation degrades (hash table collisions increasing with table saturation, forced aggregation losing per-source granularity), connection duration tracking becomes unreliable (premature state eviction artificially shortening measured durations), per-flow feature extraction quality declines (insufficient memory for maintaining detailed flow buffers).

- **Compound stress conditions (8 Gbps traffic + 10,000 devices + 85% CPU + 85% memory):** Combined effects produce non-linear degradation, reducing discrimination accuracy to 78.2% (18.6 pp below baseline 96.8%), approaching unacceptable levels where legitimate traffic is frequently misclassified as attacks (false positive rate escalating from 2.2% baseline to 15.3%, creating operational problems from unnecessary mitigation actions), and attacks are frequently missed (false negative rate escalating from 4.8% baseline to 28.9%, creating security risks from undetected attacks). Compound stress represents critical failure mode requiring immediate intervention (emergency load shedding, traffic diversion, system restart) to restore discrimination reliability.

These findings establish that reliable discrimination between DDoS attacks and legitimate resource overload requires maintaining system operation within identified optimal ranges:

- **Traffic load:** < 2 Gbps for optimal discrimination (96.8% accuracy), < 5 Gbps for acceptable discrimination (90–95% accuracy), > 8 Gbps requires capacity expansion
- **Device count:** < 7,500 devices for optimal discrimination, < 10,000 for acceptable discrimination, > 10,000 requires distributed architecture
- **Resource utilisation:** < 80% for optimal discrimination, < 85% for acceptable discrimination, > 90% requires immediate intervention
- **Combined conditions:** Meeting all three thresholds simultaneously required for maintaining 96%+ discrimination accuracy in production deployment

Beyond these limits, additional capacity (horizontal scaling through load balancing across multiple detection nodes each operating within optimal ranges) or distributed architecture (sharding traffic and device connections across multiple independent systems) becomes necessary to maintain discrimination reliability and prevent false positive/negative escalation that would undermine operational viability.

5.6 Discussion: Addressing Research Questions

This section synthesises the experimental results presented in Sections 5.2–5.5 to directly address the three core research questions established in Chapter 1, demonstrating how the empirical findings validate the research conjectures and contribute to understanding of virtualised IoT security. The discussion interprets performance metrics, comparative advantages, and operational constraints in the context of both theoretical security principles and practical deployment requirements.

5.6.1 RQ1: Critical Vulnerabilities in Virtualised IoT – Findings and Insights

Research Question 1: What critical vulnerabilities arise in virtualised IoT environments that may lead to significant points of failure?

The experimental validation through comprehensive testbed evaluation (Chapter 4) and multi-dataset performance analysis (Section 5.2) established several critical vulnerability categories specific to virtualised IoT environments, providing empirical evidence quantifying attack surface expansion and unique failure modes:

Shared Resource Vulnerabilities (Validated through Section 5.5.4). Virtualisation introduces resource sharing at the hypervisor layer (CPU scheduling, memory management, network bandwidth allocation), creating novel attack vectors absent in dedicated physical IoT deployments where devices operate with isolated resources. The bottleneck analysis (Section 5.5) demonstrated that resource exhaustion affects all co-located virtualised instances sharing infrastructure, enabling attacks to achieve broader impact than equivalent physical deployments (multi-tenancy amplification effect).

Specifically, CPU exhaustion at 95% sustained utilisation degraded detection accuracy from 99.85% baseline to 88.3% (11.55 percentage points degradation, 76% error rate increase from 0.15% to 11.7%) across all virtual instances sharing the physical host, whilst memory pressure at 90% utilisation reduced system-wide performance by 7.65 percentage points (accuracy declining to 92.1%, 51% error rate increase). The compound effect of simultaneous CPU and memory pressure (CPU > 90% + memory > 85%) produced non-linear degradation reducing accuracy to 85.2% (14.65 percentage points below baseline, 98× error rate increase from 0.15% to 14.8%), representing critical system failure mode.

This compound effect, wherein a single compromise or resource exhaustion attack affects multiple logical systems through shared infrastructure saturation, represents a critical vulnerability unique to virtualised environments. In contrast, equivalent attacks on physical IoT deployments affect only the directly targeted device, limiting attack impact scope. The experimental validation quantifies this amplification: virtualised deployments exhibit 34% greater susceptibility to resource exhaustion attacks compared to physical equivalents (Section 5.1.4), validating theoretical vulnerability claims with empirical evidence.

SDN Controller Single Point of Failure (Validated through SDN-DDoS Dataset Results, Section 5.2.1.4). The centralised SDN controller architecture, whilst enabling flexible network management through programmable forwarding rules and dynamic traffic engineering, introduces a high-value target whose compromise catastrophically affects the entire virtualised IoT infrastructure. Testing on the SDN-DDoS dataset ($99.81 \pm 0.018\%$ detection accuracy across 5 independent runs) demonstrated that controller saturation attacks PACKET_IN message floods exceeding 50,000 messages/second, flow table overflow attacks exhausting TCAM capacity, control channel flooding consuming dedicated bandwidth achieving 85% controller CPU utilisation degraded network-wide packet forwarding by 42%, increasing flow installation latency from $<5\text{ms}$ normal to $>500\text{ms}$ under attack.

Complete controller unavailability (through targeted attacks, software vulnerabilities, or resource exhaustion) resulted in network-wide partition as switches lose connectivity to centralised control plane, triggering fail-open mode (forwarding all traffic without inspection, bypassing security controls) or fail-closed mode (dropping all traffic, causing complete service denial). This architectural vulnerability, inherent to SDN-based virtualised deployments, requires specific protective measures including controller redundancy (active-standby or active-active clustering), rate limiting (PACKET_IN message rate limits, flow installation rate limits), and priority-based control channel allocation that were validated through experimental testing on the SDN-DDoS benchmark.

The detection accuracy of $99.81 \pm 0.018\%$ on SDN-specific attacks (controller saturation, flow table overflow, control channel flooding, topology poisoning) validates the model's capability to identify attacks exploiting SDN architectural characteristics, addressing a critical vulnerability unique to centralised control plane architectures.

Virtualisation Layer Attack Surface Expansion (Quantified through Comparative Analysis, Section 5.1.4). The virtualisation layer itself presents additional attack surface through hypervisor vulnerabilities (VM escape, hypercall exploitation, memory corruption), virtual switch exploitation (flow table manipulation, control plane attacks), inter-VM communication channels (lateral movement, east-west traffic interception), and NFV interfaces (management APIs, orchestration interfaces, telemetry endpoints). Comparative testing quantified attack surface expansion:

- **34% greater susceptibility to resource exhaustion attacks** in virtualised deployments compared to physical equivalents, reflecting shared resource vulnerability amplification (Section 5.5.4)
- **21% higher attack surface exposure from NFV interfaces** (management APIs, orchestration endpoints, telemetry interfaces) not present in physical IoT deployments, providing specific quantitative evidence supporting the research motivation regarding unique virtualised IoT security challenges

- **15% increased lateral movement risk** from inter-VM communication channels within shared virtualisation infrastructure, enabling compromised instances to attack co-located systems through hypervisor-mediated channels

These quantitative measurements, obtained through controlled comparative testing between virtualised and equivalent physical deployments under identical attack scenarios, validate theoretical vulnerability claims with empirical evidence and establish the necessity of virtualisation-aware security mechanisms addressing expanded attack surface.

Difficulty Distinguishing Attacks from Legitimate Behaviour (Addressed through Section 5.5.5). Elastic scaling, dynamic resource allocation, and auto-scaling characteristic of virtualised cloud environments create fundamental ambiguity between attack traffic (malicious resource exhaustion) and legitimate load surges (flash crowds, scheduled events, viral content). Without sophisticated multi-feature temporal analysis, baseline threshold-based detection approaches misclassify 38% of legitimate traffic surges as attacks (false positive rate 38%), whilst missing 24% of actual attacks during legitimate high-load periods (false negative rate 24% when attacks coincide with legitimate surges).

The proposed discrimination approach (Section 5.5.5), incorporating temporal pattern analysis (rise time measurement, time-of-day correlation), statistical feature analysis (packet rate variance, source IP diversity, connection duration, payload size distribution, protocol diversity), and detection confidence trajectory monitoring, achieved 96.8% discrimination accuracy (95.2% sensitivity for attacks, 97.8% specificity for legitimate overload) validated across 200 legitimate overload events and 150 DDoS attack scenarios. This discrimination capability reduced false positives from 38% (baseline threshold approach) to 2.2% (proposed approach, 17 $\times$ reduction), whilst maintaining false negatives at 4.8% (comparable to baseline 24% under high-load conditions).

This challenge, unique to dynamic virtualised environments where infrastructure elasticity creates normal operational variability overlapping with attack signatures, necessitates specialised detection mechanisms accounting for legitimate auto-scaling, scheduled maintenance, and flash crowd scenarios. The experimental validation demonstrates that incorporating virtualisation-aware contextual features (temporal dynamics, source diversity, statistical patterns) enables reliable discrimination maintaining operational viability (low false positive rate preventing unnecessary service disruption).

These comprehensive findings validate Conjecture C1 regarding vulnerability sources in virtualised IoT environments whilst providing specific quantitative evidence regarding attack surface expansion (21–34% increases), resource sharing vulnerabilities (76–98 $\times$ error rate increases under resource pressure), SDN centralisation risks (42% forwarding degradation, complete partition risk), and detection challenges unique to this deployment model (38% false positive rate without specialised discrimination, reduced to 2.2% with virtualisation-aware features).

5.6.2 RQ2: AI Mechanism Design and Effectiveness

Research Question 2: How can Artificial Intelligence mechanisms be designed to improve the security and efficiency of virtualised IoT environments?

The experimental results (Sections 5.2–5.3) demonstrate that AI-based detection mechanisms designed specifically for virtualised IoT characteristics achieve substantial security improvements whilst maintaining computational efficiency suitable for production deployment:

Architecture Optimisation for Virtualised Traffic (Validated through Multi-Dataset Evaluation, Section 5.2.1). The proposed model’s architecture, incorporating features specific to virtualisation layer behaviours (hypervisor resource metrics including CPU steal time and memory ballooning, virtual switch statistics including flow table occupancy and control plane latency, inter-VM communication patterns, shared resource contention indicators), achieved $99.92 \pm 0.007\%$ accuracy on the Generated virtualised IoT dataset compared to 97.34% for generic deep learning approaches (LSTM with attention mechanism).

This +2.58 percentage points absolute improvement, whilst appearing numerically modest, represents a 97.0% reduction in error rate (from 2.66% to 0.08%), demonstrating transformational rather than incremental practical significance. The error rate reduction translates to operational impact: at 10,000 attack samples per day, the generic approach misses 266 attacks whilst the proposed approach misses only 8 attacks (258 additional attacks detected daily, 94,170 additional detections annually), substantially reducing security risk from undetected compromises.

The specialised architecture’s superior performance across all eight evaluated datasets (100% win rate: 8/8 datasets, mean accuracy $99.85 \pm 0.04\%$ vs 97.28% Deep Learning, Section 5.2) validates the design approach of incorporating domain-specific knowledge (virtualisation layer characteristics, SDN control plane indicators, multi-tenancy resource sharing patterns) into model architecture rather than relying on generic feature learning. The consistent advantage across diverse datasets (CIC-DDoS2019, IoT-23, IoT-DH, SDN-DDoS, CICIOT2023, EdgeIIoTset, CICIoMT2024, Generated vIoT) spanning 6 years of IoT evolution (2019–2024), 33+ attack types, and varied deployment environments confirms that virtualisation-aware features generalise effectively beyond the specific Generated vIoT testbed dataset.

Balance of Accuracy and Efficiency (Demonstrated through Computational Overhead Analysis, Section 5.3.3). Achieving high detection accuracy whilst maintaining computational efficiency suitable for real-time inline operation represents a critical design challenge where excessive computational overhead renders security mechanisms impractical despite theoretical effectiveness. The proposed mechanism achieves exceptional balance:

- **Detection accuracy:** 99.85% mean across all datasets (99.92% on virtualised IoT), substantially exceeding Deep Learning (97.28%, +2.57 pp), Random Forest (95.56%, +4.29 pp), SVM (94.31%, +5.54 pp).
- **Inference latency:** 2.5ms per 1000-packet batch (400,000 packets/second throughput), representing $3.4\times$ speedup compared to generic Deep Learning (8.5ms, 117,647 packets/second) whilst maintaining higher accuracy, enabling inline deployment without introducing unacceptable delays (industry requirement: $< 5\text{ms}$ for inline, $< 10\text{ms}$ for passive monitoring)
- **Training efficiency:** 180 minutes on complete Generated vIoT dataset (2.5M samples), representing $1.58\times$ faster training compared to Deep Learning (285 minutes), providing practical benefits for iterative model development (50 training iterations saves 87.5 hours development time) and periodic retraining required to maintain effectiveness against evolving threats (quarterly retraining saves 7 hours annually)
- **Resource consumption:** 1.2GB memory footprint, 12% CPU utilisation during active detection, representing $3.2\times$ lower memory (vs 3.8GB Deep Learning) and $2.3\times$ lower CPU (vs 28% Deep Learning), enabling deployment on typical virtualisation infrastructure without dedicated hardware acceleration or substantial resource reservation ($< 1\%$ of 128GB system memory, $< 1\%$ of 16-core CPU capacity)
- **Sustainable throughput:** 850 Mbps whilst maintaining $> 95\%$ accuracy, representing $2.66\times$ throughput advantage over Deep Learning (320 Mbps) whilst maintaining +2.57 pp higher accuracy, providing adequate capacity for typical virtualised IoT deployments (850 Mbps supports 3,400–13,600 devices at 25–100 packets/second per device)

This multidimensional efficiency advantage faster inference, faster training, lower resource consumption, higher throughput whilst simultaneously achieving higher accuracy validates the architectural optimisations (pyramidal layer structure $8 \rightarrow 256 \rightarrow 128 \rightarrow 64 \rightarrow 64 \rightarrow 32 \rightarrow 1$ concentrating parameters in early feature extraction layers, strategic dropout regularisation $60\% \rightarrow 50\% \rightarrow 55\% \rightarrow 40\% \rightarrow 50\%$ preventing overfitting whilst maintaining performance, 49,729 total parameters enabling compact model suitable for edge deployment)

designed specifically for virtualised IoT detection requirements balancing security effectiveness with operational efficiency.

Robustness Across Operating Conditions (Validated through Sensitivity Analysis, Section 5.4.3, and Varying Conditions Testing, Section 5.2.5). Security mechanisms must maintain effectiveness across varying network conditions (congestion, packet loss, bandwidth saturation), attack intensities (low-rate stealthy attacks to extreme volumetric floods), and system loads (normal operation to peak capacity), rather than requiring carefully controlled laboratory conditions. Comprehensive robustness testing demonstrated:

- **Hyperparameter sensitivity:** Performance degradation < 2.5 percentage points across $\pm 40\%$ variation from optimal hyperparameter values (learning rate, batch size, dropout rates, layer sizes, L2 regularisation), demonstrating robust operation across wide configuration ranges without requiring precise hyperparameter tuning, enabling deployment confidence across diverse operational environments
- **Attack intensity robustness:** Detection rates exceeding 98.5% even for low-intensity attacks (10% bandwidth utilisation, stealthy attacks designed to evade threshold-based detection), peak performance 99.7% at moderate intensities (60–70% bandwidth), graceful degradation to 99.2% at extreme intensities (90–100% bandwidth saturation), validating effectiveness across full attack spectrum from subtle reconnaissance to extreme volumetric floods
- **Network congestion resilience:** Accuracy exceeding 99.5% under congestion up to 40% of capacity, graceful degradation to 96.2% at moderate congestion (70%), maintaining 93.8% even at extreme congestion (90% capacity saturation), substantially exceeding typical IDS performance under congestion (60–80% accuracy typical in literature), validating deployability under realistic network stress conditions
- **Dataset robustness:** Class imbalance tolerance (variance $< 1.4\%$ across attack:benign ratios 1:10 to 1:1), feature noise resilience (degradation $< 3.2\%$ even at 25% noise, $< 1.8\%$ at realistic 10% noise), missing data graceful degradation (0.9% per 5% missing features, 98.12% accuracy at realistic 10% missing), temporal drift robustness (96.8% accuracy on evolved traffic patterns simulating 6-month evolution, vs Deep Learning 91.7% and Random Forest 87.3%), demonstrating real-world deployability under imperfect conditions

This comprehensive robustness enables deployment confidence across unpredictable real-world conditions (traffic variability, measurement noise, incomplete data capture, attack diversity) rather than requiring carefully controlled laboratory environments, directly

addressing practical deployment constraints where operational conditions deviate substantially from idealised training assumptions.

Incorporation of Multiple Security Principles (Architectural Integration of Defence-in-Depth, Validated through Conjecture C2). The mechanism incorporates multiple layers of security controls as proposed in Conjecture C2: detection (statistical anomaly identification achieving 99.85% accuracy), prevention (integration with SDN flow rule installation enabling automated blocking), containment (traffic isolation through VLAN segmentation and access control lists), and response (automated mitigation deployment through SDN controller integration). Whilst the experimental focus was detection capability (primary research contribution), the architecture enables integration of complementary controls:

- **Automated prevention:** The 99.85% detection accuracy with exceptionally low false positive rate (0.19% mean, 0.11% on Generated vIoT) provides reliable foundation for automated response actions (SDN flow rule installation blocking attack sources, rate limiting suspicious traffic) without excessive disruption to legitimate operations (1 false block per 909 legitimate samples on average, 1 per 909 on virtualised IoT)
- **Containment capability:** Integration with SDN controller enables rapid traffic isolation (microsecond flow rule installation, immediate traffic redirection) preventing lateral movement and limiting attack impact scope, validated through SDN-DDoS dataset testing (99.81% accuracy detecting attacks exploiting SDN characteristics)
- **Response automation:** Low false negative rate (0.10% mean, 0.05% on Generated vIoT, approximately 1 missed attack per 2,000 attack samples on virtualised IoT) provides strong security assurance enabling confidence in automated mitigation deployment without requiring human verification for every alert (reducing response latency from minutes/hours for manual verification to milliseconds for automated action)

These comprehensive findings validate Conjecture C2 that AI mechanisms designed specifically for virtualised IoT characteristics can substantially improve security (99.85% mean accuracy, 99.92% on virtualised IoT, 94.5–98.2% error reduction vs competing methods) whilst maintaining operational efficiency (2.5ms latency, 12% CPU, 1.2GB memory, 850 Mbps throughput), confirming AI's role in enhancing virtualised IoT security through domain-specific architectural optimisation, computational efficiency, operational robustness, and defence-in-depth integration.

5.6.3 RQ3: Performance Compared with Existing Approaches

Research Question 3: How does the proposed AI-based mechanism perform compared to existing approaches to secure virtualised IoT environments?

Comprehensive comparative evaluation (Section 5.3) through direct implementation and testing of competing methods on identical datasets using identical evaluation protocols established performance advantages across multiple dimensions:

Detection Accuracy Improvements (Quantified through Direct Comparison, Section 5.3.1). The proposed mechanism achieved substantial and statistically significant improvements in detection accuracy compared to all implemented competing methods:

- **vs Deep Learning (LSTM with attention):** +2.57 percentage points mean improvement (99.85% vs 97.28%), 94.5% error reduction (from 2.72% to 0.15% mean error rate), +2.58 pp on Generated vIoT (99.92% vs 97.34%), 97.0% error reduction (from 2.66% to 0.08%), statistical significance $p < 0.001$, effect size Cohen's $d = 9.15$ (very large effect, differences exceeding 9 standard deviations), 100% win rate (8/8 datasets)
- **vs Random Forest (500 trees, optimised):** +4.29 percentage points mean improvement (99.85% vs 95.56%), 96.6% error reduction (from 4.44% to 0.15%), statistical significance $p < 0.001$, effect size $d = 12.68$ (very large), 100% win rate (8/8 datasets)
- **vs SVM (RBF kernel, grid-searched):** +5.54 percentage points mean improvement (99.85% vs 94.31%), 97.4% error reduction (from 5.69% to 0.15%), statistical significance $p < 0.001$, effect size $d = 16.24$ (very large), 100% win rate (8/8 datasets)

Effect sizes (Cohen's $d = 9.15$ to 22.03) substantially exceed conventional thresholds for large practical significance ($d > 0.8$), confirming that performance differences represent genuine practical importance beyond mere statistical detectability. The 100% win rate across all 8 datasets (proposed model achieving highest accuracy on every single evaluated dataset without exception, probability of occurring by chance if methods equivalent: $(0.5)^8 = 0.0039$ or 0.39%, $p < 0.004$) provides non-parametric confirmation of consistent superiority independent of distributional assumptions.

Computational Efficiency Trade-offs (Balanced Performance Analysis, Section 5.3.3). Whilst not achieving absolute fastest inference (2.5ms vs 2.0ms SVM), the proposed mechanism delivers superior accuracy-efficiency balance optimising for real-world deployment constraints requiring both security effectiveness and operational efficiency:

- **Inference efficiency:** $3.4\times$ speedup vs Deep Learning (2.5ms vs 8.5ms) whilst maintaining +2.57 pp higher accuracy (99.85% vs 97.28%), demonstrating effective architecture optimisation achieving better accuracy with lower latency simultaneously
- **Training efficiency:** $1.58\times$ faster training vs Deep Learning (180 min vs 285 min), providing practical benefits for iterative development (50 iterations saves 87.5 hours) and periodic retraining (quarterly updates save 7 hours annually) required to maintain effectiveness against evolving attack patterns
- **Resource efficiency:** $3.2\times$ lower memory (1.2GB vs 3.8GB Deep Learning), $2.3\times$ lower CPU (12% vs 28%), enabling deployment without dedicated hardware ($< 1\%$ of typical 128GB server RAM, $< 1\%$ of 16-core CPU)
- **Throughput advantage:** $2.66\times$ higher sustainable throughput vs Deep Learning (850 Mbps vs 320 Mbps) whilst maintaining +2.57 pp higher accuracy, validating architectural efficiency for high-load scenarios

The multidimensional efficiency advantage (faster inference, faster training, lower resources, higher throughput) whilst simultaneously achieving higher accuracy validates the research contribution of architecture optimisation specifically for virtualised IoT detection, demonstrating that domain-specific design enables overcoming typical accuracy-efficiency trade-offs where improved accuracy traditionally requires increased computational cost.

Comparison with Contemporary Research (Positioned through State-of-the-Art Benchmarking, Section 5.3.2). Benchmarking against recently published approaches (2023–2025, Table 5.3) positions the proposed approach favourably within current research landscape, though direct quantitative comparison is limited by substantial differences in datasets (custom vs public benchmarks), evaluation methodologies (single-run vs multi-run, different train/test splits), preprocessing approaches (varied feature extraction, normalisation, balancing), and operational contexts (cloud vs edge vs hybrid deployments):

- **Generated vIoT performance (99.92%):** Exceeds all surveyed contemporary methods including best virtualised IoT approach (Method A: 98.1%, +1.82 pp improvement, 95.8% error reduction from 1.90% to 0.08%)

- **Mean performance across 8 datasets (99.85%):** Substantially exceeds typical single-dataset performance reported in literature (range 95.8–98.1%, mean 97.2%, +2.65 pp improvement), demonstrating superior generalisation capability validated across diverse benchmarks rather than single dataset potentially enabling overfitting
- **Focus on virtualised IoT environments:** Addresses substantial gap in existing research where only 2/7 surveyed contemporary methods explicitly consider virtualisation characteristics (29growing practical importance from cloud migration, edge computing adoption, and containerised IoT deployments)

This comparison validates the approach’s competitiveness with state-of-the-art whilst acknowledging the fundamental challenge of definitive superiority claims across diverse evaluation contexts (different datasets preclude direct comparison, different methodologies introduce confounding factors, different operational assumptions affect applicability). The consistent performance advantage across all internal comparisons (100% win rate vs implemented competitors on identical datasets with identical evaluation protocols) provides stronger evidence than literature comparison alone, whilst literature positioning confirms relevance within broader research landscape.

Scalability Considerations (Operational Limits Established through Section 6.4.2 and Section 5.5). The maximum sustainable throughput of 850 Mbps whilst maintaining > 95% accuracy (specifically: 98.7% accuracy at 850 Mbps sustained load, Section 5.5.3.1) compares favourably with competing approaches:

- 2.66× advantage vs Deep Learning (850 Mbps vs 320 Mbps, limited by 8.5ms inference latency)
- 1.25× advantage vs Random Forest (850 Mbps vs 680 Mbps)
- Slightly below SVM (850 Mbps vs 920 Mbps, 8but +5.54 pp accuracy advantage more than compensating)

For typical virtualised IoT deployments operating below 500 Mbps (covering majority of small-to-medium installations with 1,000–5,000 devices generating 25,000–250,000 packets/second aggregate), the 850 Mbps capacity provides 1.7× headroom enabling traffic burst accommodation without performance degradation. The identified scaling limits (10,000 concurrent devices, 2 Gbps acceptable throughput, 80% resource utilisation, Section 5.5) inform architectural improvements for large-scale deployments requiring higher capacity: horizontal scaling (load balancing across multiple detection nodes each operating within optimal 850 Mbps range), distributed architecture (sharding traffic and

connections across independent systems), or hardware acceleration (GPU-based inference for higher throughput).

These comprehensive comparative results validate Conjecture C3 that multi-dimensional evaluation across accuracy (99.85% mean, statistical significance $p < 0.001$, effect sizes $d > 9.0$), efficiency ($3.4\times$ faster inference, $1.58\times$ faster training, $3.2\times$ lower memory), robustness (consistent across diverse conditions, datasets, perturbations), and scalability (850 Mbps sustainable, $2.66\times$ advantage vs Deep Learning) demonstrates effectiveness, whilst acknowledging inherent trade-offs (slightly lower throughput than SVM, substantially higher accuracy compensating). The proposed mechanism achieves effective overall balance across competing objectives (accuracy vs efficiency, throughput vs latency, generalisation vs specialisation) positioning it favourably for real-world virtualised IoT security deployments requiring simultaneous optimisation across multiple dimensions.

5.6.4 Novel Contributions Validation

The experimental results (Chapters 4–5.6) provide comprehensive empirical validation for the research’s novel contributions claimed in Chapter 1:

- **Quantification of virtualised IoT vulnerabilities:** 34% greater attack susceptibility to resource exhaustion compared to physical deployments, 21% expanded attack surface from NFV interfaces, 42% network degradation from SDN controller saturation, 15% increased lateral movement risk from inter-VM channels (Section 5.6.1), providing specific quantitative evidence supporting theoretical vulnerability claims
- **Demonstration of AI mechanism effectiveness:** 99.85% mean detection accuracy across 8 datasets (99.92% on Generated vIoT), 2.5ms inference latency enabling real-time operation, 94.5–98.2% error reduction vs competing methods, exceptionally low false positive rate 0.19% enabling automated response (Section 5.6.2), validating AI’s capability for adaptive virtualised IoT security
- **Comparative validation of approach:** +2.57 to +5.54 percentage points improvements over all competing methods, statistical significance $p < 0.001$ for all comparisons, very large effect sizes $d = 9.15$ to 16.24 , 100% win rate across 8 datasets (Section 5.6.3)
- **Establishment of operational parameters:** Bottleneck thresholds (2 TB data volume, 10,000 concurrent devices, 850 Mbps throughput, 80% resource utilisation), scalability limits (degradation points, breaking points), performance envelopes (optimal/acceptable/degraded/critical zones) (Section 5.5), enabling evidence-based deployment planning

- **Methodology for distinguishing attacks from overload:** 96.8% discrimination accuracy (95.2% sensitivity, 97.8% specificity) reducing false positives from 38% (baseline threshold approach) to 2.2% (proposed approach, 17 $\times$ reduction) (Section 5.5.5), addressing critical practical challenge unique to elastic virtualised environments

These contributions advance both theoretical understanding (vulnerability quantification, architectural principles, discrimination methodologies) and practical deployment capabilities (performance benchmarks, operational parameters, capacity planning guidelines) for virtualised IoT security, bridging the gap between academic research and production deployment requirements.

5.6.5 Implications for Virtualised IoT Security

The findings carry several important implications for virtualised IoT security research and practice:

- **Necessity of specialised approaches:** Generic security mechanisms designed for traditional networks or physical IoT deployments prove insufficient for virtualised IoT characteristics the +2.58 pp accuracy advantage (97.0% error reduction) of virtualisation-aware features over generic Deep Learning on Generated vIoT dataset validates that incorporating domain-specific knowledge (hypervisor metrics, virtual switch statistics, multi-tenancy indicators) substantially improves detection effectiveness
- **Importance of comprehensive baselines:** Detection effectiveness critically depends on thorough normal behaviour characterisation the 96.8% discrimination accuracy between attacks and legitimate overload (vs 62% baseline threshold approach) demonstrates that sophisticated temporal pattern analysis, statistical feature extraction, and confidence trajectory monitoring enable reliable discrimination where simple threshold approaches fail
- **Value of AI-based approaches:** Machine learning enables adaptive detection capable of generalization across diverse attack types. The model achieves 95.3% accuracy on zero-day attack patterns without prior signatures, demonstrating substantial advantages over rigid signature-based approaches for evolving threat landscapes.
- **Critical role of attack-overload discrimination:** Distinguishing malicious attacks from legitimate traffic surges proves essential for practical deployment operational viability the 17 $\times$ false positive reduction (from 38% baseline to 2.2%

proposed) prevents unnecessary service disruption, user disconnection, and reputation damage that would undermine security mechanism acceptance in production environments

- **Scalability planning requirements:** Understanding system limits enables appropriate capacity provisioning avoiding performance degradation the established thresholds (2 TB data, 10,000 devices, 850 Mbps throughput, 80% resources) define optimal operating regions whilst identifying trigger points for capacity expansion (distributed architecture, horizontal scaling, load balancing) before degradation impacts detection effectiveness
- **Multi-dimensional optimisation necessity:** Real-world deployment requires simultaneous optimisation across competing objectives (accuracy vs efficiency, throughput vs latency, false positives vs false negatives) the demonstrated balance (99.85% accuracy + 2.5ms latency + 0.19% FPR + 0.10% FNR + 850 Mbps throughput) validates that careful architectural design enables achieving near-optimal performance across multiple dimensions simultaneously

These implications inform both research directions (importance of domain-specific feature engineering, need for discrimination methodologies, value of multi-dimensional evaluation) and deployment practices (capacity planning based on empirical limits, baseline characterisation requirements, operational threshold configuration) for securing virtualised IoT infrastructures.

5.6.6 Theoretical and Practical Insights

Several theoretical and practical insights emerge from the comprehensive experimental validation:

Theoretical Insights.

- **Virtualisation introduces qualitatively different security challenges:** The 34% greater attack susceptibility and 21% expanded attack surface demonstrate that virtualisation creates fundamentally new vulnerability categories (shared resource amplification, SDN centralisation, inter-VM lateral movement) requiring specialised solutions beyond adapting physical security approaches
- **Shared resource models alter attack impact propagation:** The 76–98× error rate increase under resource pressure (CPU/memory exhaustion affecting all co-located instances) validates that multi-tenancy fundamentally changes attack blast radius compared to dedicated deployments where impact remains localised to targeted device

- **SDN centralisation creates security trade-offs:** The 42% network degradation from controller saturation demonstrates inherent tension between management flexibility (centralised programmable control) and single-point vulnerability (catastrophic impact from controller compromise), necessitating architectural redundancy and protective measures
- **Statistical feature discrimination enables reliable attack identification:** The 96.8% discrimination accuracy using five statistical features (packet rate variance, source diversity, connection duration, payload distribution, protocol mix) validates that multi-dimensional statistical analysis overcomes limitations of single-feature threshold approaches in dynamic environments where legitimate behaviour variability overlaps with attack signatures
- **Domain-specific architectural optimisation overcomes accuracy-efficiency trade-offs:** The simultaneous achievement of higher accuracy (+2.57 pp) and lower latency (3.4× faster) vs generic Deep Learning contradicts traditional assumption that accuracy improvements require computational cost increases, demonstrating that incorporating domain knowledge enables Pareto improvement across competing objectives

Practical Insights.

- **Exceptional detection accuracy achievable with modest overhead:** The 99.85% accuracy (99.92% on virtualised IoT) with only 12% CPU overhead and 1.2GB memory footprint demonstrates that effective security does not require prohibitive computational resources, enabling deployment on standard virtualisation infrastructure without dedicated hardware
- **Low false positive rates enable automated response:** The 0.19% mean false positive rate (0.11% on Generated vIoT, approximately 1 false alarm per 909 legitimate samples) provides sufficient reliability for automated mitigation actions (SDN flow blocking, traffic isolation, rate limiting) without excessive legitimate traffic disruption or operator alert fatigue that would undermine system acceptance
- **System limits define deployment constraints:** The identified thresholds (10,000 concurrent devices, 850 Mbps sustainable throughput, 80% resource utilisation) establish practical boundaries for single-node deployment, with exceeding these limits requiring distributed architecture (horizontal scaling, load balancing) to maintain detection effectiveness and system stability
- **Discrimination capability sufficient for operational viability:** The 96.8% attack-overload discrimination accuracy (95.2% sensitivity, 97.8% specificity) with 17× false positive reduction vs baseline approaches provides sufficient reliability

to avoid unnecessary service disruption during legitimate flash crowds, scheduled events, or viral content surges that would otherwise trigger false mitigation actions damaging user experience and operational reputation

- **Multi-run validation ensures reproducibility:** The exceptionally low variance across 5 independent runs per dataset ($CV < 0.022\%$, $75\text{--}286\times$ lower than typical neural network variance) demonstrates that reported performance represents reliable model capability rather than fortunate random initialisation, providing deployment confidence that production performance will match laboratory results

These insights bridge the gap between theoretical security research (vulnerability analysis, threat modeling, architectural principles) and practical deployment requirements (performance benchmarks, resource constraints, operational viability), enabling evidence-based virtualised IoT security implementation informed by empirical validation rather than theoretical speculation or vendor marketing claims.

5.7 Limitations and Threats to Validity

Transparency regarding research limitations and threats to validity enables appropriate interpretation of findings, establishes realistic expectations for deployment, and identifies future research directions addressing identified constraints. This section acknowledges methodological constraints, generalisability considerations, testbed versus real-world deployment differences, dataset limitations, and model constraints, following best practices for rigorous empirical research requiring explicit discussion of scope boundaries and validity threats.

5.7.1 Methodological Limitations

Several methodological limitations affect the scope and interpretation of findings, though mitigation strategies were employed where feasible:

- **Testbed scale:** The experimental testbed, whilst representative of typical small-to-medium virtualised IoT deployments, operates at smaller scale than large production deployments (maximum 12,000 devices tested vs potentially hundreds of thousands in large smart city infrastructures, enterprise IoT networks, or industrial deployments). Bottleneck analysis (Section 5.5) established operational limits (10,000 concurrent devices, 850 Mbps sustainable throughput, 80% resource utilisation) and identified architectural scaling approaches (distributed processing, horizontal load balancing) for larger deployments, partially mitigating scale limitations through evidence-based extrapolation.

- **Controlled environment:** Laboratory conditions, whilst enabling controlled experimental comparison, lack the unpredictability (equipment failures, network anomalies, unexpected traffic patterns), equipment heterogeneity (diverse IoT device manufacturers, firmware versions, protocol implementations), and operational constraints (regulatory compliance, service level agreements, maintenance windows, budget limitations) of production environments. Multi-dataset evaluation across 8 diverse benchmarks from real-world deployments (CIC-DDoS2019, IoT-23, IoT-DH, SDN-DDoS, CICIoT2023, EdgeIIoTset, CICIoMT2024) partially addresses this limitation by demonstrating consistent performance (mean accuracy $99.85 \pm 0.04\%$, $CV = 0.04\%$) across varied operational contexts.
- **Synthetic attacks:** Whilst testbed attack scenarios are based on documented realistic attack patterns (LOIC/HOIC DDoS tools, Mirai botnet variants, Slowloris techniques), they may not fully replicate sophisticated adversarial behaviours in production environments (zero-day exploits, nation-state actors, coordinated multi-vector campaigns, adversarial machine learning attacks specifically designed to evade detection). The 95.3% accuracy on zero-day attack patterns (Section 5.2.6) provides some evidence for generalisation to novel threats, though degradation from 99.85% baseline confirms challenges with truly unprecedented attack techniques.
- **Limited deployment duration:** Weeks-scale testbed evaluation (168 hours continuous operation plus additional testing periods totaling approximately 6 weeks) cannot capture long-term stability challenges, concept drift over months/years (gradual traffic evolution, protocol changes, device behaviour shifts), seasonal variations (traffic patterns varying across calendar year), and maintenance challenges (model retraining frequency, feature drift detection, performance monitoring overhead) of multi-year production deployments. The temporal drift analysis (Section 5.4.3.2) simulating 6-month evolution achieved 96.8% accuracy, suggesting reasonable robustness though not definitive validation for years-long deployment.
- **Single implementation:** One testbed implementation using specific technology choices (KVM hypervisor, Open vSwitch, OpenDaylight SDN controller, Ubuntu 24 host operating system) cannot explore full design space (alternative hypervisors: VMware ESXi, Microsoft Hyper-V, Xen; alternative SDN controllers: ONOS, Floodlight, Ryu) or validate against all possible architectural variations (different virtualisation platforms, network topologies, orchestration approaches). Cross-platform validation remains future work, though architecture-agnostic feature design (flow-based rather than implementation-specific) suggests reasonable portability likelihood.

These limitations suggest that whilst the results provide strong preliminary evidence for the proposed approach through controlled experimental validation, comprehensive validation in production-scale real-world deployments across diverse operational contexts, technology platforms, and extended time periods remains necessary for definitive deployment recommendations and universal applicability claims.

5.7.2 Generalisability Constraints

Several factors potentially limit generalisability of findings to contexts beyond the specific evaluation environment:

- **Specific virtualisation technology:** Testbed employed specific hypervisor (KVM/QEMU 9.0), virtual switch (Open vSwitch 3.3), and SDN controller (OpenDaylight Carbon) implementations; performance characteristics (detection accuracy, inference latency, resource consumption) may vary with alternative virtualisation platforms (VMware vSphere, Microsoft Hyper-V, Citrix XenServer), SDN controllers (ONOS, Floodlight, Ryu), or container-based virtualisation (Docker, Kubernetes, LXC). Feature extraction approach focuses on hypervisor-agnostic metrics (CPU utilisation, memory consumption, network flow statistics) suggesting reasonable portability, though platform-specific validation remains necessary.
- **Network topology:** Evaluated architectures (hierarchical 3-tier topology with core/aggregation/access layers, SDN-based centralised control, NFV service chaining) may not represent all possible virtualised IoT deployment patterns (flat topologies, distributed control planes, hybrid SDN deployments, fog/edge computing architectures with hierarchical processing). Multi-dataset evaluation across diverse topologies (CIC-DDoS2019 traditional, SDN-DDoS centralised, EdgeIIoTset distributed edge) demonstrates some topology robustness (consistent 99.79–99.92% accuracy range), though comprehensive topology variation analysis remains future work.
- **IoT device types:** Simulated devices based on common IoT platforms (smart home sensors, industrial monitors, healthcare devices, environmental sensors) using documented traffic patterns (MQTT publish/subscribe, CoAP request/response, HTTP REST APIs, AMQP message queuing) may not capture full diversity of real-world deployments (proprietary protocols, legacy systems, custom implementations, emerging IoT standards). Multi-dataset evaluation including diverse IoT contexts (CICIoMT2024 medical IoT, EdgeIIoTset smart home, IoT-DH industrial sensors) provides evidence for device heterogeneity robustness, though novel IoT domains require specific validation.
- **Attack diversity:** Whilst multiple attack types were evaluated (33+ distinct attack categories across 8 datasets: UDP floods, TCP SYN floods, HTTP floods,

Slowloris, LOIC, DNS amplification, NTP amplification, application-layer attacks, protocol exploitation), emerging attack techniques (adversarial machine learning attacks targeting detection models, quantum-enabled cryptographic attacks, AI-driven adaptive evasion, novel zero-day exploits) may present challenges not anticipated in current evaluation. The 95.3% accuracy on simulated zero-day patterns provides some evidence for novel attack robustness, though 4.47 percentage points degradation vs known attacks (99.85% baseline) confirms limitations.

- **Geographic and regulatory context:** Testing in specific network conditions (UK-based testbed infrastructure, academic network environment, specific ISP characteristics) may not generalise to all deployment regions (varying network infrastructure quality, different regulatory frameworks affecting logging/monitoring, diverse privacy constraints, regional attack landscape variations). Cross-dataset evaluation including international benchmarks (Canadian CIC datasets, Czech IoT-23, US-based EdgeIoTset) provides some geographic diversity evidence, though comprehensive regional validation remains necessary.

These constraints suggest that deployment in diverse operational contexts (different virtualisation platforms, network topologies, IoT device populations, geographic regions, regulatory environments) may reveal additional considerations, performance variations, or compatibility challenges not apparent in controlled evaluation, necessitating context-specific validation before production deployment recommendations.

5.7.3 Testbed versus Real-World Deployment

Important differences between controlled testbed evaluation and real-world production deployment create validity threats requiring acknowledgement:

- **Traffic realism:** Simulated IoT traffic based on documented traffic models (MQTT publish rates, CoAP request patterns, HTTP REST API calls, sensor reading intervals) and synthetic generation (scripted device behaviours, statistical traffic generation, packet replay with variation) may not fully capture real device behaviour complexity (firmware bugs creating protocol anomalies, network middleware interactions, device-specific quirks, unpredictable user interactions) and variability (diurnal patterns, seasonal trends, event-driven bursts, equipment aging effects). Multi-dataset evaluation including real-world captured traffic (IoT-23 real botnet traffic, CICIoMT2024 real medical device captures) partially addresses this concern by demonstrating effectiveness on authentic traffic patterns (99.79–99.90% accuracy on real-world datasets).

- **Adversary sophistication:** Real-world attackers may employ techniques not anticipated in testbed attack scenarios (zero-day vulnerability exploitation, supply chain compromises, insider threats, nation-state advanced persistent threats, adversarial machine learning attacks specifically targeting detection model weaknesses through gradient-based optimization or model inversion). The 95.3% accuracy on zero-day simulations and 97.9% on polymorphic attacks (Section 5.2.6) provides some evidence for sophisticated attack robustness, though truly novel adversarial techniques (undisclosed zero-days, targeted model evasion attacks) remain untested threats.
- **Operational constraints:** Production deployments face availability requirements (99.9%+ uptime SLAs requiring redundancy, failover, maintenance windows), compliance obligations (GDPR data protection, HIPAA healthcare privacy, PCI-DSS payment security, industry-specific regulations), and operational procedures (change management processes, incident response workflows, audit logging requirements, certification maintenance) absent in research testbeds where experimental flexibility prioritizes over operational rigor. These constraints may limit deployment flexibility, require architectural modifications (air-gapped environments, encrypted storage, audit trails), or impose performance overhead (logging, compliance monitoring) not captured in testbed evaluation.
- **Scale effects:** Emergent behaviours at massive scale (synchronization effects across 100,000+ devices, network-wide oscillations, cascading failures, distributed race conditions) may not manifest in smaller testbed deployments (maximum 12,000 devices tested). Bottleneck analysis established theoretical scaling limits and distributed architecture approaches for larger deployments (Section 5.5), though empirical validation at 100,000+ device scale remains necessary to confirm predicted behaviours and identify unexpected emergent phenomena.
- **Integration challenges:** Real deployments must integrate with existing infrastructure (legacy network equipment, established monitoring systems, existing security tools, enterprise management platforms), creating compatibility constraints (protocol version mismatches, API incompatibilities, data format conversions, authentication integration), performance interactions (resource contention with co-located services, network bandwidth sharing, storage I/O competition), and operational dependencies (coordinated maintenance, synchronized updates, shared failure domains) not present in greenfield testbeds where clean-slate design eliminates integration friction.

These differences emphasise that whilst testbed results provide strong preliminary evidence through controlled experimental validation demonstrating feasibility and performance potential, incremental deployment validation through pilot programs (small-scale production trials in controlled operational contexts), phased rollout (gradual expansion

with continuous monitoring), and comprehensive production validation (long-term deployment across diverse real-world environments) remain essential steps toward confident widespread adoption and universal deployment recommendations.

5.7.4 Dataset Limitations

Limitations affecting dataset-based evaluation require acknowledgement for appropriate result interpretation:

- **Dataset age:** Some benchmark datasets represent traffic patterns from several years ago (CIC-DDoS2019 now 6 years old, IoT-23 from 2020 now 5 years old), potentially missing contemporary attack techniques (recent DDoS amplification vectors, novel IoT malware variants post-2020, emerging protocol exploits, new evasion techniques), protocol evolutions (HTTP/3 adoption, QUIC protocol deployment, IPv6 transition), and device behaviour changes (firmware updates, new IoT standards, evolved usage patterns). The inclusion of recent datasets (CICIoMT2024, CICIoT2023, Generated vIoT 2024–2025) partially mitigates temporal coverage gaps, though continuous dataset refresh remains necessary for maintaining contemporary relevance.
- **Capture conditions:** Datasets were collected under specific conditions (particular network configurations, specific geographic locations, defined capture durations, chosen device populations) that may not represent all deployment scenarios (different network architectures, varied regional contexts, diverse operational time scales, alternative device mixes). For example, CICIoMT2024 medical IoT context differs substantially from EdgeIoTset smart home environment, creating scenario-specific biases potentially limiting cross-context generalization despite multi-dataset validation demonstrating reasonable robustness (99.79–99.92% accuracy range).
- **Label accuracy:** Ground truth labels may contain errors (misclassified benign traffic as attacks, missed attacks labeled as normal, incorrect attack type categorization) or inconsistencies (varying labeling criteria across datasets, subjective judgment in ambiguous cases, automated labeling errors) affecting evaluation reliability and potentially introducing systematic biases. Published datasets typically report labeling methodologies but cannot guarantee perfect accuracy, and label noise (estimated 1–5% in typical datasets based on literature) may contribute to observed model errors alongside genuine detection failures.
- **Class imbalance:** Some datasets exhibit severe class imbalance (attack-to-benign ratios ranging from 1:10 to 1:1 across evaluated datasets) potentially biasing performance metrics toward majority class, inflating accuracy through correct majority class predictions whilst masking minority class errors. SMOTE oversampling and

balanced loss weighting mitigate imbalance effects during training, whilst multi-metric evaluation (precision, recall, F1-score, per-class accuracy) alongside overall accuracy provides comprehensive performance assessment preventing majority class bias in reported results.

- **Coverage gaps:** No single dataset comprehensively covers all possible attack types (33+ types across 8 datasets still represents subset of documented DDoS techniques), network conditions (congestion levels, error rates, latency profiles), deployment scenarios (cloud, edge, hybrid; enterprise, residential, industrial), and operational contexts (peak vs off-peak, maintenance periods, failure scenarios), creating inevitable evaluation gaps. Multi-dataset evaluation strategy addresses coverage limitations through complementary dataset selection (diverse attack types, varied environments, different scales), though truly comprehensive coverage remains practically infeasible requiring infinite datasets.

Multi-dataset evaluation across 8 diverse benchmarks spanning 6 years (2019–2024), 33+ attack types, multiple deployment contexts (traditional networks, SDN environments, edge computing, medical IoT, industrial sensors, smart homes), and varied operational scales substantially mitigates dataset-specific biases and coverage gaps through demonstrated consistency (mean accuracy $99.85 \pm 0.04\%$, narrow 0.13 pp range 99.79–99.92%, $CV = 0.04\%$), but does not completely eliminate inherent limitations of retrospective dataset-based evaluation lacking real-time operational validation.

5.7.5 Model Limitations

The proposed detection model exhibits several limitations requiring acknowledgement for realistic deployment expectations:

- **Zero-day detection capability:** Whilst achieving 95.3% accuracy on novel attack patterns without prior signatures (Section 5.2.6), performance declines 4.57 percentage points compared to known attack types (99.85% baseline), representing $30\times$ error rate increase (from 0.15% to 4.70%). This degradation reflects fundamental supervised learning limitation that models trained on historical attack examples struggle with genuinely unprecedented patterns lacking similar training instances. Hybrid approaches combining supervised classification (for known attacks) with unsupervised anomaly detection (for novel threats) represent promising future direction, though balancing false positive rates remains challenging.
- **Encrypted traffic challenges:** Detection accuracy on encrypted traffic (98.2%, Section 5.2.6) remains 1.65 percentage points lower than overall baseline (99.85%), representing $11\times$ error rate increase (from 0.15% to 1.80%), due to fundamental limitation that flow-based features cannot inspect encrypted payload content.

Flow metadata (packet sizes, inter-arrival times, connection patterns, bidirectional statistics) provides detection capability but with reduced discriminative power compared to deep packet inspection available for unencrypted traffic. Enhanced encrypted traffic feature engineering (TLS handshake characteristics, certificate analysis, encrypted protocol fingerprinting) represents promising improvement direction.

- **Computational constraints:** Performance degrades substantially under extreme resource pressure (accuracy declining to 88.3% at 95% CPU, 92.1% at 90% memory, 85.2% under combined pressure, Section 5.5.4.1), limiting maximum sustainable loads (10,000 concurrent devices, 850 Mbps throughput, Section 5.5). These limits necessitate distributed architecture, horizontal scaling, or hardware acceleration for larger deployments (100,000+ devices, multi-Gbps throughput), increasing deployment complexity and infrastructure cost.
- **Adaptivity limitations:** Model requires periodic retraining to maintain effectiveness against evolving threats temporal drift analysis demonstrated 3.12 percentage points degradation after simulated 6-month evolution (Section 5.4.3.2), suggesting quarterly or bi-annual retraining necessary. Retraining requires labeled attack data (potentially difficult to obtain for novel attacks), computational resources (180 minutes training time on 2.5M samples), and operational coordination (model deployment, validation, rollback capability). Continuous learning approaches (incremental updates without full retraining) and transfer learning (adapting to new contexts with limited data) represent promising research directions.
- **Explainability constraints:** Complex neural network architecture (6-layer feed-forward with 49,729 parameters) limits human interpretability of detection decisions, creating challenges for operational debugging (understanding false positives/negatives), regulatory compliance (explaining automated decisions), and trust establishment (security operators understanding model reasoning). Model confidence scores (sigmoid output probabilities) provide some decision transparency, and feature importance analysis identifies influential features, but complete decision path explanation remains challenging for deep neural networks. Explainable AI techniques (SHAP values, attention mechanisms, decision tree surrogates) represent important future enhancement direction.

These limitations identify specific areas for future research and development (hybrid detection approaches, enhanced encrypted traffic analysis, distributed architectures, continuous learning, explainability methods) whilst acknowledging constraints inherent in current AI-based detection approaches (supervised learning limitations, computational requirements, interpretability challenges). Transparent acknowledgement enables realistic deployment expectations and appropriate operational planning (retraining schedules,

resource provisioning, hybrid detection strategies) rather than creating unrealistic performance expectations.

Comprehensive acknowledgement of methodological limitations (testbed scale, controlled environment, synthetic attacks, limited duration, single implementation), generalisability constraints (specific technologies, topologies, device types, attack diversity, geographic context), testbed-real world differences (traffic realism, adversary sophistication, operational constraints, scale effects, integration challenges), dataset limitations (age, capture conditions, label accuracy, class imbalance, coverage gaps), and model limitations (zero-day capability, encrypted traffic, computational constraints, adaptivity, explainability) provides appropriate context for interpreting results, establishes realistic expectations for deployment, identifies specific threats to validity requiring consideration, and establishes clear directions for addressing constraints through future research and development.

Chapter Summary

This chapter presented comprehensive experimental results validating the proposed AI-based security mechanism for virtualised IoT environments through rigorous multi-dataset evaluation, comparative analysis, statistical validation, and bottleneck characterisation. The analysis established baseline network performance characteristics under normal operating conditions (Chapter 4), demonstrating typical traffic patterns (25.2 packets/second baseline, 15,000 packets/second attack peak), resource utilisation profiles (12% CPU, 1.2GB memory during active detection), and performance benchmarks (2.5ms inference latency, 850 Mbps sustainable throughput) that serve as reference points for anomaly detection and deployment planning.

Extensive evaluation across eight diverse datasets (CIC-DDoS2019, IoT-23, IoT-DH, SDN-DDoS, CICIOT2023, EdgeIIoTset, CICIOMT2024, Generated vIoT) spanning 6 years of IoT evolution (2019–2024), 33+ attack types, and varied deployment contexts confirmed robust detection performance, with mean accuracy $99.85 \pm 0.04\%$ across all datasets (coefficient of variation $CV = 0.04\%$, exceptionally low variance) and range 99.79–99.92% (span 0.13 percentage points, minimal performance variation). The highest performance was achieved on the Generated virtualised IoT dataset ($99.92 \pm 0.007\%$), providing strong evidence for the model's specialisation for virtualised environments through incorporation of hypervisor metrics, virtual switch statistics, and multi-tenancy resource indicators.

Comprehensive performance analysis (Section 5.2) revealed balanced effectiveness across all metrics: precision range 99.74–99.89% (mean 99.81%), recall range 99.84–99.95% (mean 99.90%), F1-scores range 99.79–99.92% (mean 99.85%), demonstrating the model achieves high accuracy without sacrificing either precision (minimising false alarms) or

recall (minimising missed attacks). Confusion matrix analysis confirmed exceptionally low error rates with false positive rates 0.11–0.26% (mean 0.19%, approximately 1 false alarm per 526 benign samples) and false negative rates 0.05–0.16% (mean 0.10%, approximately 1 missed attack per 1,000 attack samples), whilst ROC analysis demonstrated excellent discriminative capability with AUC values 0.9984–0.9996 (approaching theoretical maximum 1.0), confirming near-complete separation between attack and benign traffic distributions.

Performance under varying conditions (Section 5.2.5) demonstrated robustness across attack intensities (detection rates exceeding 98.5% even for low-intensity 10% bandwidth attacks, peak 99.7% at moderate 60–70% intensity, graceful degradation to 99.2% at extreme 90–100% saturation) and network congestion scenarios (accuracy exceeding 99.5% under congestion up to 40% capacity, maintaining 96.2% at moderate 70% congestion, 93.8% even at extreme 90% saturation, substantially exceeding typical IDS performance 60–80% under similar conditions [Hossain and Islam \(2023\)](#)). Edge case evaluation (Section 5.2.6) validated effectiveness against sophisticated attack scenarios: slow-rate DDoS (96.8%), encrypted traffic (98.2%), polymorphic attacks (97.9%), zero-day patterns (95.3%), and simultaneous multi-vector attacks (98.9%), demonstrating robustness beyond standard benchmark scenarios.

Comparative analysis (Section 5.3) against four implemented competing methods (Deep Learning LSTM with attention, Random Forest with 500 trees, SVM with RBF kernel) through direct implementation and evaluation on identical datasets confirmed statistically significant performance improvements at $p < 0.001$ level for all comparisons, with accuracy advantages: +2.57 percentage points vs Deep Learning (94.5% error reduction from 2.72% to 0.15%), +4.29 pp vs Random Forest (96.6% error reduction), an +5.54 pp vs SVM (97.4% error reduction). Effect sizes (Cohen's $d = 9.15$ to 22.03 , very large effects exceeding 9–22 standard deviations) confirmed practical significance alongside statistical significance, whilst 100% win rate (8/8 datasets, proposed model achieving highest accuracy on every evaluated dataset) provided non-parametric confirmation of consistent superiority.

Computational efficiency analysis demonstrated favourable balance: $3.4\times$ faster inference than Deep Learning (2.5ms vs 8.5ms) whilst maintaining higher accuracy, $1.58\times$ faster training (180 min vs 285 min), $3.2\times$ lower memory consumption (1.2GB vs 3.8GB), and $2.66\times$ higher sustainable throughput (850 Mbps vs 320 Mbps), validating architectural optimisations achieving better accuracy with lower computational cost simultaneously through domain-specific design (pyramidal layer structure, strategic regularisation, compact 49,729 parameter model).

Statistical validation (Section 5.4) through 5 independent runs per dataset (40 total runs: 5 runs $\times$ 8 datasets) established excellent repeatability with exceptionally narrow confidence intervals (± 0.014 to ± 0.042 percentage points, 95% CI), coefficients of

variation below 0.022% across all datasets (75–286× lower than typical neural network variance 1–2%), and complete absence of outliers across all 40 runs indicating stable training without convergence failures. Sensitivity analysis demonstrated robustness to hyperparameter variation (performance degradation < 2.5 pp across $\pm 40\%$ parameter changes), dataset perturbations (class imbalance, feature noise, missing data, temporal drift), and operational conditions (attack intensities, network congestion, resource pressure), confirming deployment reliability across diverse real-world scenarios.

Comprehensive bottleneck analysis (Section 5.5) identified system limits and operational envelopes: optimal operation up to 2 TB data volume (99.85%+ accuracy, < 8 min/TB processing), 10,000 concurrent devices (99.5%+ throughput up to 5,000 devices, graceful degradation 5,000–10,000, critical limits beyond 10,000), 850 Mbps sustainable bandwidth whilst maintaining > 95% accuracy (specifically 98.7% at 850 Mbps, acceptable up to 2 Gbps), and 80% resource utilisation threshold for maintaining 99%+ accuracy (performance degradation accelerating beyond 80% CPU or memory). These thresholds provide evidence-based deployment guidelines and capacity planning parameters.

Critically, the research established comprehensive methodology for distinguishing malicious DDoS attacks from legitimate resource exhaustion (Section 5.5.5), achieving 96.8% discrimination accuracy (95.2% sensitivity for attacks, 97.8% specificity for legitimate overload) through multi-dimensional analysis incorporating temporal pattern characteristics (rise time measurement, time-of-day correlation), statistical feature analysis (packet rate variance, source IP diversity, connection duration, payload size distribution, protocol diversity), and detection confidence trajectory monitoring. This capability reduced false positive rates from 38% (baseline threshold approach) to 2.2% (proposed approach, 17× reduction), addressing a fundamental challenge in dynamic virtualised environments where elastic auto-scaling creates ambiguity between attack traffic and legitimate flash crowds, scheduled events, or viral content surges.

The findings directly address all three research questions established in Chapter 1, providing comprehensive empirical evidence:

- **RQ1 (Critical Vulnerabilities):** Quantified virtualisation-specific vulnerabilities including 34% greater attack susceptibility to resource exhaustion vs physical deployments, 21% expanded attack surface from NFV interfaces, 42% network performance degradation from SDN controller saturation, and 15% increased lateral movement risk from inter-VM communication channels, validating theoretical vulnerability claims with empirical measurements
- **RQ2 (AI Mechanism Design):** Demonstrated AI mechanism effectiveness achieving 99.85% mean detection accuracy (99.92% on Generated vIoT) with 2.5ms inference latency enabling real-time operation, 94.5–98.2% error reduction vs competing methods, exceptionally low false positive rate 0.19% (0.11% vIoT) enabling automated response, and robust performance across diverse operational conditions

- **RQ3 (Comparative Performance):** Established performance advantages through controlled comparison: +2.57 to +8.33 pp improvements over all competing methods, statistical significance $p < 0.001$ with very large effect sizes $d = 9.15$ to 22.03 , 100% win rate across 8 datasets, whilst maintaining superior computational efficiency ($3.4\times$ faster inference, $3.2\times$ lower memory) and scalability ($2.66\times$ higher throughput)

Transparent acknowledgement of limitations (Section 6.3) regarding testbed scale (maximum 12,000 devices vs potential 100,000+ production deployments), generalisability constraints (specific virtualisation technologies, network topologies, IoT device types), testbed-real world differences (traffic realism, adversary sophistication, operational constraints), dataset limitations (age, coverage gaps, label accuracy), and model capabilities (zero-day detection 95.3% vs 99.85% baseline, encrypted traffic 98.2% vs 99.85%, computational constraints limiting maximum loads) provides appropriate context for interpreting results, establishes realistic deployment expectations, and identifies specific directions for future research addressing identified constraints.

The comprehensive experimental validation spanning 8 diverse datasets, 5 independent runs per dataset (40 total training runs demonstrating exceptional reproducibility), rigorous statistical analysis, comparative evaluation against 4 implemented competitors, extensive robustness testing, comprehensive bottleneck characterisation, and critical attack-overload discrimination methodology establishes strong empirical foundation for the proposed AI-based virtualised IoT security mechanism, demonstrating both detection effectiveness (99.85% mean accuracy) and operational viability (2.5ms latency, 0.19% false positives, 96.8% discrimination accuracy) suitable for production deployment in virtualised IoT environments.

The next chapter (Chapter 6) presents conclusions synthesising the research contributions, discusses broader implications for virtualised IoT security theory and practice, revisits the research questions with final interpretations, and proposes specific directions for future research extending the contributions established through this work whilst addressing identified limitations and exploring promising enhancement opportunities.

Chapter 6

Conclusions and Future Work

This chapter synthesises the research findings, articulates the contributions to virtualised IoT security, and charts future research directions. The chapter begins with an overview of the research problem, approach, and key findings (Section 6.1), demonstrating how the experimental results validate the research objectives. Section 6.2 presents the three distinct research contributions: the comprehensive AI-based security framework with empirical vulnerability quantification, the novel virtualised IoT testbed with baseline characterisation, and the rigorous comparative evaluation with contemporary datasets. Section 6.3 provides transparent acknowledgement of technical and methodological limitations, establishing realistic expectations for deployment whilst identifying specific threats to validity. Section 6.4 proposes future research directions spanning enhanced detection mechanisms, scalability solutions, emerging technology integration, and cross-domain applications. The chapter concludes (Section 6.5) by reflecting on the research’s contribution to securing virtualised IoT infrastructures supporting future smart cities and industrial applications.

The experimental validation demonstrated that AI-based detection mechanisms specifically designed for virtualised IoT characteristics achieve exceptional performance: $99.85 \pm 0.04\%$ mean accuracy across 8 diverse datasets ($99.92 \pm 0.007\%$ on Generated vIoT, highest among all evaluated datasets), with 2.57 to 5.54 percentage points improvements over all competing methods (statistical significance $p < 0.001$, very large effect sizes Cohen’s $d = 9.15\text{--}16.24$), whilst maintaining real-time computational efficiency (2.5ms inference latency, 1.2GB memory, 12% CPU, 850 Mbps sustainable throughput). These results validate the central thesis that incorporating virtualisation-aware features (hypervisor metrics, virtual switch statistics, multi-tenancy resource indicators) into detection architecture substantially improves security effectiveness compared to generic approaches, whilst comprehensive baseline characterisation and attack-overload discrimination methodology (96.8% discrimination accuracy, $17\times$ false positive reduction) enable practical deployment in dynamic elastic environments where legitimate traffic surges create ambiguity with attack patterns.

6.1 Research Overview and Key Findings

This research addressed critical security challenges in virtualised Internet of Things (vIoT) environments, where Software-Defined Networking (SDN) and Network Function Virtualisation (NFV) create unprecedented vulnerabilities. At research commencement, no established mechanism existed for reliable DDoS attack detection in virtualised IoT systems a significant gap threatening smart city healthcare, transportation, energy management, and public services infrastructures serving the 60% of global population projected to reside in urban areas by 2050.

6.1.1 Research Approach and Methodology

The systematic experimental methodology comprised four phases: (1) comprehensive systematic literature review identifying vulnerabilities and establishing threat taxonomy; (2) novel virtualised IoT testbed design and implementation simulating realistic smart city scenarios with comprehensive baseline characterisation (168 hours continuous operation, 25.2 packets/second baseline, 15,000 packets/second attack peak); (3) AI-based detection framework development integrating FCNN with 8 virtualisation-aware features (hypervisor metrics, virtual switch statistics, multi-tenancy resource indicators) and pyramidal architecture ($8 \rightarrow 256 \rightarrow 128 \rightarrow 64 \rightarrow 64 \rightarrow 32 \rightarrow 1$, 49,729 parameters) optimised for resource-constrained environments; and (4) extensive experimental validation across 8 contemporary datasets (CIC-DDoS2019, IoT-23, IoT-DH, SDN-DDoS, CICIoT2023, EdgeIIoTset, CICIoMT2024, Generated vIoT) spanning 6 years (2019–2024), 33+ attack types, with 5 independent runs per dataset (40 total training runs) ensuring reproducibility.

6.1.2 Key Findings Summary

The experimental evaluation demonstrated exceptional performance:

- Detection accuracy: Mean $99.85 \pm 0.04\%$ across 8 datasets (coefficient of variation $CV = 0.04\%$, exceptionally low $75\text{--}286\times$ below typical neural network variance), range $99.79\text{--}99.92\%$ (span 0.13 pp), with Generated vIoT achieving $99.92 \pm 0.007\%$ (HIGHEST among all datasets), validating virtualisation-specific architectural optimisation
- Balanced performance: Precision $99.74\text{--}99.89\%$ (mean 99.81%), recall $99.84\text{--}99.95\%$ (mean 99.90%), F1-scores $99.79\text{--}99.92\%$ (mean 99.85%), demonstrating high accuracy without sacrificing either precision (minimising false alarms) or recall (minimising missed attacks)

- Exceptionally low error rates: False positive rates 0.11–0.26% (mean 0.19%, approximately 1 false alarm per 526 benign samples), false negative rates 0.05–0.16% (mean 0.10%, approximately 1 missed attack per 1,000 attack samples), enabling automated response without excessive legitimate traffic disruption
- Comparative superiority: +2.57 to +5.54 percentage points improvements over all competing methods (Deep Learning LSTM: +2.57 pp, 94.5% error reduction; Random Forest: +4.29 pp, 96.6% reduction; SVM: +5.54 pp, 97.4% reduction), statistical significance $p < 0.001$ for all comparisons, very large effect sizes Cohen's $d = 9.15$ – 16.24 , 100% win rate (8/8 datasets)
- Computational efficiency: 2.5ms inference latency (400,000 packets/second throughput, 3.4× faster than Deep Learning 8.5ms), 180 minutes training time (1.58× faster than Deep Learning 285 min), 1.2GB memory footprint (3.2× lower than Deep Learning 3.8GB), 12% CPU utilisation (2.3× lower than Deep Learning 28%), 850 Mbps sustainable throughput (2.66× higher than Deep Learning 320 Mbps)
- Operational robustness: Detection rates exceeding 98.5% for low-intensity attacks, peak 99.7% at moderate intensity, maintaining 96.2% at 70% network congestion and 93.8% at extreme 90% congestion (substantially exceeding typical IDS 60–80% under similar conditions)
- Attack-overload discrimination: 96.8% discrimination accuracy (95.2% sensitivity for attacks, 97.8% specificity for legitimate overload) reducing false positives from 38% (baseline threshold approach) to 2.2% (proposed approach, 17× reduction), addressing fundamental challenge in dynamic virtualised environments where elastic auto-scaling creates ambiguity between attack traffic and legitimate flash crowds
- Vulnerability quantification: Established empirical evidence for virtualisation-specific vulnerabilities including 34% greater attack susceptibility to resource exhaustion vs physical deployments, 21% expanded attack surface from NFV interfaces, 42% network performance degradation from SDN controller saturation, 15% increased lateral movement risk from inter-VM communication channels
- System limits established: Identified operational envelopes (2 TB optimal data volume, 10,000 concurrent devices, 850 Mbps sustainable bandwidth, 80% resource utilisation threshold) providing evidence-based deployment guidelines and capacity planning parameters

These findings validate the research hypothesis that AI mechanisms designed specifically for virtualised IoT characteristics achieve substantial security improvements (99.85% mean accuracy, 94.5–98.2% error reduction vs competitors) whilst maintaining operational efficiency (2.5ms latency, 0.19% false positives) suitable for production deployment. However, limitations were revealed including reduced performance under extreme

conditions (95.3% accuracy on zero-day attacks vs 99.85% baseline, 98.2% on encrypted traffic), performance degradation under severe resource pressure (88.3% at 95% CPU, 92.1% at 90% memory, 85.2% under combined pressure), and scalability constraints (10,000 device limit, 850 Mbps throughput ceiling) requiring distributed architecture for larger deployments.

6.2 Research Contributions

This thesis makes three distinct contributions to virtualised IoT security. First, it establishes a comprehensive understanding of virtualisation-specific DDoS vulnerabilities through systematic analysis and empirical quantification. Second, it develops an efficient AI-based detection mechanism specifically designed to address these identified vulnerabilities. Third, it provides rigorous experimental validation through comparative evaluation against contemporary methods and extensive testing across various operational scenarios and deployment conditions.

6.2.1 Contribution 1: Comprehensive AI-Based Security Framework for Virtualised IoT with Empirical Vulnerability Quantification

The first contribution encompasses systematic vulnerability identification, comprehensive taxonomy development, and AI-based detection mechanism design with extensive empirical validation, recognising these as integrated components of unified research achievement.

Theoretical contribution. Established comprehensive taxonomy of vulnerabilities specific to virtualised IoT environments, distinguishing these from conventional IoT security challenges through empirical quantification: 34% greater attack susceptibility (resource exhaustion amplification through shared infrastructure), 21% absent in physical deployments), 42% network degradation from SDN controller saturation (centralised control plane creating catastrophic single-point failure), 15% increased lateral movement risk (inter-VM communication channels, hypervisor-mediated attacks). This taxonomy provides structured framework for understanding how virtualisation technologies simultaneously enable scalability whilst introducing fundamental attack vectors requiring specialised detection mechanisms.

Methodological contribution. Development and rigorous validation of an AI-based detection framework specifically optimised for virtualised IoT characteristics represents a significant methodological advance. The FCNN architecture (8 virtualisation-aware input features $\rightarrow$ 256 $\rightarrow$ 128 $\rightarrow$ 64 $\rightarrow$ 64 $\rightarrow$ 32 $\rightarrow$ 1 output, 49,729 parameters, strategic

dropout regularisation 60%→50%→55%→40%→50%) achieved $99.85 \pm 0.04\%$ mean accuracy across 8 diverse datasets with exceptional stability ($CV = 0.04\%$, $75\text{--}286\times$ lower variance than typical neural networks), $99.92 \pm 0.007\%$ on Generated vIoT (HIGHEST among all datasets validating virtualisation-specific optimisation), and $+2.57$ to $+5.54$ pp improvements over all competing methods with statistical significance $p < 0.001$ and very large effect sizes $d = 9.15\text{--}16.24$.

Four critical methodological innovations ensure robust and reproducible results: (1) comprehensive baseline characterisation documenting normal network behaviour (168 hours continuous operation, 25.2 pps baseline, resource utilisation profiles, latency characteristics) prior to attack simulation; (2) systematic multi-dataset evaluation across 8 benchmarks spanning 6 years (2019–2024), 33+ attack types, varied deployment contexts (traditional networks, SDN environments, edge computing, medical IoT, industrial sensors) demonstrating robust generalisation (consistent 99.79–99.92% accuracy range, span 0.13 pp); (3) rigorous statistical validation through 5 independent runs per dataset (40 total training runs) with narrow confidence intervals (± 0.014 to ± 0.042 pp, 95% CI) and complete absence of outliers confirming reproducibility; (4) comprehensive attack-overload discrimination methodology achieving 96.8% discrimination accuracy addressing practical challenge unique to elastic virtualised environments.

Practical impact. The framework provides validated approach for securing smart city and industrial IoT infrastructures with demonstrated production-ready characteristics: exceptionally low false positive rate 0.19% (0.11% on vIoT, enabling automated response without excessive false alarms), minimal false negative rate 0.10% (0.05% on vIoT, providing strong security assurance), real-time inference 2.5ms latency (400,000 packets/second throughput exceeding typical vIoT traffic rates by $26\text{--}15,873\times$), modest resource consumption 1.2GB memory + 12% CPU (deployable on standard virtualisation infrastructure without dedicated hardware), and sustainable throughput 850 Mbps (adequate for typical deployments with $1.7\times$ headroom). Systematic documentation of baseline performance metrics, operational envelopes (2 TB data, 10,000 devices, 850 Mbps, 80% resources), and capacity planning guidelines enables evidence-based deployment across diverse organisational contexts.

6.2.2 Contribution 2: Novel Virtualised IoT Testbed with Comprehensive Baseline Characterisation and Reproducible Experimental Protocol

The second contribution comprises the design, implementation and validation of a virtualised IoT testbed with extensive baseline performance characterisation. This systematic documentation of normal network behaviour prior to attack simulation provides the empirical foundation for accurate anomaly detection and threshold determination.

Architectural contribution. Testbed architecture integrates multiple virtualisation technologies (KVM/QEMU 9.0 hypervisor, Open vSwitch 3.3 virtual switch, OpenDaylight Carbon SDN controller, Docker containers, NFV components) replicating layered complexity of production vIoT deployments. Design incorporates realistic traffic generation emulating diverse IoT device behaviours (smart home sensors 25 pps, industrial monitors 100 pps, healthcare devices 50 pps, environmental sensors 10 pps with varying protocols: MQTT, CoAP, HTTP/S, AMQP), enabling precise control over network parameters (bandwidth 100 Mbps–10 Gbps, latency 1–100 ms, packet loss 0–10%) whilst maintaining experimental isolation ensuring reproducibility.

Baseline characterisation. Extensive baseline experiments (168 hours continuous operation) documented normal behaviour establishing foundation for meaningful attack detection evaluation: (1) traffic flow patterns under benign conditions (packet size distributions: mean 250 bytes, range 64–1500 bytes; inter-arrival times: exponential distribution with mean 40ms; protocol usage: 60% MQTT, 25% HTTP/S, 10% CoAP, 5% other); (2) resource utilisation metrics across operational loads (CPU: 8–15% normal, peak 45% during legitimate surges; memory: 2–4GB normal, peak 8GB; network bandwidth: baseline 258.7 Kbps, peak 15 Mbps during firmware updates); (3) latency characteristics (mean 2.5ms sensor-to-cloud, 95th percentile 4.2ms, 99th percentile 8.7ms); (4) performance thresholds distinguishing normal operational variance from anomalous behaviour (traffic volume: $\pm 30\%$ diurnal variation normal, $>50\%$ sustained increase anomalous; CPU: $<60\%$ normal, $>80\%$ sustained anomalous; latency: $<10\text{ms}$ normal, $>25\text{ms}$ sustained anomalous).

These comprehensive baseline metrics, previously absent from vIoT literature, provide essential foundation for (1) establishing appropriate anomaly detection thresholds avoiding false positives during legitimate operational variance; (2) validating attack simulation realism by comparing against documented normal patterns; (3) enabling meaningful interpretation of attack impact (e.g., 42% network degradation from SDN controller saturation quantified relative to 25.2 pps baseline throughput); (4) facilitating reproducibility through detailed normal behaviour documentation enabling other researchers to validate that their testbeds exhibit comparable baseline characteristics.

Validation and reproducibility. Testbed validity established through comparison with physical IoT deployments (traffic pattern correlation $r=0.87$ with real smart home deployment, latency distribution KS-test $p=0.23$ indicating no significant difference from real industrial sensors) and published traffic datasets (packet size distribution similarity to IoT-23 real botnet traffic χ^2 -test $p=0.18$). Comprehensive documentation of experimental protocols (complete network topology diagrams, device configuration files, traffic generation parameters, SDN flow rules, NFV service chain specifications), environmental configurations (hardware specifications: Intel Core i7-12700 12-core CPU, 32GB RAM,

NVIDIA RTX 3090 GPU; software versions: Ubuntu 24.04 host OS, KVM/QEMU 9.0, OVS 3.3, ODL Carbon; network settings: 1 Gbps link capacity, MTU 1500 bytes), and parameter settings (learning rate 0.001, batch size 64, dropout rates, L2 regularisation 0.001, training epochs 100) enables independent replication of results by other researchers.

6.2.3 Contribution 3: Comprehensive Comparative Evaluation with Contemporary Datasets and Practical Deployment Validation

The third contribution provides comprehensive experimental validation through: (1) direct experimental comparison with three contemporary competing methods (Deep Learning LSTM, Random Forest, SVM) on identical datasets using identical evaluation protocols, ensuring fair and reproducible performance assessment; (2) evaluation using recently published 2024–2025 datasets not previously explored in virtualised IoT DDoS detection contexts; (3) comprehensive bottleneck analysis and system limits characterisation identifying operational thresholds for practical deployment planning.

Direct comparative evaluation. Implementation and rigorous testing of three competing methods (Deep Learning LSTM with attention mechanism, Random Forest with 500 trees, SVM with RBF kernel) on identical datasets using identical train/test splits (70/15/15) and preprocessing pipelines (8-feature extraction, min-max normalisation, SMOTE balancing) enabled controlled comparison isolating performance differences attributable to model architecture. Results demonstrated consistent superiority: +2.57 pp vs Deep Learning (99.85% vs 97.28%, 94.5% error reduction from 2.72% to 0.15%), +4.29 pp vs Random Forest (96.6% error reduction), +5.54 pp vs SVM (97.4% error reduction), with statistical significance $p < 0.001$ for all comparisons, very large effect sizes $d = 9.15$ – 16.24 , and 100% win rate.

Computational efficiency analysis revealed multidimensional advantages: $3.4\times$ faster inference (2.5ms vs 8.5ms Deep Learning) whilst maintaining +2.57 pp higher accuracy, $1.58\times$ faster training (180 min vs 285 min enabling 87.5 hours saved over 50 development iterations), $3.2\times$ lower memory (1.2GB vs 3.8GB enabling deployment without dedicated hardware), $2.3\times$ lower CPU (12% vs 28% reducing interference with co-located workloads), $2.66\times$ higher throughput (850 Mbps vs 320 Mbps supporting larger deployments), demonstrating that domain-specific architectural optimisation overcomes traditional accuracy-efficiency trade-offs.

Contemporary dataset evaluation. Validation on recently published datasets (CIIoMT2024 featuring 2024 medical IoT traffic with modern attack patterns, EdgeIIoT-set containing 2024 smart home and industrial traffic) spanning previously unexplored

attack evolution (5G-enabled device attacks, edge computing exploitation, advanced persistent threats) demonstrated robust generalisation: CICIoMT2024 achieved $99.90 \pm 0.009\%$ (second-highest among all datasets, validating cross-domain generalisation to safety-critical medical applications), EdgeIIoTset achieved $99.88 \pm 0.011\%$ (third-highest, confirming effectiveness on edge computing scenarios). Performance maintained across temporal evolution (datasets spanning 2019–2024 demonstrating consistent 99.79–99.92% accuracy despite 6-year attack landscape evolution) validates long-term effectiveness, though temporal drift analysis revealed 3.12 pp degradation after simulated 6-month evolution suggesting quarterly/bi-annual retraining advisable.

Bottleneck analysis and deployment guidelines. Systematic characterisation of system limits and performance degradation thresholds enables accurate distinction between attack signatures and legitimate resource exhaustion, preventing false positives during high-load operations. Comprehensive testing established: (1) data volume limits (optimal 99.85%+ accuracy up to 2 TB representing typical small-to-medium deployments, acceptable 98.5%+ up to 5 TB, degraded 96%+ up to 7 TB, critical <96% beyond 7 TB requiring distributed architecture); (2) concurrent device limits (optimal 99.5%+ throughput up to 5,000 devices, graceful degradation 5,000–10,000 devices, critical limit 10,000 devices beyond which performance accelerates to 78.2% at 11,500 devices with memory exhaustion risk); (3) bandwidth thresholds (optimal 99.85%+ up to 850 Mbps sustainable throughput, acceptable 98.5%+ up to 2 Gbps, degraded 96.8% at 5 Gbps, critical 88.2% at 10 Gbps with packet loss >20%); (4) resource utilisation envelopes (optimal <70% CPU/memory maintaining 99.85%+ accuracy with headroom for bursts, acceptable 70–80% maintaining 99–99.5%, warning 80–90% frequent degradation requiring capacity expansion planning, critical >90% severe degradation to 85–92% requiring immediate intervention).

Attack-overload discrimination methodology developed comprehensive framework for distinguishing malicious attacks from legitimate resource exhaustion: temporal pattern analysis (rise time measurement: legitimate 15–60 min gradual onset vs attack <5 min sudden spike achieving 94.2% discrimination using rise time alone), statistical feature analysis (5 discriminating characteristics: packet rate variance, source IP diversity, connection duration, payload size distribution, protocol diversity achieving clear separation between legitimate balanced features vs attack skewed features), detection confidence trajectory monitoring (legitimate consistently low <0.10 probability throughout event vs attack rapid escalation to >0.95 within 2 minutes achieving 97.3% discrimination using confidence trajectory alone), combined achieving 96.8% overall discrimination (95.2% sensitivity, 97.8% specificity) with $17\times$ false positive reduction from 38% baseline to 2.2% proposed, validated across 200 legitimate overload events and 150 attack scenarios.

6.3 Limitations and Threats to Validity

Transparent acknowledgement of limitations provides appropriate context for interpreting results whilst identifying specific future research directions.

6.3.1 Technical and Methodological Limitations

Performance under extreme conditions: Detection accuracy degrades under challenging scenarios including zero-day attacks (95.3% vs 99.85% baseline, 4.57 pp degradation representing $30\times$ error rate increase from 0.15% to 4.70%), encrypted traffic (98.2% vs 99.85% baseline, 1.65 pp degradation representing $11\times$ error rate increase from 0.15% to 1.80% due to flow-based features lacking payload inspection), severe resource pressure (accuracy declining to 88.3% at 95% CPU utilisation, 92.1% at 90% memory utilisation, 85.2% under combined CPU+memory pressure representing $76\text{--}98\times$ error rate increases), and extreme network congestion (91.3% at 8 Gbps, 88.2% at 10 Gbps traffic loads substantially exceeding 850 Mbps design capacity).

Scalability constraints: Testbed evaluation maximum 12,000 devices cannot definitively validate performance at large production scales (tens/hundreds of thousands devices in smart city infrastructures). Identified limits (10,000 concurrent devices, 850 Mbps sustainable throughput, 2 TB optimal data volume) necessitate distributed architecture (horizontal scaling, load balancing, data sharding) for larger deployments, though empirical validation at 100,000+ device scale remains future work. Bottleneck analysis establishes theoretical scaling approaches but cannot confirm absence of unexpected emergent phenomena at massive scale (synchronization effects, cascading failures, distributed race conditions).

Testbed-reality gap: Simulated traffic based on documented IoT models may not fully capture real device complexity (firmware bugs creating protocol anomalies, middleware interactions, unpredictable user behaviours, equipment aging effects). Whilst multi-dataset evaluation including real-world captured traffic (IoT-23 real botnet traffic 99.79%, CICIoMT2024 real medical device captures 99.90%) partially validates generalization, laboratory environment lacks production unpredictability (equipment failures, network anomalies, operational constraints including SLA requirements, compliance obligations, integration with legacy infrastructure).

Adversarial robustness: Models show theoretical vulnerability to adversarial machine learning attacks where sophisticated attackers deliberately craft traffic patterns targeting detection weaknesses through gradient-based optimization or model inversion. Whilst such attacks currently rare in practice (requiring substantial attacker sophistication, model access, domain expertise), proliferation of AI-based security mechanisms may incentivize adversarial techniques development, representing emerging threat requiring

future investigation through adversarial training, certified defences, and meta-detection of adversarially-crafted inputs.

Adaptivity requirements: Model effectiveness degrades over time as attack patterns evolve temporal drift analysis demonstrated 3.12 pp degradation after simulated 6-month evolution (96.8% vs 99.85% baseline), suggesting quarterly/bi-annual retraining necessary. Retraining requires labeled attack data (potentially difficult obtaining for novel attacks), computational resources (180 minutes on 2.5M samples), and operational coordination (model deployment, validation, rollback capability). Continuous learning approaches (incremental updates without full retraining) and transfer learning (adapting to new contexts with limited data) represent promising but unexplored enhancement directions.

6.3.2 Generalisability and Scope Constraints

Platform specificity: Testbed employed specific technologies (KVM/QEMU 9.0, Open vSwitch 3.3, OpenDaylight Carbon, Ubuntu 24.04) potentially limiting generalisation to alternative platforms (VMware vSphere, Microsoft Hyper-V, ONOS controller, Docker/Kubernetes containerisation). Feature extraction focuses on hypervisor-agnostic metrics suggesting reasonable portability, though platform-specific validation remains necessary before cross-platform deployment recommendations.

Dataset limitations: Some benchmarks represent historical traffic (CIC-DDoS2019 now 6 years old, IoT-23 from 2020 now 5 years old) potentially missing contemporary techniques (recent DDoS amplification vectors, post-2020 IoT malware variants, emerging protocol exploits, HTTP/3 adoption, QUIC deployment, IPv6 transition). Inclusion of recent datasets (CICIoMT2024, CICIoT2023, Generated vIoT 2024–2025) partially mitigates temporal gaps, though continuous dataset refresh necessary for maintaining contemporary relevance. Ground truth labels may contain errors (estimated 1–5% noise typical in published datasets) affecting evaluation reliability, class imbalance (attack:benign ratios 1:10 to 1:1 across datasets) potentially biasing metrics despite SMOTE mitigation, coverage gaps (33+ attack types still subset of documented DDoS techniques) creating inevitable evaluation limitations.

Geographic and contextual constraints: Testing in specific network conditions (UK-based academic infrastructure, specific ISP characteristics) may not generalise across all deployment regions (varying infrastructure quality, different regulatory frameworks affecting logging/monitoring, diverse privacy constraints, regional attack landscape variations). Cross-dataset evaluation including international benchmarks (Canadian CIC, Czech IoT-23, US EdgeIoTset) provides some geographic diversity, though comprehensive regional validation necessary before universal deployment claims.

6.4 Future Research Directions

The findings suggest numerous investigation avenues addressing identified limitations whilst extending contributions to emerging contexts.

6.4.1 Enhanced Detection Mechanisms Addressing Current Limitations

Zero-day attack detection: Develop hybrid approaches combining supervised classification (for known attacks achieving 99.85% accuracy) with unsupervised anomaly detection (for novel threats currently 95.3% accuracy). Promising directions include: (1) generative models (variational autoencoders, generative adversarial networks) identifying traffic deviating from normal behaviour without requiring attack signatures; (2) one-class classification and novelty detection algorithms systematic investigation in vIoT contexts; (3) meta-learning enabling rapid adaptation to emerging threats with minimal labeled examples through few-shot learning paradigms; (4) transfer learning from general network security domains to vIoT-specific contexts.

Encrypted traffic analysis enhancement: Current 98.2% accuracy on encrypted traffic (1.65 pp below 99.85% baseline) motivates enhanced feature engineering: (1) TLS handshake characteristics analysis (certificate properties, cipher suite negotiation, session resumption patterns); (2) encrypted protocol fingerprinting (packet size sequences, timing patterns unique to protocols despite encryption); (3) metadata-based discrimination (connection establishment patterns, session duration distributions, bidirectional flow ratios); (4) traffic flow graph analysis (multi-flow relationships, temporal dependencies, communication graph topology). Integration with privacy-preserving techniques (homomorphic encryption, secure multi-party computation) enables analysis without compromising data confidentiality.

Adversarial robustness: Develop defences against adversarial machine learning attacks through: (1) adversarial training incorporating gradient-based perturbations during training improving model resilience; (2) certified defences providing provable robustness guarantees within specified perturbation bounds; (3) ensemble diversity strategies where multiple models with different architectures, training data, feature sets collectively resist coordinated evasion; (4) meta-detection identifying adversarially-crafted inputs through statistical anomalies in feature distributions or prediction confidence patterns; (5) game-theoretic frameworks modeling strategic attacker-defender interaction informing optimal defensive strategies under adaptive adversaries.

Continuous adaptation and learning: Address temporal drift (3.12 pp degradation after 6 months) through: (1) incremental learning updating models with new traffic data without full retraining (reducing computational overhead from 180 minutes to minutes while

avoiding catastrophic forgetting); (2) online learning algorithms adapting continuously to traffic evolution in real-time; (3) active learning identifying most informative samples for labeling (minimizing manual labeling effort while maximizing model improvement); (4) domain adaptation techniques transferring knowledge across deployment contexts with different traffic characteristics.

6.4.2 Scalability and Distributed Architecture

Distributed detection frameworks: Address identified limits (10,000 devices, 850 Mbps) through distributed architecture: (1) hierarchical detection (lightweight edge models performing initial screening forwarding suspicious traffic to centralized sophisticated analysis, balancing detection latency, bandwidth consumption, computational efficiency); (2) horizontal scaling (load balancing across multiple detection nodes each operating within optimal 850 Mbps range enabling linear capacity increase); (3) data sharding strategies (partitioning by source IP, time window, flow hash distributing load whilst maintaining detection effectiveness); (4) consistent hashing and distributed state management enabling >100,000 device scale.

Edge computing integration: Investigate hierarchical approaches: (1) model partitioning optimization determining which detection components execute at edge vs cloud layers balancing latency, bandwidth, computation; (2) edge-based local analysis without cloud transmission addressing privacy concerns while reducing network load; (3) federated learning training models across distributed edge deployments whilst preserving data privacy through secure aggregation protocols resistant to poisoning attacks; (4) dynamic workload distribution adapting to edge resource availability and network conditions.

Performance optimization at scale: Systematic investigation of: (1) GPU/TPU acceleration reducing inference latency below 2.5ms enabling higher throughput; (2) model quantization and pruning reducing memory footprint below 1.2GB enabling deployment on resource-constrained edge nodes; (3) batch processing optimization trading latency for throughput in high-load scenarios; (4) approximate computing techniques trading minor accuracy reductions for substantial efficiency gains in resource-critical deployments.

6.4.3 Integration with Emerging Technologies and Standards

5G/6G network integration: Ultra-low latency (1ms), massive device density (1M devices/km²), network slicing create opportunities and challenges: (1) network slicing leveraging for traffic isolation (security-critical IoT in dedicated slices with guaranteed resources); (2) ultra-low latency constraints on detection algorithm design (requiring <1ms inference vs current 2.5ms); (3) massive scale adaptation (handling 1M+ devices/km² vs current 10,000 device limit); (4) virtualised network function integration

(deploying detection as native VNF within 5G/6G architecture); (5) mobile edge computing exploitation for distributed detection at network edge.

Blockchain for security enhancement: Investigate hybrid architectures: (1) blockchain-based authentication and authorization (decentralized identity management, tamper-proof access control); (2) immutable audit trails for security events enabling forensic analysis and compliance verification; (3) smart contracts for automated security policy enforcement and incident response; (4) distributed consensus for collaborative security decisions across multiple stakeholders; (5) lightweight blockchain variants specifically designed for IoT addressing computational overhead and latency concerns whilst maintaining security benefits.

Explainable AI and trust: Address interpretability limitations (current 49,729-parameter neural network lacks transparency) through: (1) SHAP (SHapley Additive exPlanations) values quantifying individual feature contributions to predictions; (2) attention mechanisms visualizing which traffic characteristics drive detection decisions; (3) decision tree surrogates approximating neural network behaviour with interpretable rules; (4) counterfactual explanations identifying minimal traffic modifications changing classifications; (5) uncertainty quantification providing confidence estimates alongside predictions enabling human oversight for low-confidence decisions.

6.4.4 Cross-Domain Applications and Specialised Deployments

Smart healthcare adaptation: Healthcare IoT (remote patient monitoring, connected medical devices, hospital infrastructure) requires domain-specific adaptations: (1) safety-critical latency requirements (millisecond response for critical care devices vs current 2.5ms potentially insufficient for life-support systems); (2) patient privacy protections (HIPAA compliance, federated learning avoiding raw data sharing); (3) clinical workflow integration (minimizing false positives disrupting medical procedures); (4) medical device-specific attack patterns (insulin pump manipulation, pacemaker attacks, infusion pump exploitation); (5) regulatory compliance frameworks (FDA cybersecurity guidance, medical device certification requirements).

Industrial IoT (IIoT) specialization: Industrial environments exhibit distinct characteristics: (1) deterministic communication patterns (periodic sensor readings, synchronized control commands enabling physics-informed detection models); (2) specialized industrial protocols (Modbus, OPC UA, PROFINET requiring protocol-specific feature extraction); (3) stringent real-time requirements (microsecond control loops in manufacturing vs current 2.5ms potentially excessive); (4) operational technology (OT) security integration with information technology (IT) security addressing OT/IT convergence vulnerabilities; (5) safety-critical implications where security incidents directly impact

physical safety (chemical plant control, power grid operation, autonomous vehicle coordination).

Critical infrastructure protection: Energy grids, water systems, transportation, telecommunications require high-assurance mechanisms: (1) formal verification methods providing provable security guarantees beyond empirical 99.85% accuracy; (2) graceful degradation architectures maintaining basic functionality even under sustained attack; (3) resilience metrics quantifying infrastructure robustness to cascading failures; (4) defense-in-depth integration coordinating multiple security layers (physical, network, application, data); (5) regulatory compliance addressing sector-specific requirements (NERC CIP for energy, TSA for transportation, sector-specific NIST frameworks).

6.5 Concluding Remarks

This research addressed critical security challenges in virtualised IoT environments through AI-based DDoS detection achieving $99.85 \pm 0.04\%$ mean accuracy across 8 diverse datasets (99.92% on Generated vIoT, HIGHEST among all datasets) with exceptional computational efficiency (2.5ms latency, 1.2GB memory, 12% CPU, 850 Mbps throughput), exceptionally low error rates (0.19% false positives, 0.10% false negatives), and robust performance across diverse conditions while maintaining 100% win rate against all competing methods (+2.57 to +5.54 pp improvements, $p < 0.001$, $d = 9.15\text{--}16.24$).

The three contributions (1) comprehensive AI-based security framework with empirical vulnerability quantification (34% attack susceptibility, 21% attack surface, 42% network degradation, 15% lateral movement risk); (2) novel virtualised testbed with comprehensive baseline characterisation (168 hours continuous operation documenting normal behaviour prior to attack simulation); (3) rigorous comparative evaluation with contemporary datasets and practical deployment validation (direct experimental comparison, bottleneck analysis establishing operational envelopes: 2 TB data, 10,000 devices, 850 Mbps, 80% resources, attack-overload discrimination 96.8% accuracy with $17\times$ false positive reduction) advance both theoretical understanding and practical deployment capabilities.

Six critical enhancements substantially strengthened the research: (1) consolidation of first two contributions improving conceptual clarity; (2) comprehensive baseline characterisation establishing empirical foundations for threshold determination; (3) direct comparative evaluation with three implemented competitors (LSTM, Random Forest, SVM) ensuring reproducible performance assessment; (4) integration of contemporary 2024–2025 datasets demonstrating temporal robustness; (5) comprehensive bottleneck

analysis identifying operational limits for deployment planning; (6) attack-overload discrimination methodology addressing practical false positive challenges. These enhancements, whilst undertaken within full-time employment constraints, produced research of substantially greater experimental rigor and practical applicability.

However, acknowledged limitations provide context for achievements: reduced performance under extreme conditions (95.3% zero-day vs 99.85% baseline, 98.2% encrypted traffic, 85–92% under severe resource pressure), scalability constraints requiring distributed architectures for large-scale deployments (10,000 concurrent device limit, 850 Mbps throughput ceiling), testbed-to-production gap despite multi-dataset validation, adversarial robustness vulnerabilities requiring defensive mechanisms, temporal drift necessitating periodic retraining (3.12 pp degradation after 6 months), and generalisability constraints across diverse platforms, operational contexts, and deployment scales.

Future research directions span enhanced detection mechanisms (zero-day improvements, encrypted traffic analysis, adversarial robustness, continuous adaptation), scalability solutions (distributed architectures, edge computing integration, performance optimisation), emerging technology integration (5G/6G networks, blockchain, explainable AI), and cross-domain applications (healthcare, industrial IoT, critical infrastructure). These directions address identified limitations whilst extending contributions to evolving technological landscapes.

The broader societal context accelerating IoT deployment in critical smart city and industrial applications serving growing urban populations (60% globally by 2050) underscores research importance. DDoS attacks disrupting healthcare, transportation, energy grids pose risks transcending technical concerns to impact citizen safety and wellbeing. This research contributes to essential securing of these systems, demonstrating that AI-based mechanisms can achieve exceptional detection effectiveness (99.85% accuracy, 0.19% false positives, 0.10% false negatives, 2.5ms latency) whilst maintaining practical operational viability (96.8% attack-overload discrimination preventing false positives during legitimate surges, established operational limits enabling evidence-based deployment planning).

In conclusion, this thesis establishes that AI-based detection mechanisms specifically designed for virtualised IoT characteristics substantially enhance security (94.5–97.4% error reduction vs competing methods, statistical significance $p < 0.001$, very large effect sizes $d = 9.15\text{--}16.24$) whilst maintaining computational efficiency ($3.4\times$ faster inference, $3.2\times$ lower memory, $2.66\times$ higher throughput vs Deep Learning) and operational feasibility (exceptionally low false positive/negative rates enabling automated response, comprehensive baseline characterisation and bottleneck analysis providing deployment guidelines, robust generalisation across 8 datasets spanning 6 years and 33+ attack types). The contributions provide both theoretical insights into virtualisation-specific vulnerabilities and practical tools for protecting IoT infrastructures, offering valuable

foundations for developing secure, resilient systems that future smart cities and industrial applications demand.

Appendix A

Smart City Framework and Context

A.1 Overview

This appendix provides detailed classifications of smart city domains and services referenced in Chapter 1, Section 1.1. The Internet of Things plays a fundamental role in enabling smart city infrastructures across multiple domains, including governance, environment, mobility, economy, living standards, and healthcare. These tables offer comprehensive context for understanding the application domain of this research but are too detailed for inclusion in the main introduction.

The security mechanisms developed in this thesis are directly applicable to protecting the IoT infrastructures that underpin these smart city services, where availability, reliability, and resilience against DDoS attacks are critical requirements.

A.2 I.1 Smart City Resource Domains

The smart city resource domain can be classified into three different subdomains based on several characteristics. The table below shows the various subdomains and the main features of each part (Hamamurad et al. (2021)).

TABLE A.1: Smart City Resource Domain Classifications

Sub-domain	Characteristics
Supra-structure	Governance, management, organisational structure, performance assessment, auditing, change administration, political commitment, vision, mission, strategy, and oversight, including regulation. Comprehensive planning, growth strategies, roadmaps, action plans, and standardised processes. Business modelling, optimisation, process coordination, and solutions.
Infrastructure	IT application networks and infrastructure; Internet access, connectivity, and high-speed communication services; Data storage systems, data servers, and cloud-based services.
Infostructure	Computing, instrumentation, information systems, communication infrastructure, connectivity, accessibility, interoperability, integration, innovation, and security; Data sensing, collection, management, analysis, simulation, open data, security, and privacy considerations.

A.3 I.2 Smart City Service Classifications

The smart city resource domain classifications are a base framework for the services domain. The services framework in the smart city domain is related to the proportions of intelligent environments expressing raw resources, involving aspects of natural emissions, environmental conservation, and sustainable use of resources ([Hamamurad et al. \(2021\)](#)). Table A.2 shows the various subdomains and the main features of each part of the service domain.

TABLE A.2: Smart City Sub-domains and Characteristics

Sub-domain	Characteristics
Environment	- Sustainable management of natural and power resources - Eco-friendly construction, solar energy, and energy-efficient street lighting - Expansion of power and energy sources - Pollution prevention, environmental preservation, and waste and water management
Economy	- Fostering creativity and innovation - Encouraging entrepreneurship and showcasing technology - Local, regional, and global interconnectedness - Branding, advertising, and marketing services - Enhancing workforce productivity and collaboration - Financial support and mobile money transfers
Governance	- Political determination, stakeholder participation in decision-making, and supply-demand choices - Urban operations and strategic planning - Administration of public utilities and supplementary services - Utilising ICT, e-procurement, e-governance, and promoting transparency through open data
People	- Healthcare, online learning, continuous education, and human resources management - Collaboration among academic institutions, government, and civil society - Diverse communities, racial harmony, smart solutions, and social initiatives - Embracing adaptability, creativity, and engagement in public administration
Mobility	- Robust IT infrastructure and seamless integration of communication technologies - Accessibility at local, regional, and national levels - Multi-modal transportation options and infrastructure development - Facilitating societal movement - Promoting sustainability, safety, and growth in transportation systems
Living	- Access to quality healthcare and medical services - Social inclusion, community resilience, public security, and emergency response - Urban development, public space maintenance, housing quality, facilities management, and building services

A.4 I.3 Security Implications for Smart City IoT Infrastructures

The diverse service requirements outlined in Tables A.1 and A.2 highlight the critical need for robust security mechanisms in smart city IoT deployments. Key security challenges include:

A.4.1 I.3.1 Service Availability Requirements

Smart city services have varying availability requirements:

- **Critical Services** (99.99% uptime): Emergency response, healthcare monitoring, energy grid management
- **High-Priority Services** (99.9% uptime): Traffic management, public transport, security surveillance
- **Standard Services** (99% uptime): Environmental monitoring, waste management, non-critical sensors

DDoS attacks pose a direct threat to these availability requirements, potentially causing:

- Service disruptions affecting public safety
- Economic losses due to infrastructure downtime
- Degraded quality of life for citizens
- Cascading failures across interdependent systems

A.4.2 I.3.2 Virtualisation in Smart City Deployments

Modern smart city architectures increasingly rely on virtualisation technologies:

- **Network Function Virtualisation (NFV)**: Virtualised network services reduce hardware costs but introduce new attack surfaces
- **Software-Defined Networking (SDN)**: Centralised control enables flexible management but creates single points of failure
- **Cloud-Based Processing**: Scalable data analytics require robust security at the virtualisation layer
- **Edge Computing**: Distributed processing nodes must be protected against distributed attacks

A.4.3 I.3.3 Application Context for This Research

This research focuses on detecting and mitigating DDoS attacks in virtualised IoT environments, with particular relevance to smart city infrastructures due to:

1. **Scale and Heterogeneity:** Smart cities deploy thousands of heterogeneous IoT devices across multiple domains, creating complex attack surfaces.
2. **Critical Dependencies:** Many smart city services depend on continuous IoT data streams; interruptions can have immediate public safety implications.
3. **Virtualisation Prevalence:** Cost pressures drive smart city deployments toward virtualised infrastructures, introducing the security challenges addressed in this thesis.
4. **Real-Time Requirements:** Services such as traffic management and emergency response require low-latency detection and mitigation of attacks.
5. **Interconnected Systems:** Smart city domains are increasingly interconnected; an attack on one system (e.g., traffic management) can cascade to others (e.g., emergency services).

A.4.4 I.3.4 Relevance to Research Questions

The smart city context directly informs the three research questions addressed in Chapter 1:

- **RQ1 (Vulnerabilities):** Smart city virtualised infrastructures exhibit unique vulnerabilities due to their scale, heterogeneity, and resource-sharing architectures.
- **RQ2 (AI Mechanisms):** The proposed AI-based detection framework must handle the diversity of traffic patterns across multiple smart city service domains.
- **RQ3 (Performance):** Detection mechanisms must achieve high accuracy while maintaining low latency to meet real-time requirements of critical smart city services.

A.5 I.4 Statistical Context: Urban Growth and IoT Deployment

As referenced in Chapter 1, Section 1.1:

- By 2050, more than 60% of the world’s population will live in urban areas [Yaqoob et al. \(2017\)](#); ?.
- Forecasts predict more than 25 billion connected IoT devices by 2030, with a significant proportion deployed in smart city infrastructures.
- Global spending on smart city initiatives is projected to exceed \$327 billion by 2025, with cybersecurity representing a critical investment area.
- DDoS attacks targeting critical infrastructure increased by 154% between 2020 and 2023, highlighting the growing threat landscape.

These trends underscore the urgency of developing robust security mechanisms for virtualised IoT environments in smart cities.

A.6 I.5 Summary

This appendix provides the detailed smart city framework and context that supports the motivation and application domain described in Chapter 1. The security challenges inherent in smart city IoT deployments, particularly those using virtualised infrastructures, directly motivate the research contributions of this thesis.

The AI-based DDoS detection mechanism developed in this work addresses critical gaps in protecting virtualised IoT environments that underpin essential smart city services in governance, environment, mobility, economy, living standards, and healthcare domains.

Appendix B

Additional IoT Architecture Details

This appendix provides supplementary information on IoT architectures, including detailed historical context and additional architectural considerations not included in the main body of Chapter 2.

B.1 IoT Architecture

In recent days, industrial firms have proposed several applications related to the IoT. Most of these proposals adopt cyberphysical systems (CPS) and machine-to-machine communications (M2M). These fields also deal with sensitive data. The IoT platform uses the conventional Internet through wireless sensor networks and mobile networks to provide flexibility and scalability (Al-Qaseemi et al. (2017)). Furthermore, IoT applications deal with sensitive data, especially those related to E-healthcare and E-government. Consequently, the IoT platform faces security issues and challenges that are easy to achieve in traditional Internet-based systems. Therefore, new algorithms and technologies are needed to achieve greater satisfaction with security requirements (Rahul et al. (2017)). The security in the traditional Internet requires essential properties such as confidentiality, integrity, and identification. Fulfilling these requirements on the standard Internet is more manageable than achieving them on IoT platforms, as the traditional Internet platform deals with high performance devices (Rahul et al. (2017)). The IoT platform consists of devices with limited resources based on the conventional Internet. Still, it applies to critical and sensitive areas of the national economy, for example, intelligent transportation and health care systems. Thus, the security of information and networks on the IoT platform requires greater availability and dependability. In general, the IoT architecture consists of three different main layers, and these layers simplify into five layers (Mazon-Olivo and Pan (2022)). The main layers of the IoT architecture are

(perception, network, and application), and modern architecture contains the following five layers (perception, access gateway, network, support or middleware, and application); the following points address the definition of each level here:

- The perception layer: is known as a physical or sensing layer as it contains devices and sensors that collect data from material things. This level works as collectors for raw data through the embedded sensors to pass them to the next layer (Mazon-Olivo and Pan (2022)).
- The access gateway layer: is known as the managing layer, as it works as a manager for the perception layer. The access gateway layer is primarily responsible for managing communications in the physical environment and helping to exchange data between systems and things (Antão et al. (2018)).
- The network layer: is known as the connection layer, consisting of connection technologies varying from wired and wireless connections. This layer is responsible for providing an open path between the physical layer and the application layer, allowing the environment to provide services, send decisions from the application layer to the physical layer, and collect data accordingly (Mazon-Olivo and Pan (2022)).
- The middleware layer: is known as a support layer. This layer works as a support layer on the network layer side to help connect low-level communication technologies with high-level technologies and provides greater flexibility and unity between hardware and application levels (Al-Qaseemi et al. (2017)).
- The application layer: is known as the platform and integration layer. This layer is responsible for delivering the functionality from the infrastructure to the business layer. This layer contains processing and business sublayers (Mazon-Olivo and Pan (2022)).

each has security and management issues. Figure B.1 shows the IoT architecture. The following subsections will illustrate the cloud concept in the IoT aspect and the need for IoT virtualisation.

B.2 Network Functions Virtualisation (NFV)

The necessity for robust, high-bandwidth network connectivity has experienced substantial escalation in recent years. This growth has been further propelled by the proliferation of devices connected to the Internet, encompassing smartphones, tablets, laptops, sensor networks, and Machine-to-Machine (M2M) connectivity. Network traffic has surpassed the capacities afforded by existing mobile service providers' networks. In light of projections indicating continued growth in network traffic, it is imperative that Network




FIGURE B.1: IoT Architecture

Service Providers (NSPs) allocate resources to bandwidth-focused infrastructure to meet the growing demand (Mijumbi et al. (2016)). Although research indicates that such investments yield minimal capital returns, network upgrades are inextricably linked to network infrastructure. This relationship, coupled with the exponential cost of network equipment, may result in reduced profit margins for NSPs when launching upgrades or new services. The challenges faced by NSPs extend beyond the cost of high-priced hardware and include the increase in energy costs and the scarcity of personnel possessing the requisite expertise to design, implement, and orchestrate increasingly complex hardware-based infrastructure (Hawilo et al. (2019))(Han et al. (2015)). In addition, the maintenance of the network infrastructure constitutes a primary concern for service providers. The ramifications of these issues extend beyond revenue loss and manifest themselves as delays in time-to-market and general impediments to innovation within the telecommunications sector. Consequently, network operators strive to diminish or eliminate their reliance on proprietary hardware. In pursuit of these objectives, NSPs explore the incorporation of virtualisation technology within the telecommunications industry. Virtualisation technology has emerged as a means for Information Technology (IT) specialists to increase capital investments with higher capital returns while facilitating the decoupling of hardware and software, allowing multiple isolated software programmes to share the underlying hardware (Han et al. (2015)). As an initial measure, a consortium of seven telecommunication operators established an industry specification group for Network Function Virtualisation (NFV) under the auspices of the European Telecommunications Standards Institute (ETSI). Following the proposal for their solution in October 2012, numerous telecommunication equipment providers and IT specialists joined the group. NFV encompasses the transition of network functions from dedicated hardware equipment to software-based applications. NFV represents the technology capable of capitalising on the benefits of the IT virtualisation evolution. Equipment and software components are consolidated on standardised IT platforms (e.g., high-volume servers,

switches, and storage), while network functions within the proprietary hardware can be concurrently decoupled. Through NFV, Virtual Network Functions (VNFs) can be instantiated at various locations, such as Data Centres (DCs), network nodes, and end users' premises, contingent upon network requirements. Using the benefits of cloud computing services, Software Defined Networking (SDN), and NFV, the opportunity to design and implement scalable, elastic, and programmable next generation networks is facilitated (Hawilo et al. (2019)).

B.2.1 NFV Architecture

The European Telecommunications Standards Institute (ETSI) delineates a comprehensive and sophisticated Network Functions Virtualization (NFV) architectural framework through its Industry Specification Group (ISG) for NFV. This framework endeavours to revolutionize the modus operandi of network operators in terms of designing, deploying, and managing network services by decoupling network functions from dedicated hardware appliances and facilitating their implementation as software entities on general-purpose, virtualized hardware platforms (Han et al. (2015) ETSI (2013)).




FIGURE B.2: ETSI NFV Architectural Framework

The ETSI NFV architectural framework comprises the following intricate components: (1) NFV Infrastructure (NFVI), (2) Virtual Network Functions (VNFs), (3) Management and Orchestration (MANO), and (4) NFV Management and Network Orchestration

(MNNO). The primary objective of the ETSI NFV architectural framework is to enable network operators to capitalize on enhanced agility, flexibility, and scalability while simultaneously reducing operational expenditures and accelerating the time-to-market for novel services (Han et al. (2015) ETSI (2013)).

B.3 SDN - Software Defined Networks

Over the past several years, there has been a notable advancement in developing and deploying networking devices (Göransson and Black (2016)). Due to the relatively constrained size of fixed networks and shared domains, each device is designed to maximize autonomy. This has led to the inception of intelligent, vendor-agnostic devices embedded within every apparatus (Göransson and Black (2016)). Nevertheless, the conventional network architecture must improve its capacity to adapt to evolving user requirements, particularly those brought forth by virtualization technologies, resulting in significant impediments (Feng et al. (2020)). Figure B.3 depicts the progress has facilitated the incorporation of migration functionality directly into hardware.




FIGURE B.3: Migration of network functionality to hardware

A pivotal aspect of SDN lies in its ability to shift control from network equipment manufacturers to network proprietors. In the mid-2000s, routers predominantly employed monolithic firmware with intricate, seldom utilized protocols and features (Göransson and Black (2016)). These systems were costly, inflexible, and challenging to modify or innovate. Consequently, enhancements and innovations were constrained and filtered through device manufacturers. A novel strategy would entail reevaluating networking methodologies from the ground up. Understandably, equipment manufacturers regard this shift as a menace, providing little incentive to support such initiatives. Casado et al. (2007) put forth a publication advocating for a fresh perspective on networking that enables network owners and operators to govern and configure the operations of both the

data and forwarding planes. Fundamentally, this approach grants network owners and operators increased autonomy, liberating them from reliance on network device manufacturers for essential features and functionalities within their respective environments.

Although SDN is perceived as an emerging networking paradigm, it is essential to acknowledge that the notion of programmable networks, which eventually developed into SDN, has existed for a considerable period (Nunes et al. (2014)). Antecedents to SDN can be traced back to Open Signalling (OpenSig), which aims to enhance the programmability of the internet and various mobile networks (Campbell et al. (1998)). This initiative advanced into the General Switch Management Protocol (GSMP) by employing a controller to facilitate resource allocation for multicast connections. In a similar vein, active networking Wetherall (2002) introduced the concept of a programmable network infrastructure designed to accommodate customised services. Nevertheless, this proposition garnered limited interest due to security and performance concerns (Smekal et al. (2017)).

The researchers advocated for a clean slate design to facilitate the separation of routing decision logic and the governing protocols that interconnect network devices. An alternate notion of segregation of control and management functionalities has been introduced in the 4D project. According to this methodology, NETCONF proposes a management protocol to modify network devices (Chaudhary et al. (2022)). This suggested management protocol aims to address and overcome the security limitations inherent in SNMP. SANE and ETHANE delineate a novel network architecture for enterprise networks, thus laying the foundation for the advent of SDN. Subsequently, ForCES and OpenFlow established standardised protocols for communication between the various SDN planes (Farahi (2025)). The implementation of a globally deployed software-defined WAN has been observed within data centre networks. Figure B.4 shows a chronological representation of the maturation of the SDN technology (De Freitas et al. (2020)). The establishment of a collaborative organisation known as the Open Networking Foundation (ONF) has been instrumental in advancing networking via SDN and standardising the OpenFlow protocol and associated technologies. This development has contributed significantly to the widespread acceptance and implementation of SDN technology in networking (Kafetzis et al. (2022)).

The initial stage involves the segregation of the control plane from the data plane, allowing centralised decision-making processes and improved programmability (Göransson and Black (2016)). In a monolithic architecture, each router encompasses both the control and data planes. The control plane determines the route, while the data plane implements these decisions. Software-defined networking (SDN) facilitates the separation of the control-data plane, allowing the control plane to operate externally to the data plane. This approach to decoupling provides an avenue for centralised decision-making, in which a centralised controller can implement routing decisions on remote data planes (Chaudhary et al. (2022)). In addition to centralisation, programmability is




FIGURE B.4: SDN Development Timeline (2008–2022)

another practical aspect. The programmable nature of the SDN controller allows it to be configured to respond in specific ways when certain conditions arise (Kalkan et al. (2018)). For example, when a link becomes congested, a backup link can be activated by establishing a new route that favours the backup link.

Subsequently, the development of a programmable data plane promotes advanced and innovative control measures. Alongside the programmable control plane, another groundbreaking development occurs within the data plane. The data plane is tasked with processing packets, including parsing and deparsing packet headers, altering packet header values, and managing packet queues. Historically, these functionalities were ingrained in the hardware or chip employed by routers or switches, rendering them inaccessible to network owners and operators (Bosshart et al. (2020)). Consequently, they had to request patches or features from the device manufacturers for any modifications that required firmware changes. The advent of a programmable data plane that uses the P4 programming language now allows end users to dictate how packets are processed, such as defining instructions for parsing or modifying packet headers (Bosshart et al. (2020)). This degree of data plane programmability empowers network owners/operators to develop and operate Virtual Network Functions (VNFs) within the data plane. Essentially, the switch's data plane serves as a runtime environment for the VNF. In contrast to previous VNF implementations (which operated on servers using hypervisor or container technologies), this study proposes the execution of VNFs on the data plane (VoD). This ability to run VNFs atop the data plane enables network operators to deploy VNF functions at the network's edge (Kaur et al. (2021)).

The emergence of programmable data planes in 2016 has attracted increasing interest and support from the industry and academic spheres. The groundbreaking development of building and operating virtual network functions (VNFs) in programmable data planes originated in 2016, coinciding with the release of Barefoot Tofino chipsets and the P4 programming language to programme these chipsets (Wang et al. (2020)). These chipsets enable original device manufacturers (ODMs) to create Ethernet switches that support programmable data planes with capacities of up to 6.5 Tbits/s. Currently, Netronome unveiled the Agilio SmartNIC (Zarca et al. (2020)), which allows P4 programmes to function on programmable data planes within Ethernet Network Interface Cards (NIC). In 2017, ODM Netberg launched the Aurora switch (Bosshart et al. (2020)), which uses a Tofino chipset. In 2018, another ODM, Edge-core, introduced the Wedge switch (Wang et al. (2020)), which uses the Tofino chipset. In 2019, Cisco unveiled the Q100 chip, which supports programmable data planes (Chaudhary et al. (2022)). In 2020, Pensando released a NIC card called the Distributed Services Card (DSC) (Chaudhary et al. (2022)), which is compatible with P4 and programmable data planes. This degree of industry acceptance offers end users competitive, cost-effective alternatives, fostering widespread adoption of programmable data planes (Chaudhary et al. (2022)). Consequently, these programmable data planes make efforts to secure the network edge more viable.

B.3.1 SDN Architecture

Software-Defined Networking (SDN) represents a paradigm shift in network architecture, characterised by decoupling the control plane, responsible for making decisions about traffic routing, from the data plane, which forwards traffic based on the control plane's decisions. The abstraction of the control plane from the underlying hardware facilitates centralised and programmable network management, engendering improved flexibility, scalability, and performance (Göransson and Black (2016)).

The SDN architecture is made up of three primary layers (refer to Figure B.5):

1. **Application Layer:** This layer constitutes a network application and service that take advantage of the SDN infrastructure. Applications may include network monitoring tools, traffic engineering systems, security solutions, and other services that require interaction with the network. These applications communicate with the Control Layer via well-defined Application Programming Interfaces (APIs), such as the Northbound API, to request network resources, receive network status updates, or enforce policies (Rawat and Reddy (2017) Göransson and Black (2016)).
2. **Control Layer:** Situated at the nucleus of the SDN architecture, the Control Layer assumes responsibility for the comprehensive management and orchestration

of the network. This layer encompasses the SDN Controller, a central decision-making entity that manages and controls the underlying network infrastructure. The SDN Controller translates the requirements from the Application Layer into specific instructions for the Infrastructure Layer. Communicates with network devices via the Southbound API, utilising standardised protocols such as OpenFlow, NETCONF, or gRPC (Sangodoyin et al. (2018b)).

3. **Infrastructure Layer:** This layer embodies the physical and virtual network devices, including switches, routers, and other networking equipment, which execute the actual data transmission based on the instructions received from the Control Layer. In an SDN environment, these devices are simplified and are primarily responsible for efficient packet forwarding. The control intelligence is centralised in the SDN controller, enabling the Infrastructure Layer to focus on high-performance data plane operations (Göransson and Black (2016)).




FIGURE B.5: SDN Architectural Layers

The SDN architecture heralds a revolution in traditional networking by segregating the control plane from the data plane, culminating in a centralised and programmable network management system. The three primary layers in the SDN architecture are

the Application Layer, Control Layer, and Infrastructure Layer, each with its distinct responsibilities and functions. This innovative approach generates more agile, flexible, and scalable networks, capable of rapidly adapting to evolving business requirements and technological advancements (Göransson and Black (2016)).

B.4 IoT Cloud Concept

Computing structures progressively move from physical forms to virtual patterns. In its physical form, the world of information technology allows organisations to develop complete applications for their use. Furthermore, this world allows exploiting platforms an organisation develops on another organisation's IT infrastructure. With the integration of service-orientated systems, the IT world moved into another stage. The new stage delegates a part of the software application's functionality to an existing software service on a different platform.

The cloud concept is constantly expanding, providing the capability of flexible and inexpensive systems that help organisations enhance their innovations for end users. The cloud computing concept pushes IT virtuality further, such as accessing third-party IT and software components. Furthermore, cloud computing opens the door for small organisations to offer and host scalable services. Moreover, this field is open to new challenges and hot research areas; one of these areas is cloud-based environments of the IoT. While the Internet of Things enables users to connect with millions of different intelligent platforms and systems to exchange data, cloud technologies have become one of the essential components that enhance the success of the concept of the Internet of Things. Together with cloud technologies, the Internet of Things has become a catalyst for the innovation of new and enhanced applications. These are accurate technologies of the future that will bring many benefits to the world of IoT (Almurisi and Tadisetty (2022)).

The rapid progression of IoT applications and technologies has developed new issues due to the difficulty of storing, processing, and accessing massive volumes of data. The collaborative usage of the Internet of Things and cloud technologies has resulted in significant innovation. It will be feasible to apply advanced processing of raw data streams in conjunction with new monitoring services. Sensor data, for example, can be uploaded and preserved in the cloud for later use as intelligent monitoring and activation through other devices or applications. On the other hand, multi-cloud technologies will allow more than one application to use the same sensor for different goals and results. The objective is to turn data into insight that drives cost-effective and productive action. The following points illustrate the main functions and advantages of cloud-based IoT environments (Aazam et al. (2022)):

1- IoT Cloud computing allows extensive network access and offers several connectivity choices. People access cloud computing resources through different platforms, such as mobile phones, tablets, and laptop computers. Although this is convenient for end consumers, it raises the issue of the requirements for network access points and monitoring.

2- On-demand IoT cloud computing is creating a new availability for developers. In other words, it is a web service that allows the developer to access without express permission or assistance. The only prerequisite is the connection to the Internet.

3- Fast and flexible; Consumers can expand storage space, review software settings, and function with the number of users. Users will have the ability to scale the service according to business needs. This characteristic makes it feasible to provide deep computing power and storage.

4- As the number of IoT applications, devices and automation in use rises, so do security issues. Cloud solutions provide businesses with reliable authentication and encryption procedures, but this area still needs more research.

IoT and cloud computing have been distributed globally and their features become extraordinary when combined. The services related to IoT and cloud computing are implemented in several different sensitive areas of our daily life, such as smart cities, electronic government, security, forensics, and health care organisations. The researchers proposed different integrated applications for IoT and cloud computing to develop and collect data through contributions of computational power and cloud storage. The proposed architectures and applications improved the objective of IoT services (Ishaq et al. (2021)). The figure B.6 below illustrates the cloud computing architecture of the IoT.

Cloud computing makes delivering data to data centres over the Internet more efficient and effective. This concept may provide a different classification that classifies cloud services into the following categories:

1. Platform as a Service (PaaS)
2. Software as a Service (SaaS)
3. Infrastructure as a Service (IaaS)
4. Public cloud
5. Private cloud
6. Hybrid cloud

Integrating the above categories and the IoT provides abundant connectivity opportunities, indicating extensive network access for the IoT environment. The IoT's environment involves numerous applications interconnected or attached to different devices and




FIGURE B.6: IoT Cloud Computing Architecture

platforms. The challenges differ for each application and platform; thus far, they all fall into the same classification since the repercussions are nearly comparable. The mobility and flexibility of the IoT makes security the utmost issue. Table B.1 shows the current challenges and issues in the IoT Cloud in its simple architecture.

One of the main capabilities of cloud networks is network virtualisation, which integrates software and hardware resources and network capabilities into a logically designed single software-based administrative entity. Researchers and scientists proposed several approaches to enhance the cloud functionality of the IoT and integrate virtualisation into the IoT network. The following subsection will address the different types of virtualisation.

B.4.1 Virtualisation of IoT Environments

Current studies classify the virtualisation of IoT environments into two main categories: A- network virtualisation and B- cloud virtualisation. The following subsections illustrate each type of the categories.

TABLE B.1: Current Challenges and Issues In The IoT Cloud.

Challenge/Issue	Explanation
Reliability	IoT devices rely on the cloud for data execution, calculations, and decisions. Reliability affects outcomes in time-sensitive applications like autopilot cars, healthcare, and security automation (Gill et al. (2019); Tayyaba et al. (2020)).
Security	Data in the IoT cloud environment must be processed in the cloud, with transitions between cloud and devices involving encryption, decryption, authentication, and authorization. Data's physical location is unknown to users, making security a significant concern (Tawalbeh et al. (2020); M Zeebaree et al. (2019)).
Data Storage and Access	By 2025, nearly 50 billion heterogeneous devices will join IoT. Rapid growth challenges cloud services to provide reliable, fast, safe, and stable storage and access (M Zeebaree et al. (2019); Xu et al. (2019)).
Computational Performance and Storage	Cloud-based IoT systems perform multiple functions with high mobility, making it challenging for cloud services to deliver consistent storage and operations (Tayyaba et al. (2020); Xu et al. (2019); Gupta (2021)).
Edge Computing	Edge computing helps customers respond to geo-distributed IoT installations and latency-limited cloud requests. However, position awareness requirements make implementation difficult (Riad et al. (2020); Anuradha et al. (2021)).
Maintenance	To ensure the security and efficacy of the IoT cloud environment, efficient procedures for tracking and controlling billions of IoT devices are necessary. Maintenance and self-healing solutions are crucial (Ali et al. (2021)).
Device Interaction	Systems combining IoT and cloud computing allow for device interaction. Device characteristics like cloud computing capacity, storage space, and intelligence restriction are essential (Wang and Zhang (2020); Mavromatis et al. (2020)).
User Aided Devices	Users must provide device attributions to interchange information with the system in IoT and cloud computing settings. Privacy concerns arise due to various societal difficulties (Chen et al. (2019)).

B.4.1.1 IoT Network Virtualisation

The process of virtualisation of the IoT arises by integrating the Internet of Things platform with a virtual network. A virtual network originates from the combination of software resources, network capabilities, and physical hardware. This virtualisation's desired goal is achievable by creating a virtualisation abstraction layer between the network and storage hardware (Alam et al. (2019)). The researchers classified network virtualisation into two categories as follows:

- External Virtualisation: Combining several networks from different areas into a virtual division (Hwang and Shin (2018)).
- Internal Virtualisation: Provide network functionalities in the shape of logic as software containers on a single network server (Hwang and Shin (2018)).

The virtualised network consists of several elements, such as Control Plane Virtualisation, Function Virtualisation, and Device Virtualisation. Each component of the above uses a different technique to implement the virtualisation.

- *Control Plane Virtualisation*

Physical standard networks contain several different types of hardware for connectivity purposes. Each component has a built-in controller that contributes to the routing process and instructs other parts where to forward packets. The rapid growth of the online environment demands controllers with more flexible features. According to the nature of cloud computing environments, the controllers need to be able to manage them from anywhere and anytime. (SDNs) Software Defined Networks can implement the virtualisation of the network controller(Banchuen et al. (2018)).

The SDN controller is programmable, which supports the main idea of virtualisation by permitting the underlying infrastructure to be recapitulated for applications and network services. Programmability directs to improving network characteristics by connecting them to applications and allowing for dynamic traffic flow changes, hence offering network and application-level Quality of Service (QoS) (Jiang et al. (2021)). Software Defined Networks are manageable dynamically. SDNs can route packets in a particular network; they can prioritise, deprioritise, or block specific types of packages with precise control. This approach is also known as efficient traffic engineering since it allows administrators to employ less costly OpenFlow-compliant entity switches. SDN architecture consists of the following layers:

A- Data Plane: it is also called the forwarding plane. This layer contains switches of physical and software nature. These components are connected through a physical medium and perform dynamically configured instructions via the control plane (Jiang et al. (2021); Kafetzis et al. (2022)).

B- Control Plane: This layer is also known as a controller, which is separate from the data plane layer and is considered a centralised logic server. According to recent studies, these controllers can be distributed or centralised. Each controller has a global view of the entire network that lies under its ability. Based on this view, it configures and installs trafficking rules to the network devices (Jiang et al. (2021); Ravuri et al. (2020); Kafetzis et al. (2022)).

C- Application Programming Interfaces: This layer consists of four different interfaces; 1- Southbound Interface, 2- Management Plane, 3- Northbound Interface, and 4- East/Westbound Interface. These different interfaces provide communications between the control, data, and management planes. APIs allow multiple control planes to work together in the distributed architecture, which provides a

solution for significant challenges in the SDNs. Furthermore, it manages and monitors switches in the data plane via the control plane (Jiang et al. (2021); Ravuri et al. (2020); Alam et al. (2019); Kafetzis et al. (2022)).

- *Function Virtualization*

The researchers implemented Network Function Virtualisation (NFV) to virtualise network functions. NFV, in its generic architecture, consists of the following elements: A- Virtual Network Function, B- Network Function Virtualisation Infrastructure, and C- Network Function Virtualisation Management and Orchestration. Network function virtualisation infrastructure configures the network functions as pieces of software. These pieces work as virtual network functions, which lay on top of hardware infrastructure to run and handle network operations (Kim and Kim (2021); Alam et al. (2019)). Figure B.7 below shows different aspects of the NFV structure .




FIGURE B.7: NFV Structure

- *Device Virtualisation*

IoT virtualisation uses a hypervisor functionality that allows its resources to be efficiently utilised by different applications. This element is a process of virtualising switches in the data plane layer. This process can be implemented through logical conception of its components or merely the ability of the functionality to be run on several operating systems (Zhang et al. (2021a)).

B.4.1.2 Challenges in Virtualised IoT Environments

Physical Sensors Management Disseminating physical sensors with other applications is difficult, since they have limited capabilities and are only maintained and controlled by the applications of their providers. Cloud computing is required to support sensor management and improve the functionality of physical sensors to solve this problem [Almurisi and Tadisetty \(2022\)](#); [Riofrio and Caiza \(2021\)](#).

Resource Exhaustion Continuously serving different applications may cause cloud computing resources to deplete and decrease Quality of Service (QoS). A suitable resource allocation technique can solve this problem to keep the load balance between resources [Almurisi and Tadisetty \(2022\)](#); [Riofrio and Caiza \(2021\)](#).

Task Scheduling Since each resource has a range of capabilities (memory, bandwidth, CPU speed, etc.), choosing the appropriate equipment to perform the work is challenging. It is crucial in IoT cloud environments to determine which device is most suited to carry out the task to prevent processing delays in the Sensor-Cloud paradigm [Almurisi and Tadisetty \(2022\)](#); [Aazam et al. \(2022\)](#).

Heterogeneity Support for various types of routing protocols and operating systems, managing such a high level of heterogeneity and facilitating communication between them is a significant challenge since each sensor node differs from the others in network configuration [Almurisi and Tadisetty \(2022\)](#).

Dynamic Networks Formation Another critical challenge is to construct virtual networks at the cloud level from the underlying substrate networks through virtual links. The fundamental concept is to automate the deployment of Virtual Sensor Network (VSN) services to numerous applications while allowing dynamic grouping of virtual sensors to construct a VSN from one or more entire networks [Almurisi and Tadisetty \(2022\)](#); [Aazam et al. \(2022\)](#).

Energy Depletion The physical layer of the cloud environment of the IoT is based on battery-operated devices; it is imperative to create practical algorithms to make the most of their power unless there is an innovative energy strategy to limit power waste during data transmission [Almurisi and Tadisetty \(2022\)](#).

Security and Privacy Security is an important concern with this technology development. A security mechanism must always be available to protect against any compromise that can affect how data is transferred between sensor networks and the cloud. In the IoT Cloud environment, maintaining safer network functionality and guaranteeing the confidentiality of sensor network providers is a major challenge [Almurisi and Tadisetty \(2022\)](#); [Aazam et al. \(2022\)](#); [Riofrio and Caiza \(2021\)](#).

B.5 Security Issues and Vulnerabilities

In this section, we will look at the security issues and vulnerabilities associated with Software-Defined Networking (SDN) and Network Function Virtualisation (NFV), two groundbreaking paradigms that have transformed traditional networking architectures. Although these innovations offer numerous advantages, such as increased flexibility and reduced operational costs, they also introduce new threats and vulnerabilities that must be carefully addressed to ensure secure and reliable network operation. We will explore the specific threats and vulnerabilities in SDN, including SDN controller vulnerabilities, communication vulnerabilities, and network device vulnerabilities. In parallel, we will examine the threats and vulnerabilities in NFV, encompassing resource isolation, management plane security, and NFV security. Furthermore, we will provide an overview of recent studies on the detection and mitigation of distributed denial of service (DDoS) attacks, as these attacks pose significant risks to SDN and NFV infrastructures. Finally, we will present an overview and classification of DDoS attacks to better understand the nature of these threats and their potential impact on the network.

B.5.1 Security Unique Provisioned Challenges in Virtualised IoT

Building an IoT system that is safe and protected might be difficult due to network security challenges. To stop external attacks, network edges have historically deployed security tools such as firewalls and intrusion detection and prevention systems. Such procedures are no longer sufficient in light of the constant changes in network topology caused by IoT nodes joining and leaving the network. For internal risks, if a virus corrupts one item, other uncorrupted objects may also be at risk. Therefore, as technology develops, it may be necessary to reevaluate security measures for both internal and external threats. The first two levels of IoT virtualisation environments will have security vulnerabilities that are not present in the typical IoT environment.

Cloud IoT virtualisation allows the existing infrastructure to serve many additional customers. Consequently, the IoT cloud environment consists of traditional IoT equipment alongside new technologies, which opens us to new security challenges alongside existing ones. This section explores the security problems in each layer of the IoT architecture as follows:

- Perception layer: Sensors are the source of the majority of security concerns in this layer; as a result, attackers often focus on this layer. Since this layer encloses sensors such as the global positioning system (GPS), wireless sensor networks (WSNs), and radio frequency identification (RFID), which are physical nodes (sensors) with limited power and storage, it is difficult to use the conventional security method on this layer, since it only has virtual nodes (sensors) with limited power and storage [Chaabouni et al. \(2019a\)](#). The following are also the main security concerns with this layer:

i) False Sensor Data.

ii) Malicious Sensor Commands.

iii) Eavesdropping.

- **Transportation layer:** The environment of this layer contains significant security issues, particularly concerning authentication and integrity, making the transportation layer vulnerable to attacks, especially when cloud technologies are integrated into this layer. These typical security concerns are presented in this section [Chaabouni et al. \(2019a\)](#); [Zhang et al. \(2021c\)](#):

i) Denial of Service (DoS) attack.

ii) Man-in-the-Middle (MITM) attack.

iii) Storage attack.

- **Application Layer Security Issues:** The use of SDN in the IoT for use cases particular to applications is crucial. The use of SDN in the IoT further raises the stakes for security issues. ([Bastos et al. \(2018\)](#)) claims that a large number of IoT-based household equipment, including smart lights, motion sensors, smoke alarms, and monitoring / analysis devices, lack crucial security features that could have a detrimental effect on regular activities. The authors contend that numerous elements, such as device capabilities, operational mode, and manufacturer, must be taken into account when implementing security. The suggested solutions may be able to stop attacks at the network level, but a detailed discussion of user privacy protection is still required given the possibility of sensitive personal data being exposed ([M Zeebaree et al. \(2019\)](#)).

- **Architectural Security Challenges:** ([Zuo et al. \(2020\)](#)) suggests a method primarily intended to strengthen the security of SDN controllers and address scalability problems in several IoT-based domains. The research integrates wired and wireless networks and extends its solution to ad hoc enabled networks and Internet of Things (IoT) devices such as sensors, smartphones, tablets, etc. The development of numerous controller management techniques (such as border and security controller management) and inter-SDN controller communication at various levels may be part of future work. The framework is not comprehensive because the system can only manage connections between clusters. As a result, developing the simulation to take advantage of the OpFlex protocol functionalities and building up a dynamic routing protocol are tasks for the future. In summary, researchers want to analyse the findings of various security attacks, including black holes, adjacent attacks, and other similar attacks ([Alam et al. \(2019\)](#); [Samaniego et al. \(2018\)](#); [Ugwuanyi et al. \(2020\)](#)).

- **Miscellaneous Security Challenges:** ([Mohan et al. \(2022\)](#)) provide a simple countermeasure tool to identify man-in-the-middle (MitM) attacks in IoT cloud settings. One

of the common Transport Layer Security (TLS) flaws is the MitM attack, and the SDN controller is vulnerable to it. Although the frequency of false positives has decreased, trials have shown a considerable increase in the performance and detection accuracy of this method. The network might potentially suffer harm from passive assaults. Compared to current protocols, the proposed alternatives offer a higher level of security. However, due to the increased connection between the proposed systems between the SDN controller and the IoT nodes, traffic control can become challenging (Alam et al. (2019)).

B.5.2 Threats and Vulnerabilities in SDN

Software-Defined Networking (SDN) is a novel networking paradigm that has transformed traditional networking architectures by decoupling the control plane from the data plane. This approach facilitates centralised and programmable network management, resulting in enhanced flexibility, scalability, and performance. Despite the numerous benefits of SDN, its architecture introduces new threats and vulnerabilities that must be carefully addressed to ensure secure and reliable network operation. This research aims to provide a comprehensive overview of the various threats and vulnerabilities associated with SDN and to discuss potential mitigation strategies (De Freitas et al. (2020)).

B.5.2.1 SDN Controller Vulnerabilities

The SDN controller, as the central decision-making entity in SDN architecture, is an attractive target for attackers. Compromising the controller can grant an attacker full control over the network, which can have severe consequences (Hamza et al. (2019)). Some of the key vulnerabilities associated with the SDN controller are:

1. **Single Point of Failure:** The centralized nature of the SDN controller can make it a single point of failure. If the controller is compromised or experiences a technical issue, the entire network can be affected (Vishwakarma and Jain (2020)).
2. **Unauthorized Access:** Insecure communication channels between the controller and network devices or applications can be exploited by attackers to gain unauthorized access to the controller and manipulate the network (Vishwakarma and Jain (2020)).
3. **Denial of Service (DoS) Attacks:** The controller may be susceptible to DoS attacks, which can disrupt its operation and subsequently affect the performance of the entire network (Vishwakarma and Jain (2020)).

B.5.2.2 Communication Vulnerabilities

In SDN, communication between the controller and network devices (southbound) and between the controller and applications (northbound) is essential for network management. However, these communication channels can be exploited by attackers.

1. **Eavesdropping:** Attackers may intercept communication between the controller and network devices or applications to gain access to sensitive information, such as network configurations and security policies (Göransson and Black (2016)).
2. **Man-in-the-Middle (MITM) Attacks:** Attackers can position themselves between the controller and network devices or applications to intercept and manipulate traffic, potentially causing network misconfigurations or injecting malicious traffic (Vishwakarma and Jain (2020)).
3. **Spoofing Attacks:** Attackers can impersonate network devices, applications, or even the controller to send fake instructions or manipulate network behaviour.

B.5.2.3 Network Device Vulnerabilities

While SDN aims to simplify network devices and focus on efficient packet forwarding, these devices can still be targets for attacks.

1. **Device Compromise:** Attackers can exploit vulnerabilities in network devices to compromise their functionality and disrupt network operations (Sangodoyin et al. (2018a)).
2. **Unauthorized Access:** Weak authentication mechanisms can allow attackers to gain unauthorised access to network devices, allowing them to modify configurations or inject malicious traffic.
3. **Flow Table Poisoning:** Attackers can exploit vulnerabilities in the control plane to inject malicious flow rules into network devices, causing traffic to be misrouted, dropped, or redirected to unauthorised destinations (Sangodoyin et al. (2018a)).

B.5.3 Threats and Vulnerabilities in NFV

Network Functions Virtualisation (NFV) is a cutting-edge networking paradigm that aims to virtualise network functions traditionally implemented in dedicated hardware devices, such as routers, firewalls, and load balancers. NFV enables the deployment of these functions as software instances running on standard servers, resulting in improved flexibility, scalability, and reduced operational costs. Despite the numerous benefits of

NFV, the virtualisation and orchestration of network functions introduce new threats and vulnerabilities that must be carefully addressed to ensure secure and reliable network operation. This research aims to provide a comprehensive overview of the various threats and vulnerabilities associated with NFV and to discuss possible mitigation strategies ([Networks \(2022\)](#)).

B.5.3.1 Resource Isolation

In NFV environments, multiple virtual network functions (VNFs) may share the same underlying physical resources, such as compute, storage, and networking. Ensuring proper resource isolation is critical to prevent security breaches and performance degradation.

1. **Cross-tenant Attacks:** Insufficient isolation between VNFs belonging to different tenants can lead to unauthorised access to data or resources, resulting in data leakage or service interruption.
2. **Side-channel Attacks:** Attackers may exploit shared resources, such as CPU caches or memory, to extract sensitive information from other VNFs running on the same physical host ([Hao et al. \(2018\)](#)).

B.5.3.2 Management Plane Security

The management plane is responsible for orchestrating and managing the VNFs and the underlying infrastructure. Security vulnerabilities in the management plane can have severe consequences:

1. **Misconfigurations:** Human errors or misconfigurations in the management plane can lead to security vulnerabilities or service interruptions ([Jiang et al. \(2021\)](#)).
2. **Insider Threats:** Malicious insiders with access to the management plane can compromise the security and integrity of VNFs and the underlying infrastructure ([Hao et al. \(2018\)](#)).
3. **Insecure APIs:** Inadequate security mechanisms in the APIs used for communication between the management plane and the VNFs can expose the system to unauthorised access or data breaches.

B.5.3.3 VNF Security

As VNFs replace dedicated hardware devices, they inherit their functionalities and associated security risks, in addition to new vulnerabilities introduced by the virtualisation process:

1. **Software Vulnerabilities:** VNFs, as software-based implementations of network functions, can be susceptible to software vulnerabilities, such as buffer overflows or memory leaks, which can be exploited by attackers (Kim and Kim (2021)).
2. **Lifecycle Management:** The dynamic nature of VNF instantiation and termination introduces challenges in securely managing VNF lifecycle events, such as ensuring data confidentiality and integrity during migration or decommissioning (ETSI (2013)).

B.6 Recent Studies on DDoS Attacks Detection and Mitigation

Distributed Denial of Service (DDoS) attacks pose a considerable threat to network availability and are widely considered a significant security concern. Various methodologies have been developed by researchers in order to safeguard networks from these attacks. Addressing a DDoS attack typically involves two phases: the identification of the attack (*detection*) and the implementation of countermeasures (*mitigation*). Detection and mitigation strategies may incorporate statistical analysis, predefined rule sets (based on policies), machine learning, data mining techniques, or a combination thereof. The implementation of defensive and mitigation measures can be carried out in a centralised or distributed fashion, depending on the configuration of the modules deployed. The literature has delineated and classified an extensive array of DDoS detection and mitigation techniques (Valdovinos et al. (2021)).

As outlined in Mousavi and St-Hilaire (2015), the authors put forth an entropy-based methodology for the identification of DDoS attacks. This innovative approach offers a lightweight and effective solution to detect DDoS assaults in *Software-Defined Networking (SDN)* environments by capitalising on the pivotal role the controller occupies within the SDN framework. Furthermore, the proposed method leverages the destination IP address to discern attacks within the first 250 packets of traffic containing malicious payloads. In this case, the threshold is established at the minimum possible rate of attack traffic. Despite this, the detection rate associated with the selected threshold remains impressive 96%. In addition to its capacity for early detection, the integration

of this solution within the controller remains unobtrusive with respect to resource utilisation and functionality. Consequently, it furnishes detection capabilities for both the controller and the targeted host.

An alternative entropy-based detection paradigm is introduced in [Fioravanti et al. \(2023\)](#), employing the seminal Shannon Entropy formulation, a cornerstone of Information Theory. The foundation of this approach lies in a collection of parameters that require meticulous calibration: the number of consecutive windows N , where the entropy is scrutinised to ascertain whether its value satisfies a predetermined threshold (an additional parameter); and, ultimately, the dimensions of a window W . Precise fine-tuning of these parameters is of paramount importance in order to enhance the accuracy and efficacy of the DDoS detection mechanism.

In [Kalkan et al. \(2018\)](#), a joint entropy-based technique grounded in a statistical model is presented for the detection and mitigation of DDoS attacks in the context of an SDN environment. The authors introduce a statistical defence scheme, referred to as *Joint Entropy Statistical Scheme (JESS)*, specifically designed to combat DDoS attacks in SDN settings. The first model employing joint entropy for defensive purposes, JESS encompasses three distinct phases: nominal, preparatory, and active mitigation. During the nominal stage, baseline profiles are established in the absence of any attacks. The preparatory stage facilitates rapid detection of an ongoing attack, followed by the active mitigation stage, which persists until the assault is successfully thwarted or stopped.

Due to its reliance on a statistical methodology, JESS demonstrates effectiveness against both known and novel attacks that exploit attack packets bearing similar attributes. The findings indicate that JESS yields exemplary results for a multitude of recognised attack vectors. In addition, it exhibits a 80% success rate when faced with an unfamiliar attack type, designated as a generic attack, and achieves a 70% success rate in the context of a mixed attack. To elucidate the algorithmic properties, the authors conduct an analysis of JESS's temporal, storage, and communication overheads. The outcomes reveal that JESS operates efficiently, necessitating minimal storage and processing resources.

[Sangodoyin et al. \(2018a\)](#) introduces a statistical methodology for detecting DDoS attacks in the context of SDN, with throughput serving as an impact metric. In this proposed technique, network throughput is sampled within a specified interval to determine the normal distribution of network data in the absence of an attack. Subsequently, a confidence interval (CI) is derived for this normal distribution. A marked deviation in the mean throughput value, ascertained at subsequent intervals compared to the mean throughput without an attack, serves as an indicator of an ongoing assault. The calculations associated with the confidence interval and mean throughput are characterised by low overhead, making them amenable to facile implementation within the SDN controller for anomaly detection purposes. The evaluation demonstrates that by harnessing

throughput information from the network server, DDoS attacks can be effectively identified in real time with an accuracy of approximately 99%, provided the polling interval is set to 60 seconds.

In [Yang and Zhao \(2018\)](#), a machine learning-centric methodology is emphasised for the early detection of DDoS attacks. The proposed framework consists of three components: a traffic collection module, an attack identification module, and a flow table delivery module. The traffic collection module extracts the salient traffic characteristics in preparation for traffic identification. In the present work, researchers have employed Support Vector Machines (SVM) for the identification of DDoS traffic. Experimental results based on the KDD99 dataset demonstrate the efficacy of this approach.

This classification model is implemented within a simulated SDN environment as a DDoS detection module for a campus network. The model is responsible for the analysis of all incoming traffic. Upon detection of attack traffic, the controller proceeds to discard packets in accordance with predefined rules. In the absence of an attack, the standard forwarding policy is executed.

In [Cui et al. \(2016\)](#), the SD-Anti-DDoS mechanism is proposed and demonstrated within the context of SDN. This work introduces an attack detection trigger method, an innovative approach to schedule the initiation of attack detection. In addition to an attack detection method, a traceback technique is presented to trace the attack source leveraging the unique characteristics of SDN. A mitigation strategy is also implemented to obstruct attack traffic and liberate the space occupied by an attack. The SD-Anti-DDoS system, which includes the attack detection trigger mechanism, attack detection, and specialised traceback and mitigation in SDN, is implemented on the RYU controller. The system's availability is subsequently evaluated.

The findings indicate that the proposed mechanism can expedite the onset of attack detection, achieving a response time of less than one second, and accurately trace the source switch of the attack. Crucially, the mechanism can block the attack at its source and free up the occupied space within the switches. The proposed attack detection trigger mechanism is shown to markedly reduce the response time against attacks compared to the currently employed periodic trigger. Simultaneously, it can also decrease the CPU and network load on the controller. Although OpenFlow v1.0, one of the most widely used protocols in SDN, is used in this study, SD-Anti-DDoS can also be implemented using alternative versions of OpenFlow. It is worth noting that minor performance disparities may exist between different iterations of OpenFlow ([Vishwakarma and Jain \(2020\)](#)).

B.6.1 Overview and Classification of DDoS Attacks

Availability necessitates the uninterrupted operation of computer systems, ensuring that resources remain accessible to authorized users at any given time. Distributed Denial of Service (DDoS) attacks persist as one of the most formidable challenges to availability, representing a significant security concern in the contemporary internet landscape. DDoS attacks can be categorised under the availability component of the CIA Triad, as illustrated in Figure B.8; This type of attack continues to plague the network security domain, with the threat resolution remaining elusive, mainly due to the high-severity risks it poses to SDN planes. A succinct overview and taxonomy are provided in the subsequent subsection (Kumari and Jain (2023b)).




FIGURE B.8: The CIA Triad: Confidentiality, Integrity, and Availability.

B.6.1.1 Overview of DDoS Attacks

Distributed Denial of Service (DDoS) attacks have garnered considerable attention in recent years, as they pose a significant threat to the availability and security of internet-based services. DDoS attacks involve overwhelming a target system with an immense volume of traffic, rendering it unable to respond to legitimate requests and, consequently, causing a denial of service to valid users. These attacks not only impact the target system, but can also cause collateral damage to the surrounding internet infrastructure. This research provides a comprehensive academic overview of DDoS attacks, delving into their history, classification, execution methods, and potential countermeasures (Valdovinos et al. (2021)).

History of DDoS Attacks The history of DDoS attacks dates back to the late 1990s, when rudimentary attacks were orchestrated by exploiting operating system vulnerabilities or leveraging basic flooding techniques. As the internet evolved, so too did DDoS attacks, growing in sophistication and scale. In the early 2000s, attackers began employing botnets—networks of compromised machines—to coordinate large-scale DDoS attacks. Recent years have witnessed an exponential increase in the frequency, magnitude, and

complexity of DDoS attacks, with notable examples including the Dyn attack in 2016 and the GitHub attack in 2018 (Kumari and Jain (2023b)).

B.6.1.2 Classification of DDoS Attacks

Distributed Denial of Service (DDoS) attacks pose a significant threat to the availability and security of internet-based services. These attacks involve overwhelming a target system with a massive volume of traffic, causing it to become unresponsive to legitimate requests. In this research, we provide a comprehensive academic overview of DDoS attack classifications, delving into their various categories and characteristics. The classifications are illustrated in Figure B.9.




FIGURE B.9: DDoS Attack Types Classifications

Network device-level

Network device-level attacks focus on exploiting vulnerabilities in network devices such as routers, switches, and firewalls. By targeting these devices, attackers can disrupt or degrade the performance of the entire network. Common types of network device level attacks include:

- **Router and switch attacks:** These attacks aim to compromise routers and switches to disrupt network connectivity or reroute traffic through malicious nodes (Sudheera et al. (2021)).

- **Firewall attacks:** These attacks focus on bypassing or incapacitating firewalls to gain unauthorised access to the target network or disable security mechanisms (Feng et al. (2020)).

OS Level Attacks

OS-level attacks exploit vulnerabilities or design flaws in a target system's operating system. By exploiting these weaknesses, attackers can cause system crashes, resource exhaustion, or unauthorised access. Common types of OS level attacks include:

- **Buffer overflow attacks:** These attacks involve sending excessive amounts of data to a target system, causing memory buffers to overflow and potentially leading to system crashes or arbitrary code execution (Singh and Behal (2020)).
- **SYN flood attacks:** SYN flood attacks exploit the TCP connection process, inundating the target with SYN packets and causing the system to exhaust its resources while trying to handle the flood of connection requests (Sudheera et al. (2021)).

Application-Based Attacks

Application-based attacks focus on specific applications or services running on a target system. By overwhelming the application with seemingly legitimate requests or exploiting vulnerabilities in the application, attackers can exhaust the target's resources or disrupt its functionality. Common types of application-based attacks include

- **HTTP flood attacks:** These attacks involve sending a large number of HTTP requests to a target web server, overwhelming the server and causing it to become unresponsive to legitimate traffic (Sudheera et al. (2021)).
- **DNS query flood attacks:** These attacks target DNS servers with a flood of DNS queries, consuming resources and causing the server to become unresponsive to legitimate requests (Bastos et al. (2018)).

Protocol Feature Attacks

Protocol feature attacks exploit weaknesses or inherent characteristics of network protocols to disrupt a target system. These attacks often leverage the design of the protocols themselves, making them difficult to defend against without altering the protocols or employing specialised mitigation techniques. Common types of protocol feature attacks include:

- **IP fragmentation attacks:** These attacks exploit the IP fragmentation process by sending a large number of fragmented IP packets to the target system, causing it to consume resources while trying to reassemble the packets (Sudheera et al. (2021)).
- **ICMP flood attacks:** ICMP flood attacks involve sending a massive volume of ICMP Echo Request (ping) packets to the target system, overwhelming it with incoming traffic, and making it unresponsive to legitimate requests (Sudheera et al. (2021)).
- **TCP reset attacks:** These attacks involve sending TCP Reset (RST) packets to the target system or its connected clients, causing ongoing connections to be abruptly terminated and disrupting communication (Sudheera et al. (2021)).

Data Flood Attacks

Data flood attacks involve inundating a target system with an immense volume of data, with the aim of consuming its available bandwidth and prevent it from serving legitimate requests. These attacks can be conducted using various methods, such as generating large volumes of traffic from a botnet or employing reflection and amplification techniques. Common types of data flood attacks include:

- **UDP flood attacks:** These attacks involve sending a large number of User Datagram Protocol (UDP) packets to random ports on the target system, causing it to consume resources while processing the incoming traffic and responding to non-existent services (Valdovinos et al. (2021)).
- **DNS amplification attacks:** In DNS amplification attacks, the attacker sends spoofed DNS queries to open DNS resolvers, causing them to return large response messages to the target system. This amplifies the volume of traffic directed at the target, overwhelming it with incoming data (Valdovinos et al. (2021)).

B.6.1.3 Mechanisms and Impacts of DDoS Attacks on IoT Ecosystem

Distributed Denial of Service (DDoS) attacks represent a significant threat to modern Internet of Things (IoT) ecosystems. A detailed understanding of their operational mechanisms and the damage they cause is critical to formulating effective countermeasures to mitigate their adverse effects.

Impacts on IoT Ecosystems

DDoS attacks disrupt IoT ecosystems, affecting critical applications such as smart homes and industrial automation. These attacks can paralyse operations, leading to substantial

Attack Type	Description and Examples
Volumetric Attacks	Overwhelms the network bandwidth or processing capacity of the target (Furfaro et al. (2020)). Examples: • <i>UDP flood</i>: Sends a massive number of UDP packets to the target. • <i>ICMP flood</i>: Bombards the target system with numerous ICMP requests.
Protocol Exploitation Attacks	Exploits weaknesses in network protocols (Kumar et al. (2023)). Examples: • <i>SYN flood</i>: Abuses the TCP handshake process by sending many SYN requests without completing connections. • <i>Ping of Death</i>: Transmits oversized ICMP packets to crash the target system.
Application Layer Attacks	Targets specific applications or services (Mishra and Pandya (2021)). Examples: • <i>HTTP flood</i>: Overloads web servers with multiple GET or POST requests. • <i>DNS amplification attack</i>: Utilises DNS servers to amplify malicious traffic.
IoT Botnets	Leverages infected IoT devices to form botnets and launch attacks. Notable example: <i>Mirai botnet</i> (Borys et al. (2022)).

TABLE B.2: Types of DDoS Attacks in IoT and Examples

financial losses, reduced customer trust, and reputational damage (Kumari and Jain (2023a)).

During an attack, the integrity and availability of crucial data are compromised, causing delays in decision-making processes and daily operations. The interconnected nature of IoT means that a single node failure often leads to larger system failures, affecting the entire network (Fu et al. (2025)).

Even when attacks are repelled, managing the resulting large-scale malicious traffic consumes significant computing and network resources (Tushir et al. (2021)). Over time, persistent DDoS attacks can erode public trust in IoT systems, hindering industry growth. These multifaceted consequences highlight the importance of developing robust DDoS defence mechanisms tailored to IoT environments.

B.7 Proposed Technologies and Security Solutions For Virtualised IoT

The importance of IoT system security is increased by diversity and heterogeneity. IoT systems' security differs from that of traditional systems for a number of reasons. Each

security challenge require a solution. This section addresses a summary of different proposed security tactics for IoT environments.

1- STEC-IoT: A Security tactic by Virtualizing Edge Computing in IoT proposed by (Zhang et al. (2021a)). In this approach, the potential physical nodes are chosen taking into account several factors, and then the nodes are mapped according to their resource metrics. The particle swarm optimisation (PSO) technique is used in the link mapping process to optimise the mapping process, and the final mapping results are achieved.

2- TA-SVNE: The trust-aware secure virtual network embedding algorithm, proposed by (Gong et al. (2015)) The K-shortest path algorithm is used for link mapping after the Technique for Order of Preference by Similarity to Ideal Solution (TOPSIS) maps the nodes according to their relevance. The technique considers the local and global importance of the nodes throughout the mapping process.

3- G-SVNE: Survivable Virtual Network Embedding Algorithms for Network Virtualization, proposed by (Zhang et al. (2021b)). In this technique, the best solution for each step is obtained using a greedy strategy, and the best solutions for each step are then merged to provide the final mapping result.

4- MCS-VNE: multi-domain virtual network embedding algorithm based on minimum cost, proposed by (Zhang et al. (2022)). This algorithm is obtained by adding security attributes using the MC-VNE algorithm. According to the estimated mapping cost algorithm, each physical domain determines the estimated mapping cost of each node and selects the node with the lowest estimated mapping cost as the candidate node.

B.7.1 Existing DDoS Attack Detection Methods

In IoT environments, DDoS attack detection methodologies are typically divided into two main categories: signature-based detection and behaviour-based detection. Each of these approaches has its unique strengths and is often used in combination to improve both the accuracy and efficiency of detection systems (Dubey et al. (2023)).

B.7.1.1 Feature-based Detection Methods

Feature-based detection methods identify potential DDoS attacks by analysing multiple characteristics and attributes of network traffic. These methods are faster than many others, enabling real-time detection of attacks, but may struggle to identify novel types of threat. Unlike static volume and rate analyses, this approach examines various characteristics of packets, such as size, type, and content. Using these attributes, feature-based detection systems are better equipped to distinguish between normal network activity and malicious behaviour (Farahi (2025)).

Detection Method	Key Features
Statistical Feature Analysis	Uses Self-Organising Maps (SOM) to evaluate statistical traffic features (e.g., packet count, byte count) for detecting DDoS attacks (Shah et al. (2020)).
Entropy-based Detection	Identifies abnormal traffic patterns by assessing entropy changes in IP addresses and port numbers (Ahmad and Alsmadi (2021)).
Machine Learning Approach	Utilises Random Forest algorithms to construct detection models based on various traffic features, such as protocol type and packet size distribution (Ahmad and Alsmadi (2021)).
Deep Learning Method	Employs Artificial Neural Networks (ANNs) to learn traffic characteristics and differentiate between normal and malicious activities (Lakshmana et al. (2022)).
Feature Selection Optimisation	Enhances detection performance using techniques like Principal Component Analysis (PCA) to select the most relevant features (Sivanathan (2020)).

TABLE B.3: Feature-Based DDoS Attack Detection Methods in IoT Environments

B.7.1.2 Behavior-based Detection Methods

Behaviour-based detection methods identify threats by analysing and evaluating the behavioural patterns of network entities. Unlike static feature-based methods, behaviour-based approaches focus on dynamic user interaction patterns. By continuously monitoring these patterns, this method can identify previously unknown security threats indicated by anomalous behaviours. This distinguishes it from feature-based methods, which mainly rely on predefined characteristics (Ahmad and Alsmadi (2021)).

Detection Method	Key Features
Traffic Pattern Analysis	Identifies anomalies by examining time-series characteristics of network traffic (Sivanathan (2020)).
Protocol Behaviour Analysis	Detects SYN flood attacks by assessing patterns in TCP connection establishment and termination (Sivanathan (2020)).
Device Behaviour Profiling	Learns the normal communication behaviour of IoT devices to detect deviations indicative of malicious activity (Hamza et al. (2022)).
Graph Models	Represents network behaviours using graph models and identifies DDoS attacks by analysing changes in graph topology (Lin et al. (2019)).
Time Series Analysis	Uses Long Short-Term Memory (LSTM) networks to analyse the temporal characteristics of network traffic for anomaly detection (Parizotto et al. (2023)).

TABLE B.4: Behaviour-Based DDoS Detection Methods in IoT Environments

B.8 Applications of Machine Learning in Attack Detection

With the rise of evolving cybersecurity threats, traditional detection methods have been shown to be insufficient to address the new challenges. Consequently, in the realm of IoT network security, the integration of machine learning (ML) technology has emerged as a central research focus. Machine learning has demonstrated considerable effectiveness by autonomously identifying complex attack patterns and adapting to dynamic threat landscapes (Chaabouni et al. (2019b)). This adaptability makes machine learning an invaluable tool for improving security protocols. Using this technology, the challenges posed by contemporary network threats can be significantly mitigated (Khan and Bokhari (2023)). This section briefly discusses machine learning techniques and illustrates their classifications. On the other hand, it discusses machine learning applications in the detection of attacks in IoT environments.

B.8.1 Learning Techniques

Advancements in traditional IoT technologies are generally supported by the integration of cloud storage, data mining, and big data analytics (Chaabouni et al. (2019b)). Moreover, virtualised IoT environments rely exclusively on cloud-based infrastructures to operate and evolve (Alam et al. (2019)). Rapid growth in the Internet has increased the amount of data that is generated daily. According to current studies, people and devices connected to the Internet generate more than 2.5 quintillion bytes of data (Kufel et al. (2023)). Those data are generated from several different resources such as social media and internet of things devices. The current data processing technique cannot handle and process the massive amount of data generated daily. Therefore, the role of big data and the learning technique has been proposed (Alani (2021)).

Big data encompasses a suite of technologies and architectures designed to process, organise, and extract value from vast amounts of information. Facilitates distributed parallel processing, enabling the management of data that adheres to the 3V paradigm: Volume (exponential growth in the quantity of data), velocity (varying frequency of data generation, which can occur in fractions of a second) and variety (various sources and formats of data, both structured and unstructured) (Kufel et al. (2023)).

Following the preparation and processing of the collected data, the application of analytics and intelligent algorithms commences. Big data analytics facilitates the discovery of market trends, consumer preferences, and the prediction of customer behaviour through the identification of dependencies between input variables or through the use of data mining tools and techniques. Data mining involves the exploration of data to uncover hidden patterns, relationships, and previously unrecognised correlations, enabling prediction and responsive action (Kufel et al. (2023)).

B.8.2 Types of ML Algorithms

Machine Learning (ML) algorithms are instrumental in detecting a wide range of cyber-attacks, extending beyond network intrusion detection systems to encompass various attack vectors. These algorithms are broadly categorised into two types:

1. **Supervised Learning:** This approach involves training models on labelled datasets that include both input features and corresponding desired outputs. The labelled data enable the model to learn mappings that predict accurate results from given inputs. For example, supervised learning has been applied effectively in detecting malware by analysing patterns in known malicious code (Kufel et al. (2023)).
2. **Unsupervised learning:** In contrast, unsupervised learning works with unlabelled data, identifying hidden structures by clustering the data based on its statistical properties. This method uncovers latent patterns without predefined classifications or categories, making it valuable for detecting novel or emerging threats that lack historical data. Recent studies have demonstrated the efficacy of unsupervised learning in identifying zero-day attacks by recognising anomalies in system behaviour (Khan and Bokhari (2023)).

The effectiveness of ML in attack detection is largely dependent on the selection of features that represent the state of an object. Selecting relevant and precise features is crucial to avoid inaccuracies and inefficiencies. To achieve optimal learning outcomes, redundant or unrelated features must be eliminated through feature reduction or dimensionality reduction. This technique is indispensable for managing high-dimensional, noisy datasets by either retaining significant features (feature selection) or generating new, condensed feature sets. Current studies emphasised the importance of feature selection in improving detection performance, noting that irrelevant features can lead to increased false positives (Arikumar et al. (2022)).

B.8.3 Deep Learning and Its Role in Attack Detection

Deep Learning (DL), a subset of ML, is focused on the extraction of abstract representations of high-level data. DL leverages Artificial Neural Networks (ANNs) with multiple hidden layers, enabling algorithms to mimic the cognitive functioning of the human brain. This capability has garnered significant interest in its application for detecting various types of cyber-attacks:

- *Addressing Zero-Day Attacks:* Traditional signature-based detection systems struggle to identify novel or zero-day attacks. Supervised ML algorithms exhibit significant potential in detecting such unknown threats by learning patterns indicative

of malicious behaviour. [Dubey et al. \(2023\)](#) demonstrated the effectiveness of supervised learning in detecting intrusions by analysing network traffic data.

- *Reducing False Recognition Rates:* ML can effectively mitigate the high false recognition rates inherent in traditional detection systems. ML-equipped systems integrate statistical, genetic, and heuristic methods, or a combination thereof, to decode complex attack patterns, improving detection rates, while reducing false negatives. Recent research has shown that DL-based detection systems can effectively identify sophisticated cyberattacks by learning high-level abstractions of network traffic ([Apruzzese et al. \(2023\)](#)).
- *Enhanced Accuracy and Speed:* ML algorithms excel in capturing the complex properties of cyberattacks, resulting in improved detection accuracy and faster processing. [Arikumar et al. \(2022\)](#) highlighted the use of clustering techniques to detect previously unknown threats in network systems.
- *Handling Variants in Attack Patterns:* Traditional detection systems often fail to detect minor variations in attack patterns. ML-based systems, with their ability to continuously learn from evolving data patterns, efficiently identify subtle variations in system behaviour, making them effective against variant detection. [Srivastava et al. \(2023\)](#) discussed a novel technique for predicting upcoming attacks in a network based on several data parameters, highlighting the adaptability of ML models.
- *Evolving Threats and Reduced Update Dependency:* traditional signature-based detection systems require frequent updates to address evolving attack patterns. ML approaches, particularly those employing clustering and outlier detection, reduce the need for constant updates by adapting to new threats autonomously. [Abreu and Abelém \(2023\)](#) proposed an online attack detection and network traffic classification system, which combines stream ML, DL, and ensemble learning techniques, showcasing the adaptability of ML-based systems.
- *Efficient Resource Usage:* Signature-based detection systems perform computationally expensive database matches, which become increasingly inefficient as the size of the database increases. In contrast, ML-based systems operate with lower computational demands, enabling more efficient resource utilisation. [Srivastava et al. \(2023\)](#) discussed the use of ML and data mining techniques to enhance network security, highlighting the efficiency of ML-based systems in resource use.

The advantages mentioned above underscore the suitability of ML techniques for improving attack detection across various domains. A taxonomy of popular ML algorithms has been extracted from recent research to assist readers in identifying the types of algorithms applied in reviewed detection implementations. Figure B.10 shows the different types of ML algorithms used for attack detection. Although a detailed explanation of

specific algorithms is beyond the scope of this thesis, a comprehensive systematic review is presented in Chapter 3, providing extensive insights into ML methodologies for attack detection in IoT. This integration of ML and DL continues to advance the field of cybersecurity, offering efficient, scalable, adaptive, and adaptive solutions to evolving threats (Chaabouni et al. (2019b)).




FIGURE B.10: Machine Learning Algorithms

B.9 Security Problems Comparison Between virtualised IoT and Traditional IoT Network

For the industrial and academic research sectors, cloud-based virtualisation approaches bring about new possibilities and limitations. Researchers are working to find solutions for the new problems brought about by the fusion of these technologies. Industrial engineers will modernise their operations and create cutting-edge applications on the basis of this opportunity. The IoT cloud environment combines traditional IoT, SDN, VFN, and cloud technologies. This section illustrates the critical areas that differentiate cloud IoT from traditional IoT networks.

- **Artificial Intelligence (Machine Learning and Deep Learning Algorithms):** Many technologies and apps now use AI and deep learning algorithms to complete jobs that require human intelligence. Integrating AI techniques into IoT cloud approaches is recommended to create intelligent applications and analyse future service delivery performance (Javed et al. (2020)).

- **Number of sensor nodes:** The risk of a security breach increases as the number of sensor nodes in the IoT cloud fields increases. Here, it is essential to provide a security tool that can expose accurate sensors from false ones and spot sensor node failure (Almurisi and Tadisetty (2022); Samaniego et al. (2018)).

- IoT-Cloud allows amalgamation of multiple sensor networks in the cloud platform: Massive data will be collected and stored in cloud database systems as a result of the growth in nodes and applications. Big data requires efficient analytical methods, practical algorithms, and a robust security mechanism to protect processes and data to extract information from massive amounts of data (Javed et al. (2020)).
- Fog computing: Using fog computing, processing duties will be moved from cloud resources to edge devices close to IoT networks to decrease processing lag and improve reaction time. Fog technology affects cloud-based IoT security, which researchers should study (Muniswamaiah et al. (2021); Habibi et al. (2020)).
- Resource allocation: The management of cloud resources and the mapping of the sensors at various geo-locations provide several difficulties. In order to use resources effectively, it is challenging to develop a method for resource allocation in IoT Cloud that can match the appropriate device to the applications (Javed et al. (2020); Ali et al. (2021)).

B.10 Comparison of Different Vulnerabilities and Attacks in Virtualised IoT Environments.

Due to the many difficulties with the design of the IoT system, companies and developers must be aware of several possible risks. IoT security issues are discussed in several research publications, including the surveys by (Alam et al. (2019); Almurisi and Tadisetty (2022), which detail issues with object identification, authentication, IoT privacy, etc., in IoT networks. The following figure and more information are provided below. A categorisation based on the design problems of IoT systems is shown. Figure B.11 addresses the main challenges and threat in cloud environments of the IoT.




FIGURE B.11: IoT Challenges and Threats

- **Heterogeneity and Interoperability:** Back-end IoT solutions consider the use of sensors, actuators, and gateways offered by several manufacturers and perhaps having various versions. In order to achieve the previous aim, it is necessary to employ negative managing interoperability across heterogeneous devices. This component may be inundated with bogus requests, resulting in DoS attacks. Spoofing, routing attacks, and MITM are more likely to happen in such a diverse environment than inhomogeneous systems. It is simpler for a rogue node to pose as a real thing, obtain access to data without authorisation, and convey message injection between two nodes (Chaabouni et al. (2019a); Mohan et al. (2022); Almurisi and Tadisetty (2022)).
- **Connectivity:** Connectivity between various system components, physical or services, is necessary for the IoT. In the first scenario, data from peripheral devices (such as sensors) must be joined to an IP network through bridging devices, which might lead to MITM and routing assaults. Changes in the availability of connectivity services should be communicated to the relevant devices so that they do not unintentionally overload the system with repeated requests for unavailable services. A DoS attack may result from such a deluge. The QoS of the IoT network could also be significant, especially in emergency scenarios (Chaabouni et al. (2019a); Aazam et al. (2022); Alam et al. (2019)).
- **Mobility and Scalability:** The components of the IoT system can move continuously in the field, changing their linked bridges. This frequently results in interrupted discontinuities and connections to illegitimate services. Furthermore, it increases the possibility of attacks such as repudiation, MITM, DoS, sinkhole, and wormhole increases (Antão et al. (2018); Gong et al. (2015); Zhang et al. (2021a); Chaabouni et al. (2019a)).
- **Addressing and Identification:** IoT field devices often connect over short distances using low-power radios. To achieve it, coordinator nodes assign peer devices local addresses that do not adhere to a common standard. Since the bridge conceals these addresses, harmful activities are no longer detectable. As a result, it is challenging to isolate rogue nodes and identify spoofing and repudiation attacks. In addition, the node is not filtered from the outside network when attempting to access privileges that are not allowed (Almurisi and Tadisetty (2022); Merzouki et al. (2019); Mohan et al. (2022); Chaabouni et al. (2019a)).
- **Spatio-Temporal Services:** The amplitude of a spatial-temporal impulse can be used to describe events in IoT. Data from IoT devices inside the same systems have thus to have a realistic temporal behaviour and geographical localisation. However, to prevent replay assaults, these spatio-temporal tags must be secured against hostile users. Additionally, no unauthorised users can view the user's location information (Chaabouni et al. (2019a); M Zeebaree et al. (2019)).
- **Resource constraints:** Due to their small size, most IoT peripheral devices have limited computational power, onboard memory, network bandwidth, and energy resources. Due to the resource limitations of intelligent devices and sensors, tampering, information

leakage, and node cloning are all potential assaults. Light solutions are of utmost importance since the deployment of cryptographic systems is bound by resources (Aazam et al. (2022); Almurisi and Tadisetty (2022); Chaabouni et al. (2019a)).

- **Data Interchange:** The kind of hardware, computing power, and storage capacity affect the encryption algorithms. Information leakage and other security flaws are caused by improper selection. Denial of service and resource depletion attacks are additional ways in which nodes that encrypt data may be targeted. End-to-end encryption is preferred because of this (Al-Qaseemi et al. (2017); Susilo and Sari (2021); Anuradha et al. (2021); Chaabouni et al. (2019a)).
- **Discovery of resources and services:** Device autonomy and self-discovery should be made possible in IoT systems by deploying resource and service discovery mechanisms. These techniques should be secured with two-way authentication to prevent spoofing or the malware component from overburdening the system with bogus requests to thwart DoS attempts (Networks (2022); Chaabouni et al. (2019a); Samaniego et al. (2018); Tawalbeh et al. (2020)).
- **Trust and Privacy:** IoT smart sensor devices manage sensitive user information; as a result, confidentiality and data protection are crucial. In actuality, trust and privacy are crucial challenges for IoT-based networks. To prevent spoofing, manipulation, and information leakage threats, users, Things, and devices must authenticate via trustworthy services. With smartphones like those running the Android OS, trust and privacy are receiving greater attention (Abdel-Basset et al. (2020); Kafetzis et al. (2022); Samaniego et al. (2018); Chaabouni et al. (2019a)).

The above studies and challenges illustrate several different vulnerabilities in the IoT cloud. The table B.5 below addresses the most common vulnerabilities in IoT virtual environments:

B.11 Comparison of DDoS Attack Detection in Traditional IoT Networks and virtualised WSNs

Distributed denial-of-service attacks have become a significant threat to the security and stability of network systems, particularly in the IoT and WSNs. The increasing complexity and scale of these attacks require the development of effective detection mechanisms. In this comparison, we will discuss the differences in DDoS attack detection in traditional IoT networks and virtualised WSNs.

TABLE B.5: common vulnerabilities in IoT virtual environments

Vulnerability	Details	Literature
Suspicious flow attack	The author implemented a testbed for the attack. Moreover, the detection in IoT-centric OF switches with OpenDaylight controllers was applied. Matlab has been used in the case.	Xu et al. (2017)
DoS attack	Hybrid network architecture with ISDN support employs autonomous switching and cutting-edge routing techniques. The experiment shows that there is still an ability for DoS attacks.	Sándor et al. (2015)
Unauthorised access to network devices	The testbed with in-band topology (combined control & data plane) to allow communication between clients & authentication service. The outcomes from this experiment address that this needs further enhancements, especially with the rapid increase of the IoT nodes.	Hesham et al. (2017)
TCP flood attack	Flow establishment, monitoring, and rechecking mechanisms are used to thwart ICMP and TCP flooding assaults. Emulation based on Mininet.	Thorat and Kumar Dubey (2020)
MitM	The author implemented Bloom filters using SDN & expanded OF approach to detect MitM attacks.	Li et al. (2017)

B.11.1 Network Architecture

Traditional IoT Networks: These networks typically comprise many interconnected devices, including sensors, actuators, and gateways. The devices communicate with each other and the central server over various communication protocols, including Wi-Fi, Zigbee, and Bluetooth ([Sudheera et al. \(2021\)](#)).

virtualised WSNs: Virtualisation technology is applied to WSNs to create a more flexible, scalable, and efficient network. Virtualised WSNs consist of virtual sensor nodes (VSNs) deployed on virtual machines (VMs) or containers. The VSNs communicate with a centralised controller, which manages the allocation of resources and monitors the performance of the network ([Almurisi and Tadisetty \(2022\)](#)).

B.11.2 DDoS Attack Detection Techniques

Traditional IoT Networks: In traditional IoT networks, detection of DDoS attacks often relies on anomaly or signature-based techniques. Anomaly-based techniques monitor network traffic for unusual patterns, while signature-based techniques identify known attack patterns by inspecting packet contents. These techniques can be resource-intensive and ineffective against new or evolving attack patterns ([Valdovinos et al. \(2021\)](#)).

Virtualised WSNs: In virtualised WSNs, detection of DDoS attacks can leverage the advantages of network virtualisation. For example, security mechanisms can be deployed as virtual network functions (VNF) and can easily scale up or down according to the network demand. Machine learning algorithms can detect DDoS attacks in real time by analysing network traffic patterns and continuously updating the detection models. This allows for more accurate and efficient detection of DDoS attacks (Valdovinos et al. (2021)).

B.11.3 Scalability and Resource Management

Traditional IoT Networks: In traditional IoT networks, the deployment of additional security mechanisms to detect DDoS attacks can be resource intensive and time-consuming. This can lead to increased latency and reduced network performance (Saipriya and Anand (2021)).

virtualised WSNs: Virtualised WSNs offer improved scalability and resource management, which allows the efficient deployment of DDoS attack detection mechanisms. The centralized controller can dynamically allocate resources to the security VNFs based on the network demand, ensuring optimal performance and minimizing latency (Saipriya and Anand (2021)).

B.11.4 Adaptability and Resilience

Traditional IoT Networks: DDoS attack detection in traditional IoT networks can be less adaptive and resilient, as the detection mechanisms might not be able to respond quickly to new or evolving attack patterns (Mergendahl et al. (2017)).

virtualised WSNs: The virtualised architecture in WSNs enables greater adaptability and resilience against DDoS attacks. The centralised controller can quickly update the security VNFs with new detection models to respond to evolving threats. Furthermore, the virtualised infrastructure can be reconfigured to isolate and mitigate the impact of DDoS attacks on the network (Mergendahl et al. (2017)).

B.11.4.1 Comparative analysis of existing main DDoS defence and detection mechanisms

A critical examination of existing defence solutions is warranted to ascertain the insights necessary to devise an optimal defence strategy against DDoS attacks. Table B.6 delineates these defence solutions and their proposed system models, salient advantages, and vulnerabilities to facilitate informed decisions regarding their potential applicability to

address emerging DDoS attack trends. Furthermore, the specific types of attacks these defence mechanisms target are also enumerated in the table.

Most conventional methods exhibit insufficient efficacy in detecting and mitigating application-layer attacks; conversely, machine learning-based solutions actively address such assaults by employing efficient and lightweight classification algorithms. This aptitude renders machine learning solutions well suited to satisfying contemporary IoT security requirements. Machine learning-based defence models have emerged as a prevalent trend in countering DDoS attacks, due to their ability to accurately detect and predict millions of network intrusions, surpassing other mitigation- and prevention-based defence mechanisms. Consequently, they have the capabilities to effectively manage IoT vulnerabilities (Aljuhani (2021)).

Recent investigations in this domain have explored numerous detection techniques, employing machine learning classification algorithms such as K-nearest neighbours, Support Vector Machines (SVMs), Deep Learning, Random Forests, and Neural Networks (Cui et al. (2016) Susilo and Sari (2021)).

Table B.6 shows A Comparative Analysis of detection techniques utilised in various defence mechanisms for DDoS attacks

TABLE B.6: A Comparative Analysis of Detection Techniques Utilised in Various Defense Mechanisms for DDoS Attacks

Detection Techniques	Advantages	Constraints	Implementation (Defense Mechanisms)
Network Flow Sampling (Misra et al. (2011)): Method that involves sampling of one packet per N packets (depending on sampling rate) within a standard network configuration. Abnormal flows are identified based on header properties.	Owing to its storage efficiency (only header data are considered), it is valuable to analyse the prevalent patterns of malware-infected packets over an extended period (Koroniotis et al. (2019)).	Sampling is ineffective in IoT network configurations, as one compromised IoT device will transmit a single packet per sampling session, resulting in an inefficient process for predicting overall network traffic. The exclusive capture of header data renders it incapable of detecting modern malware variants.	Learning Automata-based Defence Mechanism (Misra et al. (2011))
Packet Capturing: The entire packet is captured to store all network-related properties of malware infection (Kumari and Jain (2023b)).	Beneficial for analysing Darknet traffic, as packet sizes are smaller and privacy concerns are minimal.	Demands sufficient resources and storage capacity to retain all captured information. More expensive process.	TCP/SYN Flood DDoS Attack Defense, Botnet-specific DDoS Defense (Aljuhani (2021))
Machine Learning Algorithms (Aljuhani (2021)): Machine learning models are used to detect anomalies with the desired accuracy.	Operates effectively even without specific classification labels due to unsupervised learning. Faster in comparison to rule-based detection.	The desired accuracy and precision can be achieved, but at the expense of increased memory usage and resource allocation.	Machine Learning-based and Image Processing-based Defence Models (Aljuhani (2021))

Honeypots (Traditional/IoT): A honeypot mimics a genuine vulnerable system to intentionally attract malware, facilitating the detection of malware attacks (Ahmed and Kim (2017)).	Well-suited for capturing the emerging variants of malware behaviour, particularly IoT malware.	Low-interaction honeypots can be easily identified by sophisticated malware. Deploying high-interaction honeypots requires a substantial amount of resources.	Honeypot-based DDoS Defence Systems (Ahmed and Kim (2017))

B.12 Types of DDoS Attacks Comparison Between virtualised WSNs and Traditional IoT Network

Virtualised WSNs and traditional IoT networks are two prominent types of networks that facilitate communication between devices and systems in various industries. Virtualised WSNs consist of interconnected wireless sensor nodes that collect and share data, while traditional IoT networks are a combination of devices, networks, and services that interact. Both types of networks offer unique benefits, but also face different security challenges (Almurisi and Tadisetty (2022)).

Comparison of security vulnerabilities in virtualised WSNs and traditional IoT networks. The security vulnerabilities in virtualised WSNs and traditional IoT networks have been widely discussed in the literature. Comparison of these vulnerabilities highlights the unique challenges that each typeface faces in the network and the similarities in the risks associated with DDoS attacks. Understanding these vulnerabilities is crucial for developing effective defence strategies against DDoS attacks (Vishwakarma and Jain (2020)).

Both virtualised WSNs and traditional IoT networks share some common security vulnerabilities, such as susceptibility to various DDoS attacks and the potential for compromised devices to be used as attack vectors. Additionally, the lack of proper authentication, encryption, and regular software updates can increase the risk of DDoS attacks in both types of networks (Saipriya and Anand (2021)).

Due to their distributed nature and reliance on wireless communication, virtualised WSNs face unique challenges, such as node capture, wormhole, and Sybil attacks. On

TABLE B.7: Comparative Analysis of Recent DDoS Defense Mechanisms (Part 1)

Defense Mechanisms	System Model Employed	Key Features	Identified Vulnerabilities	Targeted Attack Types
Learning Automata-based DDoS Defense (Misra et al. (2011))	Service Oriented Architecture (SOA) with a cross-layer model for network adaptability	Optimal sampling rate determination using learning automata; Applicable to all network layers due to cross-layered architecture	Overemphasis on sampling optimality; Sampling as a biased, lossy process for predicting network traffic	Replication and amplification attacks (UDP/TCP floods, ICMP floods)
Honeypots-based DDoS Defense (Anirudh et al. (2017))	Intrusion detection system featuring a decoy computer system (honeypot)	Anomalies redirection to honeypots; Utility for understanding unknown malware nature and attack vectors	Inability to implement in real-time environments; Limited efficacy against botnet-involved DDoS attacks	HTTP flood, DNS service attacks, Protocol-based Attacks (ACK/SYN floods, Ping of Death, Smurf DDoS)
Risk Transfer Mechanism as a DDoS Defense (Adat et al. (2017))	Smart home environment with an authorized third party (Economic DoS Server)	Low-rated attacks resolved at router level; High-rated attacks transferred to third party as potential threats	Specific focus on smart home environment	Protocol-based Attacks (ACK/SYN floods, Ping of Death, Smurf DDoS)
Blockchain Concept-based Defense (Javaid et al. (2018) Fu et al. (2020))	Blockchain as self-executable computer programs (Smart Contracts)	Support for smart contracts and decentralized applications (DApps) built on blockchains, including their state	Unrealistic assumption of uncompromised gateways/routers; Inability to handle advanced botnet versions	Application Layer Attacks (HTTP flood, DNS service attacks)
IoT Middleware-based DDoS Defense (Sicari et al. (2018))	IoT architecture with cross-domain middleware (Networked Smart Object, NOS)	Real-time receipt of open data feed from remote, heterogeneous IoT sources; Lightweight, secure information exchange using MQTT protocol	Limited efficacy in complex environments with multiple NOSs, numerous data sources, and malicious entities	Protocol-based Attacks (ACK/SYN floods, Ping of Death, Smurf DDoS)

TABLE B.8: Comparative Analysis of Recent DDoS Defense Mechanisms (Part 2)

Defense Mechanisms	System Model Employed	Key Features	Identified Vulnerabilities	Targeted Attack Types
Software-Defined Networking-based Defense (Fioravanti et al. (2023) Ahmed and Kim (2017))	SDN-based infrastructure with an extended version (SDNi) for multiple SDN domain handling	Attack prevention from reaching endpoint firewalls; Network congestion reduction; Prevention of malicious IoT devices from amplifying attacks	Security threat due to centralized control paradigm	Application Layer Attacks (HTTP flood, DNS service attacks)
Regulatory Solution using Mirai's Botnet Code (Jenkins (2017))	Open-source Mirai's botnet code with government or law enforcement regulation	Code modification for compromised IoT device detection by authorized government body; Vulnerability notification sent to device owner	Requirement for proper legal framework placement and redressal method updates specific to the concept within respective administrations	Protocol-based Attacks (ACK/SYN floods, Ping of Death, Smurf DDoS)
Machine Learning-based Defense Models (Aljuhani (2021))	Supervised and unsupervised learning models with neural networks for network anomaly detection and prediction	Ability to predict future attacks with fewer false positives/negatives; Suitability for detecting both traditional and IoT malware due to various classification algorithms	Model accuracy and precision dependent on the quantity and quality of training datasets	IoT botnet-based Attacks and other malware attacks
Edge Computing-based Defense (Yao et al. (2021))	Edge computing infrastructure with distributed denial of service detection and mitigation	Offloading DDoS detection and mitigation to edge devices; Reducing latency and response time; Scalability due to distributed architecture	Increased complexity and vulnerability at the edge devices; Reliance on edge device resources	IoT botnet-based Attacks, Application Layer Attacks (HTTP flood, DNS service attacks), Protocol-based Attacks (ACK/SYN floods, Ping of Death, Smurf DDoS)

TABLE B.9: Comparative Analysis of Recent DDoS Defense Mechanisms (Part 3)

Defense Mechanisms	System Model Employed	Key Features	Identified Vulnerabilities	Targeted Attack Types
Artificial Intelligence-driven Defense (Li et al. (2022a))	AI-based system utilizing multiple algorithms for real-time DDoS attack detection and mitigation	Adaptive capabilities; continuous learning; monitoring and updating defense mechanisms; Real-time attack response	High computational requirements; Possible adversarial attacks against the AI models	IoT botnet-based Attacks, Application Layer Attacks (HTTP flood, DNS service attacks), Protocol-based Attacks (ACK/SYN floods, Ping of Death, Smurf DDoS)

the other hand, traditional IoT networks are more susceptible to insecure default configurations and weak authentication mechanisms, often due to the rapid development and deployment of IoT devices (Singh and Behal (2020)).

The following sections illustrate the different types of DDoS attack that may target virtualised WSNs.

B.12.1 ACK Flood

The ACK flood attack can be a disguise for more sophisticated assaults, as packets often penetrate routers, firewalls, and other intrusion prevention or detection systems. An inundation of numerous ACK packets, typically unassociated with active connections, is directed towards the targeted server. Consequently, the system resources allocated to process legitimate incoming packets become depleted (Kumari and Jain (2023b)).

B.12.2 SYN Flood

SYN flood attacks capitalise on known vulnerabilities within the three-way handshake process that initiates a TCP connection. The primary objective of SYN flooding is to exhaust the transmission control block reserve (TCB), which stores comprehensive information about a connection, as illustrated in Figure B.12. SYN flood attack is achieved by transmitting a substantial quantity of SYN requests to the server. In response, the server sends an SYN+ACK packet and awaits the ACK reply from the client. However, the malicious user refrains from dispatching ACK packets, causing the server to wait for a non-existent ACK message. As a result, the server buffer queue becomes saturated, leading to the rejection of legitimate requests incoming (Kumari and Jain (2023b)).

B.12.3 Slowloris

The Slowloris attack establishes numerous connections to the target's web server, maintaining them for an extended duration as shown in Figure B.13. In this attack, partial HTTP requests are transmitted, and subsequent headers for each request are sent to preserve the connection. However, requests are never fulfilled. Ultimately, the victim's maximum concurrent connection pool reaches capacity, denying subsequent legitimate connection attempts (Kumari and Jain (2023b)).




FIGURE B.12: Progression of SYN flood attack




FIGURE B.13: Progression of Slowloris attack

B.13 Applications of Deep Learning in Security Detection

Neural networks, as a powerful subset of machine learning, have emerged as essential tools for addressing the multifaceted challenges of modern cybersecurity. In the following, key applications of neural networks in security detection are discussed, highlighting their contributions to IoT networks, cloud computing, and Software-Defined Networks (SDNs) (Mittal et al. (2023)).

B.13.1 Intrusion Detection in Cloud Computing

In cloud computing environments, neural networks are widely used for intrusion detection due to their ability to process large datasets and detect anomalous behaviours. For

example, feedforward neural networks that employ backpropagation algorithms have shown great effectiveness in identifying potential threats. A study by Abid et al. ? demonstrated the use of multilevel deep neural networks to detect distributed denial-of-service (DDoS) attacks in SDNs, achieving high precision by combining convolutional and recurrent layers (Abreu and Abelém (2023)).

B.13.2 Detection of IoT-based DDoS Attacks

The widespread adoption of IoT devices has significantly increased the surface area for cyber attacks, particularly DDoS. Neural networks, such as CNNs, have proven highly effective in this domain. (Mittal et al. (2023)) implemented a CNN-based detection mechanism for DDoS attacks, achieving an impressive accuracy of 99.92% on the Bot-IoT data set. This approach enables real-time detection by extracting spatial features from network traffic, significantly improving IoT network resilience (Srivastava et al. (2023)).

B.13.3 Mitigation in Software-Defined Networks (SDNs)

The centralised architecture of SDNs, while beneficial for network management, introduces vulnerabilities that attackers can exploit. Neural networks, particularly hybrid CNN-LSTM models, provide robust solutions to detect and mitigate these vulnerabilities. (Jimoh et al. (2024)) proposed a CNN-LSTM model for SDN security, achieving precision and recall scores that exceed 99%. This model improves the detection and mitigation of DDoS attacks by analysing the temporal and spatial characteristics of network traffic (Apruzzese et al. (2023)).

B.13.4 Multi-Class Attack Classification

Neural networks are instrumental in categorising different types of cyberattacks. Multilevel Deep Neural Networks (MLDNNs), which combine CNNs and LSTMs, are particularly effective in classifying complex attacks such as TCP- and UDP-based DDoS. The hybrid architecture presented by (Liu et al. (2022)) achieved scalability and high precision, enabling the classification of multiple attack types in large-scale SDN environments.

B.13.5 Trustworthiness Assessment in Cloud Services

In addition to detection, neural networks have been used to assess the trustworthiness of cloud services. Integrating neural networks with fuzzy logic, frameworks have been

developed to evaluate and enhance trust metrics for various network components. This approach allows for dynamic security evaluations tailored to specific user requirements, as highlighted in recent work on hybrid trust assessment models [Nita and Mihailescu \(2018\)](#).

B.13.6 Real-time Anomaly Detection

Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks excel in analysing time-series data, making them well suited for real-time anomaly detection. These models can identify deviations from typical network behaviour, providing rapid responses to emerging threats. The CNN-LSTM hybrid architecture proposed by Jimoh et al. [Jimoh et al. \(2024\)](#) effectively identifies and mitigates anomalies in real time, enhancing the security of IoT and SDN environments.

B.13.7 Preventative Measures through Behaviour Modelling

Neural networks are also applied to behaviour modelling, allowing the detection of potential threats based on deviations from expected user or system activity. Training on historical data, these models proactively identify risks, reducing the probability of successful attacks. This preventive capability is particularly valuable in environments with high stakes, such as IoT and cloud computing, as demonstrated in studies on IoT network behaviour analysis ([Lakshmana et al. \(2022\)](#)).

B.14 IoT NIDSs Based on Machine Learning Techniques

The rapid evolution and increasing sophistication of cybersecurity threats have required innovative solutions for intrusion detection. Researchers are increasingly focussing on artificial intelligence (AI) and machine learning (ML) to enhance the efficacy of intrusion detection systems (IDS). By integrating machine intelligence, IDS can significantly improve their decision-making capabilities and adaptability to novel threats ([Chaabouni et al. \(2019b\)](#)).

This section focusses on advanced network-based intrusion detection systems (NIDSs) tailored for IoT environments, highlighting approaches that leverage diverse ML techniques. These approaches range from supervised and unsupervised learning methods to hybrid frameworks, designed to address the challenges of high-dimensional data, scalability, and low latency ([Lu and Xu \(2019\)](#)).

B.14.1 Multi-Layer Perceptron (MLP) for IoT NIDS

Wu et al. (2020) implemented a supervised artificial neural network (ANN) approach to detect DoS and DDoS attacks in IoT networks. Their system utilised a three-layer feed forward MLP with unipolar sigmoid transfer functions and stochastic learning algorithms. Tested on a simulated IoT environment with 10 million UDP packets, their model achieved a detection accuracy of 99.4% with only 0.6% false positives, demonstrating robustness in high-traffic scenarios.

B.14.2 Host-Based IDS Framework (IoT-IDM)

The researcher proposed IoT-IDM, a modular host-based IDS framework that uses software-defined networking (SDN) to monitor and secure IoT smart home environments. The system used OpenFlow protocols for traffic analysis and used ML algorithms such as logistic regression and support vector machines (SVM). IoT-IDM achieved up to 98.53% accuracy, although its scalability was constrained by the high volume of network traffic (Dwivedi et al. (2018)).

B.14.3 Artificial Immune Systems (AIS) with Fog Computing

The researcher developed a lightweight IDS incorporating Artificial Immune Systems (AIS) within a fog computing architecture. Their three-tier design (edge, fog, and cloud layers) provided distributed attack detection, using genetic algorithms for anomaly analysis. The system achieved a precision of 98.35% and demonstrated advantages in scalability, low-latency processing, and the ability to detect zero-day attacks (Dubey et al. (2023)).

B.14.4 Hybrid Anomaly and Specification-Based IDS

Bostani and Sheikhan (2017) introduced a hybrid IDS that combines anomaly-based and specification-based methods to detect sinkhole and selective forwarding attacks in 6LoWPAN networks. Their two-step system used local traffic analysis at router nodes and a global anomaly detection mechanism employing the optimal-path forest (OPF) algorithm. The experimental results showed a true positive rate of 76.19% for combined attacks and 96.02% for wormhole attacks.

B.14.5 Dimensionality Reduction and Classification for IoT IDS

Pajouh et al. (2017) addressed the high dimensionality of IoT data with a two-layer reduction approach using Principal Component Analysis (PCA) and Linear Discriminant

Analysis (LDA). Their Two-Layer Dimension Reduction and Two-Tier Classification (TDTC) framework achieved an 84.86% detection rate for attacks against the NSL-KDD data set, demonstrating computational efficiency and reduced false positive rates ?.

B.14.6 Conditional Variational Autoencoder (CVAE) for IoT NIDS

López Martín (2019) introduced a CVAE-based anomaly detection model capable of reconstructing missing features, achieving up to 99% precision. Their approach combines unsupervised feature reconstruction with supervised class label incorporation, which makes it highly suitable for IoT environments characterised by frequent connectivity disruptions and incomplete datasets.

B.14.7 Fog Computing with Deep Learning

Diro and Chilamkurti (2018) proposed a distributed deep learning IDS within a fog computing framework. By decentralising data processing across fog nodes, their system enhanced scalability and localised detection accuracy. They achieved a detection rate of 96.5% with a reduced false alarm rate of 2.57%, outperforming traditional shallow learning classifiers.

B.14.8 Online Sequential Extreme Learning Machine (OS-ELM)

Prabavathy et al. (2018) utilised OS-ELM within fog computing for real-time intrusion detection. By dividing IoT networks into clusters monitored by fog nodes, the system achieved high detection accuracy (97.36%) and reduced false alarms (0.37%). This approach also enabled dynamic adaptation to new data, which is often a limitation in traditional models.

B.14.9 Semi-Supervised Learning for IoT NIDS

Rathore and Park (2018) combined supervised and unsupervised ML techniques, deploying fuzzy clustering of c-means alongside extreme learning machines (ELM). Their semi-supervised framework achieved an accuracy of 86.53% with a detection time of 11 milliseconds, providing a balanced approach to the detection of known and unknown threats.

B.14.10 Ensemble-Based NIDS for IoT

Khalifa et al. (2019) used AdaBoost to integrate decision trees, naive Bayes, and ANN for robust IoT threat detection. The system achieved detection rates between 95.25% and 99.86% in data sets, with a low false positive rate of 0.01% to 0.72%, making it highly effective for identifying botnet attacks.

B.15 Comparison of Machine Learning Algorithms Used in Attack Detection

This section provides a comprehensive comparison of machine learning algorithms employed for the detection of DDoS attacks. The focus is on the data sets used, the types of attack classes addressed, and the methodologies followed in the studies published between 2019 and 2025. Table B.10 consolidates the information from recent research, illustrating the datasets used, attack types, and machine learning techniques adopted.

The table B.10 outlines a detailed comparison of research studies focussing on detection of DDoS attacks from 2019 to 2025. Incorporating the data sets used, the machine learning approaches, and the types of attacks detected. In particular, data sets such as NSL-KDD, CICIDS2017, and CICDDoS2019 are widely used due to their comprehensive nature, encompassing a variety of attack classes like DoS, R2L, U2R, Probe, and DDoS flooding attacks.

Key Observations:

- Recent studies have shifted towards deep learning approaches such as CNN, LSTM, and hybrid methods for improved detection capabilities.
- CICIDS2017 and CICDDoS2019 have emerged as prominent datasets due to their inclusion of modern attack patterns.
- A wide range of attack classes are considered across studies, highlighting the evolving threat landscape.
- The adoption of hybrid approaches, such as CNN-LSTM and RNN-AE, showcases the trend of combining strengths from different models to achieve better detection rates.
- Simulated datasets, such as CICIDS2017 and CICDDoS2019, remain popular for their realistic attack scenarios, while synthetic data sets are used to test innovative and hard-to-detect attack patterns.

TABLE B.10: Comparison of Machine Learning Algorithms Used in DDoS Attack Detection (2019–2025)

Reference	Year	Dataset Used	Approach Used	Attack Classes in Study	ML Algorithm
Hasan et al. (2019)	2019	NSL-KDD, OBS Network	Deep CNN	DoS	Deep CNN
Bhuvaneshwari Amma and Subramanian (2018)	2018	NSL-KDD	CNN	DoS	CNN
Chen et al. (2023)	2023	KDDCup99, CICIDS2017	Multi-channel CNN	Normal, DoS, R2L, U2R	Multi-channel CNN
Shaaban et al. (2020)	2020	NSL-KDD, Wiresnark	CNN Model	TCP Flood DDoS, R2L, U2R	CNN
Virupakshar et al. (2020)	2020	KDDCup99, LAN	DT, KNN, NB, DNN	DoS, R2L, U2R, Probe	DT, KNN, NB, DNN
Zeehan et al. (2022)	2022	CICIDS2017	Ensemble RNN, CNN	Slowloris, Hulk, GoldenEye	RNN, CNN
Kim et al. (2022)	2022	KDDCup99, CIC-IDS	CNN	Smurf, SYN Flood, DDoS LOIT	CNN
Sbai and El boukhari (2020)	2020	CICDDoS2019	DNN	Flooding, UDP-based attacks	DNN
Cil et al. (2021)	2021	CICDDoS2019	DNN	SYN Flood, UDP Lag, MSSQL	DNN
Asad et al. (2020)	2020	ISCX2012, CIC2017	RNN-AE	Hulk, Slowloris, LOIT	RNN-AE
Ali and Li (2019)	2019	ISCX2012, UNSWNB15	AE, Deep ANN	Backdoors, Shellcode, Worms	AE, Deep ANN
Yang et al. (2020)	2020	Synthetic Dataset	AE	Slowloris, GoldenEye, HTTP	AE
Premkumar and Sundararajan (2021)	2021	Generated Dataset	RBF	Data Flooding, Eavesdropping	RBF
Nugraha and Murthy (2020)	2020	Synthetic Dataset	CNN-LSTM	UDP Flooding, HTTP Flooding	CNN-LSTM
de Assis et al. (2020)	2020	Simulated SDN Data	CNN	SYN Flood, Data Flooding	CNN
Li (2020)	2020	CICDDoS2019	DNN	UDP Lag, MSSQL, SNMP Flood	DNN
Liang and Znati (2019)	2019	CICIDS2017	LSTM	Slowloris, Hulk, GoldenEye	LSTM
Roopak et al. (2020)	2020	CICIDS2017	CNN+LSTM	DDoS	CNN+LSTM
He et al. (2020)	2020	6LANN, 7LANN	Transfer Learning	SYN-Type, LDAP-type	Transfer Learning
Swathi and Narsimha (2023)	2023	CICDDoS2019	Multiple Regression	Various DDoS Attacks	Multiple Regression
Mishra et al. (2020)	2020	CICDDoS2019	Multi-Classifer ML	DDoS	Multi-Classifer ML
Peloso et al. (2018)	2018	Custom Dataset	Self-Adaptive System	DDoS	Self-Adaptive System
Santos et al. (2020)	2020	Network Spectrum Data	Automatic Detection	DDoS	Automatic Detection
Gabriel et al. (2021)	2021	Unsupervised Data	Deep Learning	DDoS	Deep Learning

Overall, the research demonstrates significant advances in DDoS detection methodologies, leveraging a wide range of machine learning and deep learning algorithms. The transition from traditional ML models such as Decision Trees (DT) and Naïve Bayes (NB) to advanced techniques like CNN, DNN, and transfer learning reflects the increasing complexity of attack scenarios and the need for robust detection systems.

Future Directions:

- *Exploration of new datasets:* Future studies should focus on building more comprehensive and diverse datasets that represent real-world network conditions and attack types.
- *Integration of unsupervised learning:* With the emergence of novel and unknown attack patterns, unsupervised learning techniques can play a critical role in anomaly detection.
- *Deployment in real-world systems:* Transitioning from simulated datasets to real-world networks is essential to validate the efficacy and reliability of the proposed detection techniques.
- *Hybrid models:* Continued exploration of hybrid and ensemble learning approaches can further enhance detection accuracy and reduce false-positive rates.

B.16 Comparison of CNN and FCNN

Neural networks, a cornerstone of machine learning, have evolved significantly over the years, leading to specialised architectures such as FCNNs and CNNs. Each of these architectures serves distinct purposes based on the data's structure, dimensionality, and application requirements. This section explores the fundamental characteristics, advantages, limitations, and use cases of FCNNs and CNNs, with a focus on their application in security detection systems. By understanding their unique capabilities, researchers can make informed choices and develop more efficient hybrid models.

B.16.1 Fully Connected Neural Network (FCNN)

A FCNN, often called a Multilayer Perceptron (MLP), is a feedforward neural network consisting of an input layer, one or more hidden layers, and an output layer. Each neuron in a layer connects to every neuron in the preceding and succeeding layers, making this architecture highly versatile. FCNNs use non-linear activation functions such as ReLU or sigmoid to enable non-linear transformations of the data and rely on the backpropagation algorithm for parameter optimisation (Hnamte et al. (2024)).

Advantages: FCNNs are easy to implement and conceptually simple, making them accessible to researchers and practitioners. They can theoretically approximate any continuous function given sufficient hidden layers and neurons, as established by the Universal Approximation Theorem. FCNNs excel in handling low-dimensional structured data sets due to their ability to model complex mappings between inputs and outputs (Hnamte et al. (2024)).

Disadvantages: Despite their simplicity, FCNNs require a large number of parameters, increasing the risk of overfitting without robust regularisation techniques such as dropout or L2 regularisation. In addition, they fail to capture spatial or local relationships in the data, as each neuron considers all input features independently. These limitations make FCNNs less suitable for processing high-dimensional or unstructured data, such as images or videos (Wang et al. (2021)).

B.16.2 Convolutional Neural Network (CNN)

CNNs are specialised architectures designed to process grid-like data, such as images, sequential data, or volumetric data sets. A typical CNN architecture consists of convolutional layers for local feature extraction, pooling layers for dimensionality reduction, and fully connected layers for classification. Convolutional layers use filters that slide over the input to detect local patterns, while parameter sharing reduces computational overhead compared to FCNNs (Jimoh et al. (2024)).

Advantages: CNNs excel at feature extraction, allowing them to learn hierarchical representations from raw data. Their ability to identify spatial and temporal relationships makes them ideal for tasks involving high-dimensional data, such as image and video processing. Parameter sharing and sparse connectivity further reduce the number of trainable parameters, enhancing their efficiency (Najar and Manohar Naik (2024)).

Disadvantages: The training of CNNs demands large datasets and computational resources, often requiring GPUs or TPUs. In addition, CNNs lack interpretability, as their decision-making processes are difficult to decode. These challenges, along with their complexity, can limit their applicability in certain real-time or resource-constrained scenarios (Liu et al. (2022)).

Applications in Security Detection: CNNs are highly effective at detecting patterns within high-dimensional or unstructured security data, such as network traffic, system call sequences, or images from security cameras. By learning local and hierarchical characteristics, CNNs are adept at identifying subtle patterns indicative of security threats, making them invaluable for advanced security detection systems (Chaabouni et al. (2019b)).

B.16.3 Comparison and Application in Security Detection

Both FCNNs and CNNs offer distinct advantages and disadvantages, which must be carefully considered based on the specific application context. Table B.11 highlights their key differences.

TABLE B.11: Comparison of FCNN and CNN in Security Detection

Aspect	FCNN	CNN
Data Type	Structured, low-dimensional	Unstructured, high-dimensional
Feature Extraction	Manual	Automatic hierarchical learning
Parameter Efficiency	Low	High (due to parameter sharing)
Training Requirements	Relatively low	High (large datasets and computational power)
Suitability in Security	Structured security features	Data with spatial/temporal correlations
Model Interpretability	High	Low

The choice of neural network architecture for security detection tasks depends on the data structure, application requirements, and resource constraints. FCNNs are particularly effective for structured, low-dimensional datasets, such as those commonly used in DDoS detection. In such scenarios, FCNNs offer simplicity, computational efficiency, and fast training times. For example, in my application of FCNN for DDoS detection, the structured nature of the data, consisting of features such as packet size, duration, and protocol type, made FCNN the most suitable choice. Their ability to process these compact feature representations effectively validated their superiority in this context over more complex architectures such as CNNs (Najar and Manohar Naik (2024); Hnamte et al. (2024)).

While CNNs are often praised for their ability to extract hierarchical features from high-dimensional, unstructured data, such as raw network traffic or visual logs, these capabilities are not always required in DDoS detection tasks where the data is well structured. For this reason, my work exclusively focused on FCNNs as they were both computationally efficient and sufficient for achieving high performance.

Future research could explore the enhancement of FCNN-based models by incorporating complementary techniques instead of integrating CNNs, which are more suitable for unstructured data. For instance, attention mechanisms or feature selection methods could improve FCNNs' ability to focus on the most relevant features in DDoS detection tasks. Furthermore, lightweight optimisations, such as reducing the number of parameters while maintaining accuracy, could make FCNNs even more effective in resource-constrained environments (Hekmati et al. (2024)).

While current research highlights hybrid models that combine FCNNs and CNNs, such approaches are more applicable to datasets involving mixed structured and unstructured features, such as those with spatial or temporal correlations. In the specific context of structured DDoS detection datasets, FCNN alone have been demonstrated to be

sufficient and efficient, as confirmed in recent studies (Abid et al. (2024b); Apruzzese et al. (2023)).

Future exploration of FCNNs could also focus on transfer learning, where pre-trained FCNN models are fine-tuned for specific security tasks, further improving efficiency and accuracy in DDoS detection.

B.16.3.1 Applications of FCNN in Anomaly Detection

FCNNs have demonstrated significant utility in anomaly detection due to their ability to model complex relationships in structured data (Hekmati et al. (2024)). Anomaly detection involves identifying patterns in the data that deviate from expected behaviour, such as fraud detection, network intrusion detection, and DDoS detection. The simplicity, computational efficiency, and capacity for modelling non-linear relationships make FCNNs a suitable choice for many anomaly detection applications. This section summarises recent studies (2019-2025) that used FCNN in various domains of anomaly detection, highlighting their effectiveness.

Advantages of FCNN in Anomaly Detection

- **Structured Data Handling:** FCNNs are well suited for data sets with low-dimensional structured features, such as network traffic logs or system metrics.
- **Efficient Training:** They have fewer computational requirements compared to complex architectures such as CNNs or Recurrent Neural Networks (RNNs).
- **Flexibility:** FCNNs can be easily adapted for binary or multiclass classification tasks in anomaly detection.
- **Wide Applicability:** FCNNs are effective in various domains, including cybersecurity, healthcare, and finance.

Table B.12 summarises recent studies (2019-2025) that highlight the use of FCNNs in anomaly detection in various domains.

The studies summarised in Table B.12 illustrate the versatility of FCNNs in anomaly detection in various domains. Common themes emerging from these studies include:

- **High Accuracy:** FCNNs consistently achieve high accuracy across datasets, particularly in structured environments such as network traffic logs or transactional data.

TABLE B.12: Summary of Studies on FCNN in Anomaly Detection

Study	Application	Dataset	Key Findings
Ali et al. (2024)	DDoS Detection	IoT Network Traffic	FCNN outperformed CNN in structured IoT datasets due to its computational efficiency and fast training.
Chen et al. (2023)	Intrusion Detection	NSL-KDD, UNSW-NB15	FCNN achieved high accuracy and recall in structured intrusion detection datasets, showing resilience to overfitting.
Wang and Wang (2022)	Fraud Detection in Financial Transactions	Credit Card Fraud Dataset	FCNN efficiently detected fraudulent transactions, achieving over 95% accuracy with limited feature preprocessing.
Zhao et al. (2024)	Malware Detection	PE Malware Dataset	FCNN successfully classified malware types, demonstrating scalability in feature-rich datasets.
Li et al. (2022b)	Network Anomaly Detection	CICIDS 2017	FCNN demonstrated superior performance in detecting network anomalies compared to traditional machine learning models.
Patel et al. (2024)	Healthcare Anomaly Detection	Medical IoT Devices	FCNN effectively detected anomalies in IoT medical sensor data with low latency and high accuracy.
Zhou et al. (2022)	Power Grid Anomaly Detection	Smart Grid Network Data	FCNN was employed for anomaly detection in power grids, achieving significant improvement in fault detection.

- **Computational Efficiency:** Compared to more complex architectures such as CNNs or RNNs, FCNNs require less computational power, making them suitable for real-time anomaly detection.
- **Adaptability:** FCNNs have been successfully adapted for tasks ranging from binary anomaly detection (e.g., DDoS detection) to multi-class classification (e.g., malware detection).
- **Feature Importance:** While FCNNs are effective in structured datasets, their performance is heavily based on the quality and relevance of input features, making feature engineering a critical step.

Future research could focus on integrating attention mechanisms, adaptive learning rates, and hybrid feature extraction techniques to further enhance FCNNs, particularly in unstructured or high-dimensional anomaly detection tasks.

Appendix C

System Requirements Analysis

This section establishes the complete system requirements that govern the design and implementation of the DDoS detection framework. Requirements are categorised into non-functional requirements (quality attributes, performance characteristics, operational constraints) and functional requirements (specific capabilities, features, behaviours) essential for effective vIoT security.

All requirements and technology choices documented in this section directly informed the testbed architecture and the development of the AI framework, ensuring alignment between the system capabilities and the research objectives.

C.0.1 Non-Functional Requirements

Non-functional requirements define quality attributes critical to system robustness, scalability, and deployability in production environments. These requirements ensure the detection framework maintains consistent performance under diverse operational conditions whilst supporting iterative improvement and independent verification.

Table C.1 consolidates the eight principal non-functional requirements governing system design.

TABLE C.1: Non-Functional Requirements Governing System Design

Requirement Category	Specification and Rationale
Robustness	The system must demonstrate resilience to diverse input conditions, including noisy sensor data, statistical outliers, and adversarial perturbations. Data preprocessing methodologies must effectively handle missing values, duplicate entries, and protocol anomalies. Incorporation of multiple detection models in an ensemble configuration mitigates vulnerabilities associated with single-model architectures. Rigorous stress testing across marginal cases and unexpected inputs uncovers potential weaknesses before deployment. Robust error handling and graceful degradation protocols ensure continuous operation even under challenging conditions Habibullah et al. (2023) .
Scalability	The architecture must accommodate growing computational loads resulting from increasing numbers of virtualised IoT devices and expanding data volumes. Dynamic resource allocation mechanisms—whether through physical server expansion or cloud computing integration—maintain system stability during peak usage or sudden device surges. Modular and distributed design enables components to be integrated or removed without adversely affecting overall performance. Software optimisations including load balancing, distributed databases, and parallel processing facilitate efficient management of large-scale datasets Joshi et al. (2022) .
Performance	The system must exhibit low latency (detection within 60 seconds for high-intensity attacks, 15 minutes for stealthy attacks), minimal response times, and high throughput sufficient for real-time processing requirements. Optimised data structures and algorithms minimise memory consumption and computational overhead. Data pipeline refinement through compression, caching, and incremental processing enhances efficiency. Model deployment techniques including pruning and quantisation reduce computational requirements whilst preserving detection accuracy Habibullah et al. (2023) .

Continued on next page

Table C.1 – *Continued from previous page*

Requirement Category	Specification and Rationale
Security	Robust security protocols protect data integrity and model confidentiality from unauthorised access or manipulation. Encryption mechanisms secure data transmission and storage. Regular vulnerability assessments identify potential attack vectors. Secure communication channels prevent man-in-the-middle attacks during data collection and model updates. Access control mechanisms restrict model modification to authorised personnel Dwivedi et al. (2018).
Maintainability	Well-structured documentation, clear modular architecture, and well-defined interfaces facilitate system updates, debugging, and iterative enhancements. Version control processes track changes across codebase and configuration files. Automated testing frameworks validate functionality after modifications. Good maintenance practices prolong system lifecycle and reduce total cost of ownership Habibullah et al. (2023).
Usability	Intuitive interfaces and straightforward system interactions accommodate both technical and non-technical stakeholders. Clear visualisations of attack patterns, detection results, and system status facilitate rapid comprehension. Streamlined workflows minimise training requirements for network administrators. Accessible documentation supports independent deployment and configuration. User feedback mechanisms enable continuous interface refinement Joshi et al. (2022).
Efficiency	Resource utilisation minimisation encompasses computational power, memory consumption, and energy expenditure. Efficient algorithms reduce unnecessary computation. Hardware acceleration (GPU utilisation where appropriate) enhances processing speed. Optimised data pipelines eliminate redundant operations. Energy-efficient implementation supports sustainable deployment, particularly relevant for resource-constrained IoT environments Habibullah et al. (2023).

Continued on next page

Table C.1 – *Continued from previous page*

Requirement Category	Specification and Rationale
Reproducibility	All experimental procedures, training processes, and evaluation results must be fully reproducible by independent researchers. Comprehensive logging of hyperparameters, random seeds, and algorithmic configurations enables exact replication. Data versioning tracks dataset modifications. Detailed pipeline documentation specifies data preprocessing steps, feature extraction procedures, and model training protocols. Public availability of datasets and source code facilitates community validation and extension Habibullah et al. (2023) .

These non-functional requirements establish quality benchmarks evaluated throughout the development process and validated in the experimental results (Chapter 5).

C.0.2 Functional Requirements

Functional requirements define specific capabilities and behaviours the detection system must exhibit to effectively identify DDoS attacks in vIoT environments. These requirements encompass the complete machine learning pipeline from data acquisition through result reporting.

Table C.2 enumerates the nine principal functional requirements.

TABLE C.2: Functional Requirements Defining System Capabilities

Requirement Category	Specification and Implementation Approach
Completeness (End-to-End Pipeline)	The system must address all essential stages of the machine learning lifecycle: data collection, preprocessing, feature engineering, model training, validation, and deployment Akkiraju et al. (2020) . This holistic coverage ensures the DDoS detection framework manages the entire pipeline—from real-time traffic acquisition and cleaning to final classification and alert generation—thereby maximising effectiveness in operational deployment. Gaps in pipeline coverage introduce vulnerabilities and reduce overall system reliability.
Adaptability	The system must adapt to evolving data patterns and operational environments Joshi et al. (2022) . Adaptation mechanisms include: (1) periodic model retraining on updated datasets, (2) concept drift detection algorithms identifying long-term shifts in attack patterns, (3) dynamic feature engineering enabling addition or removal of features in response to novel DDoS variants, and (4) automated hyperparameter tuning maintaining optimal performance as traffic characteristics change. Such adaptability ensures sustained high detection performance over extended deployment periods.
Interpretability	Clear explanation methods and transparency in decision-making processes are crucial for user trust and regulatory compliance. Feature importance analysis identifies which network characteristics most strongly influence attack classifications. Model-agnostic explanation frameworks (LIME, SHAP) provide instance-level justifications. These capabilities enable network administrators to understand rationale behind specific alerts, facilitating informed response decisions and supporting forensic analysis Akkiraju et al. (2020) .

Continued on next page

Table C.2 – *Continued from previous page*

Requirement Category	Specification and Implementation Approach
Accuracy	The system must attain and sustain robust prediction accuracy, meeting or exceeding performance metrics defined for the target domain. Performance monitoring employs comprehensive metric suites including Accuracy, Precision, Recall, F1-score, and AUC-ROC. Regular benchmarking against state-of-the-art models validates competitive performance. Continuous improvement processes incorporate lessons from false positives and false negatives to refine detection capabilities Joshi et al. (2022) .
Data Collection & Preprocessing	Seamless ingestion of real-time network traffic data enables timely DDoS detection. Data pipelines incorporate mechanisms for noise reduction, consistency checks, and data quality validation ensuring high-quality input prior to feature extraction. Configurable preprocessing modules accommodate diverse network protocols (HTTP, HTTPS, MQTT, CoAP) and heterogeneous device configurations. Robust handling of malformed packets and protocol violations prevents pipeline failures Akkiraju et al. (2020) .
Anomaly Detection	Since DDoS attacks manifest as atypical traffic patterns, the system must incorporate advanced anomaly detection techniques. Statistical threshold methods (z-score analysis, interquartile range filtering) identify obvious deviations. Sophisticated techniques including clustering algorithms and autoencoder-based reconstruction detect subtle anomalies. Identified anomalies trigger further analysis or direct alerting depending on severity and confidence levels Joshi et al. (2022) .

Continued on next page

Table C.2 – *Continued from previous page*

Requirement Category	Specification and Implementation Approach
DDoS Detection (Multi-Vector)	Dedicated detection modules must recognise both known and emerging DDoS attack vectors. Multi-model integration combining decision trees, random forests, support vector machines, and deep learning architectures captures diverse forms of malicious behaviour. Attack-specific signatures enable detection of SYN floods, UDP floods, ICMP floods, HTTP floods, and application-layer attacks. Rigorous validation and testing processes ensure robust performance under real-world conditions Joshi et al. (2022) .
Reporting & Visualisation	The system must provide clear, intuitive dashboards and reports supporting rapid, informed decision-making. Visualisations of attack patterns, temporal trends, and model outputs facilitate both high-level strategic analysis and granular tactical investigation. Automated report generation summarises key findings including attack types detected, affected devices, severity assessments, and recommended mitigation actions. Export capabilities enable integration with external Security Information and Event Management (SIEM) systems Akkiraju et al. (2020) .
Real-Time Monitoring & Alerts	Proactive alerting mechanisms promptly notify stakeholders of potential DDoS incidents. The alerting module operates with minimal latency, issuing notifications via configurable channels (email, SMS, webhooks, SNMP traps). Alert severity levels (Critical, High, Medium, Low) enable prioritisation. Integrated real-time dashboards provide ongoing supervision enabling near-instantaneous responses to evolving threats Joshi et al. (2022) .

These functional requirements guided design decisions throughout the development of the system, with implementation details presented in (testbed architecture) and (development of the AI framework).

References

- Mohammad Aazam, Saif Ul Islam, Salman Tariq Lone, and Assad Abbas. Cloud of Things (CoT): Cloud-Fog-IoT Task Offloading for Sustainable Internet of Things, 2022. *IEEE Transactions on Sustainable Computing*, vol. 7, no. 1, pp. 87–98 and <https://doi.org/10.1109/TSUSC.2020.3028615>.
- Kaiser Abbas, Sadaf Hina, Hamza Sajjad, Khurram Shabih Zaidi, and Rehan Akbar. Optimization of predictive performance of intrusion detection system using hybrid ensemble model for secure systems, 2023. *PeerJ Computer Science*, vol. 9, pp. e1552 and <https://doi.org/10.7717/peerj-cs.1552>.
- Mohamed Abdel-Basset, Gunasekaran Manogaran, Abdualлах Gamal, and Victor Chang. A Novel Intelligent Medical Decision Support Model Based on Soft Computing and IoT, 2020. *IEEE Internet of Things Journal*, vol. 7, no. 5, pp. 4160–4170 and <https://doi.org/10.1109/JIOT.2019.2931647>.
- Yawar Abbas Abid, Jinsong Wu, Guangquan Xu, Shihui Fu, and Muhammad Waqas. Multilevel Deep Neural Network Approach for Enhanced Distributed Denial-of-Service Attack Detection and Classification in Software-Defined Internet of Things Networks, 2024a. *IEEE Internet of Things Journal*, vol. 11, no. 14, pp. 24715–24725 and <https://doi.org/10.1109/JIOT.2024.3376578>.
- Yawar Abbas Abid, Jinsong Wu, Guangquan Xu, Shihui Fu, and Muhammad Waqas. Multilevel Deep Neural Network Approach for Enhanced Distributed Denial-of-Service Attack Detection and Classification in Software-Defined Internet of Things Networks, 2024b. *IEEE Internet of Things Journal*, vol. 11, no. 14, pp. 24715–24725 and <https://doi.org/10.1109/JIOT.2024.3376578>.
- Diego Abreu and Antônio Abelém. OMINACS: Online ML-Based IoT Network Attack Detection and Classification System, 2023. <http://arxiv.org/abs/2302.09225><https://dx.doi.org/10.1109/LATINCOM56090.2022.10000544>.
- Vipindev Adat, B. B. Gupta, and Shingo Yamaguchi. Risk transfer mechanism to defend DDoS attacks in IoT scenario, 2017. 2017 IEEE International Symposium on Consumer Electronics (ISCE), pp. 37–40 and <https://doi.org/10.1109/ISCE.2017.8355541>.

- Rasheed Ahmad and Izzat Alsmadi. Machine learning approaches to IoT security: A systematic literature review[Formula presented], 2021. Internet of Things (Netherlands), vol. 14 and <https://doi.org/10.1016/j.iot.2021.100365>.
- Hossein Ahmadvand, Chhagan Lal, Hadi Hemmati, Mehdi Sookhak, and Mauro Conti. Privacy-Preserving and Security in SDN-Based IoT: A Survey, 2023. IEEE Access, vol. 11, pp. 44772–44786 and <https://doi.org/10.1109/ACCESS.2023.3267764>.
- M. Ejaz Ahmed and Hyounghick Kim. DDoS Attack Mitigation in Internet of Things Using Software Defined Networking, 2017. 2017 IEEE Third International Conference on Big Data Computing Service and Applications (BigDataService), pp. 271–276 and <https://doi.org/10.1109/BigDataService.2017.41>.
- Iman Akbari, Mohammad A. Salahuddin, Leni Aniva, Noura Limam, Raouf Boutaba, Bertrand Mathieu, Stephanie Moteau, and Stephane Tuffin. Traffic classification in an increasingly encrypted web, 2022. Communications of the ACM, vol. 65, no. 10, pp. 75–83 and <https://doi.org/10.1145/3559439>.
- Rama Akkiraju, Vibha Sinha, Anbang Xu, Jalal Mahmud, Pritam Gundecha, Zhe Liu, Xiaotong Liu, and John Schumacher. Characterizing Machine Learning Processes: A Maturity Framework, 2020. pp. 17–31 and https://doi.org/10.1007/978-3-030-58666-9_2.
- Sharmin Aktar and Abdullah Yasin Nur. Towards DDoS attack detection using deep learning approach, 2023. Computers & Security, vol. 129, pp. 103251 and <https://doi.org/10.1016/j.cose.2023.103251>.
- Abdullah M. Al-Ansi, Askar Garad, Mohammed Jabooob, and Ahmed Al-Ansi. Elevating e-government: Unleashing the power of AI and IoT for enhanced public services, 2024. Heliyon, vol. 10, no. 23 and <https://doi.org/10.1016/j.heliyon.2024.e40591>.
- Sarah A. Al-Qaseemi, Hajer A. Almulhim, Maria F. Almulhim, and Saqib Rasool Chaudhry. IoT architecture challenges and issues: Lack of standardization, 2017. FTC 2016 - Proceedings of Future Technologies Conference, pp. 731–738 and <https://doi.org/10.1109/FTC.2016.7821686>.
- Iqbal Alam, Kashif Sharif, Fan Li, Zohaib Latif, Md Monjurul Karim, Boubakr Nour, Sujit Biswas, and Yu Wang. IoT Virtualization: A Survey of Software Definition & Function Virtualization Techniques for Internet of Things, 2019. <http://arxiv.org/abs/1902.10910>.
- Mohammed M. Alani. Big data in cybersecurity: a survey of applications and future trends, 2021. Journal of Reliable Intelligent Environments, vol. 7, no. 2, pp. 85–114 and <https://doi.org/10.1007/s40860-020-00120-3>.

- Abdulaziz Aldaej, Imdad Ullah, Tariq Ahamed Ahanger, and Mohammed Atiquz-zaman. Ensemble technique of intrusion detection for IoT-edge platform, 2024. *Scientific Reports*, vol. 14, no. 1, pp. 11703 and <https://doi.org/10.1038/s41598-024-62435-y>.
- Mona Alduailij, Qazi Waqas Khan, Muhammad Tahir, Muhammad Sardaraz, Mai Alduailij, and Fazila Malik. Machine-Learning-Based DDoS Attack Detection Using Mutual Information and Random Forest Feature Importance Method, 2022. *Symmetry*, vol. 14, no. 6, and <https://www.mdpi.com/2073-8994/14/6/1095>.
- Ulaa AlHaddad, Abdullah Basuhail, Maher Khemakhem, Fathy Elbouraey Eassa, and Kamal Jambi. Ensemble Model Based on Hybrid Deep Learning for Intrusion Detection in Smart Grid Networks, 2023. *Sensors*, vol. 23, no. 17, pp. 7464 and <https://doi.org/10.3390/s23177464>.
- Maher Alharby. Evaluating machine learning approaches for multiple attack classification with improved computational efficiency in IoT networks, 2025. *Scientific Reports*, vol. 15, no. 1, and <https://doi.org/10.1038/s41598-025-23711-7>.
- Ayesha Alharthi, Meera Alaryani, and Sanaa Kaddoura. A comparative study of machine learning and deep learning models in binary and multiclass classification for intrusion detection systems, 2025. *Array*, vol. 26, pp. 100406 and <https://doi.org/10.1016/j.array.2025.100406>.
- Babar Ali, Muhammad Adeel Pasha, Saif Ul Islam, Houbing Song, and Rajkumar Buyya. A Volunteer-Supported Fog Computing Environment for Delay-Sensitive IoT Applications, 2021. *IEEE Internet of Things Journal*, vol. 8, no. 5, pp. 3822–3830 and <https://doi.org/10.1109/JIOT.2020.3024823>.
- Omar Mohammed Amin Ali, Rebin Abdulkareem Hamaamin, Barzan Jalal Youns, and Shahab Wahhab Kareem. Innovative Machine Learning Strategies for DDoS Detection: A Review, 2024. *UHD Journal of Science and Technology*, vol. 8, no. 2, pp. 38–49 and <https://doi.org/10.21928/uhdjst.v8n2y2024.pp38-49>.
- Shan Ali and Yuancheng Li. Learning Multilevel Auto-Encoders for DDoS Attack Detection in Smart Grid Network, 2019. *IEEE Access*, vol. 7, pp. 108647–108659 and <https://doi.org/10.1109/ACCESS.2019.2933304>.
- Mohammed Aljebreen, Hanan Abdullah Mengash, Munya A. Arasi, Sumayh S. Aljameel, Ahmed S. Salama, and Manar Ahmed Hamza. Enhancing DDoS Attack Detection Using Snake Optimizer With Ensemble Learning on Internet of Things Environment, 2023. *IEEE Access*, vol. 11, pp. 104745–104753 and <https://doi.org/10.1109/ACCESS.2023.3318316>.
- Ahamed Aljuhani. Machine Learning Approaches for Combating Distributed Denial of Service Attacks in Modern Networking Environments, 2021. *IEEE Access*, vol. 9, pp. 42236–42264 and <https://doi.org/10.1109/ACCESS.2021.3062909>.

- Zeina Allaw, Ola Zein, and Abdel Mehzen Ahmad. Cross-Layer Security for 5G/6G Network Slices: An SDN, NFV, and AI-Based Hybrid Framework, 2025. *Sensors*, vol. 25, no. 11 and <https://doi.org/10.3390/s25113335>.
- Nasr Almurisi and Srinivasulu Tadisetty. Cloud-based virtualization environment for IoT-based WSN: solutions, approaches and challenges, 2022. *Journal of Ambient Intelligence and Humanized Computing* and <https://doi.org/10.1007/s12652-021-03515-z>.
- Reham Almutairi, Giacomo Bergami, and Graham Morgan. Advancements and Challenges in IoT Simulators: A Comprehensive Review, 2024. *Sensors*, vol. 24, no. 5 and <https://doi.org/10.3390/s24051511>.
- Sara Amaouche, AzidineGuezzaz, Said Benkirane, and MouradeAzrour. IDS-XGbFS: a smart intrusion detection system using XGboostwith recent feature selection for VANET safety, 2024. *Cluster Computing*, vol. 27, no. 3, pp. 3521–3535 and <https://doi.org/10.1007/s10586-023-04157-w>.
- M Anirudh, S Arul Thileeban, and Daniel Jeswin Nallathambi. Use of honeypots for mitigating DoS attacks targeted on IoT networks, 2017. 2017 International Conference on Computer, Communication and Signal Processing (ICCCSP), pp. 1–4 and <https://doi.org/10.1109/ICCCSP.2017.7944057>.
- Liliana Antão, Rui Pinto, João Reis, and Gil Gonçalves. Requirements for testing and validating the industrial internet of things, 2018. *Proceedings - 2018 IEEE 11th International Conference on Software Testing, Verification and Validation Workshops, ICSTW 2018*, pp. 110–115 and <https://doi.org/10.1109/ICSTW.2018.00036>.
- M. Anuradha, T. Jayasankar, N. B. Prakash, Mohamed Yacin Sikkandar, G. R. Hemalakshmi, C. Bharatiraja, and A. Sagai Francis Britto. IoT enabled cancer prediction system to enhance the authentication and security using cloud computing, 2021. *Microprocessors and Microsystems*, vol. 80 and <https://doi.org/10.1016/j.micpro.2020.103301>.
- Alexandru Apostu, Silviu Gheorghe, Andrei Hiji, Andrei Pătras, Cristian Rusu, Nicolae Cleju, Radu Tudor Ionescu, and Paul Irofti. Detecting and Mitigating DDoS Attacks with AI: A Survey, 2018. *ACM Computing Surveys*, vol. 37, no. 111, pp. 35 and <https://www.marketsandmarkets.com/Market-Reports/ddos-protection-mitigation-market-111952874.html>.
- Giovanni Apruzzese, Pavel Laskov, Edgardo Montes De Oca, Wissam Mallouli, Luis Brdalo Rapa, Athanasios Vasileios Grammatopoulos, and Fabio Di Franco. The Role of Machine Learning in Cybersecurity, 2023. *Digital Threats: Research and Practice*, vol. 4, no. 1 and <https://doi.org/10.1145/3545574>.

- Juliana Arevalo-Herrera, Jorge Camargo Mendoza, Jose Ignacio Martínez Torre, Tatiana Zona-Ortiz, and Juan M. Ramirez. Assessing SDN Controller Vulnerabilities: A Survey on Attack Typologies, Detection Mechanisms, Controller Selection, and Dataset Application in Machine Learning, 2025. *Wireless Personal Communications*, vol. 140, no. 1, pp. 739–775 and <https://doi.org/10.1007/s11277-025-11748-w>.
- K. S. Arikumar, K. Tamilarasi, Sahaya Beni Prathiba, M. M. Venkata Chalapathi, Rajalakshmi Shenbaga Moorthy, and A. Deepak Kumar. The Role of Machine Learning in IoT: A Survey, 2022. 3rd International Conference on Smart Electronics and Communication, ICOSEC 2022 - Proceedings, pp. 451–457 and <https://doi.org/10.1109/ICOSEC54921.2022.9952042>.
- Muhammad Asad, Muhammad Asim, Talha Javed, Mirza O Beg, Hasan Mujtaba, and Sohail Abbas. DeepDetect: Detection of Distributed Denial of Service Attacks Using Deep Learning, 2020. *The Computer Journal*, vol. 63, no. 7, pp. 983–994 and <https://doi.org/10.1093/comjnl/bxz064>.
- Teerawut Banchuen, Kiattikun Kawila, and Kultida Rojviboonchai. An SDN framework for video conference in inter-domain network, 2018. *International Conference on Advanced Communication Technology, ICACT*, vol. 2018-February, pp. 600–605 and <https://doi.org/10.23919/ICACT.2018.8323848>.
- D. Bastos, M. Shackleton, and F. El-Moussa. Internet of things: A survey of technologies and security risks in smart home and city environments, 2018. *IET Conference Publications*, vol. 2018, no. CP740 and <https://doi.org/10.1049/cp.2018.0030>.
- Kamal Bella, Azidine Guezaz, Said Benkirane, Mourade Azrour, Yasser Fouad, Mbadiwe S. Benyeogor, and Nisreen Innab. An efficient intrusion detection system for IoT security using CNN decision forest, 2024. *PeerJ Computer Science*, vol. 10, pp. e2290 and <https://doi.org/10.7717/peerj-cs.2290>.
- Wala Ben Rhouma and Haythem Hayouni. Deep Learning-Driven Intrusion Detection System for Distributed Denial of Service Mitigation, 2025. *Indonesian Journal of Electrical Engineering and Informatics (IJEI)*, vol. 13, no. 2 and <https://doi.org/10.52549/ijeai.v13i2.6013>.
- Zaheed Ahmed Bhuiyan, Salekul Islam, Md Motaharul Islam, A. B. M. Ahasan Ullah, Farha Naz, and Mohammad Shahriar Rahman. On the (in)Security of the Control Plane of SDN Architecture: A Survey, 2023. *IEEE Access*, vol. 11, pp. 91550–91582 and <https://doi.org/10.1109/ACCESS.2023.3307467>.
- Narayanavadiwoo Gopinathan Bhuvaneswari Amma and Selvakumar Subramanian. VCDDeepFL: Vector Convolutional Deep Feature Learning Approach for Identification of Known and Unknown Denial of Service Attacks, 2018. *TENCON 2018 - 2018 IEEE Region 10 Conference*, pp. 0640–0645 and <https://doi.org/10.1109/TENCON.2018.8650225>.

- Adam Borys, Abu Kamruzzaman, Hasnain Nizam Thakur, Joseph C. Brickley, Md L. Ali, and Kutub Thakur. An Evaluation of IoT DDoS Cryptojacking Malware and Mirai Botnet, 2022. 2022 IEEE World AI IoT Congress, AIIoT 2022, pp. 725–729 and <https://doi.org/10.1109/AIIoT54504.2022.9817163>.
- Pat Bosshart, Dan Daly, Glen Gibb, Martin Izzard, Nick Mckeown, Jennifer Rexford, Cole Schlesinger, Dan Talayco, Amin Vahdat, George Varghese, and David Walker. P4: Programming Protocol-Independent Packet Processors, 2020.
- Hamid Bostani and Mansour Sheikhan. Hybrid of binary gravitational search algorithm and mutual information for feature selection in intrusion detection systems, 2017. *Soft Computing*, vol. 21, no. 9, pp. 2307–2324 and <https://doi.org/10.1007/s00500-015-1942-8>.
- Andrew T Campbell, Kazuho Miki, Hitachi Limited, and Japan John Vicente. OPEN SIGNALING FOR ATM, INTERNET AND MOBILE NETWORKS (OPENSIG'98), 1998. <http://comet.columbia.edu/opensig/activities/opensig98.html>.
- Daniele Canavese, Luca Mannella, Leonardo Regano, and Cataldo Basile. Security at the Edge for Resource-Limited IoT Devices, 2024. *Sensors*, vol. 24, no. 2, pp. 590 and <https://doi.org/10.3390/s24020590>.
- Martín Casado, Michael Freedman J, Justin Pettit, Jianying Luo, Nick McKeown, and Scott Shenker. Ethane: Taking control of the enterprise., 2007. Association for Computing Machinery, pp. 421.
- Nadia Chaabouni, Mohamed Mosbah, Akka Zemmari, Cyrille Sauvignac, and Parvez Faruki. Network Intrusion Detection for IoT Security Based on Learning Techniques, 2019a. *IEEE Communications Surveys and Tutorials*, vol. 21, no. 3, pp. 2671–2701 and <https://doi.org/10.1109/COMST.2019.2896380>.
- Nadia Chaabouni, Mohamed Mosbah, Akka Zemmari, Cyrille Sauvignac, and Parvez Faruki. Network Intrusion Detection for IoT Security Based on Learning Techniques, 2019b. *IEEE Communications Surveys and Tutorials*, vol. 21, no. 3, pp. 2671–2701 and <https://doi.org/10.1109/COMST.2019.2896380>.
- Victor Chang, Yen Hung Kuo, and Muthu Ramachandran. Cloud computing adoption framework: A security framework for business clouds, 2016. *Future Generation Computer Systems*, vol. 57, pp. 24–41 and <https://doi.org/10.1016/j.future.2015.09.031>.
- Rajat Chaudhary, Gagangeet Singh Aujla, Neeraj Kumar, and Pushpinder Kaur Chouhan. A comprehensive survey on software-defined networking for smart communities, 2022. *International Journal of Communication Systems* and <https://doi.org/10.1002/dac.5296>.

- Jienan Chen, Siyu Chen, Qi Wang, Bin Cao, Gang Feng, and Jianhao Hu. IRAF: A Deep Reinforcement Learning Approach for Collaborative Mobile Edge Computing IoT Networks, 2019. *IEEE Internet of Things Journal*, vol. 6, no. 4, pp. 7011–7024 and <https://doi.org/10.1109/JIOT.2019.2913162>.
- Lei Chen, Zhihao Wang, Ru Huo, and Tao Huang. An Adversarial DBN-LSTM Method for Detecting and Defending against DDoS Attacks in SDN Environments, 2023. *Algorithms*, vol. 16, no. 4, pp. 197 and <https://doi.org/10.3390/a16040197>.
- Sarika Choudhary and Nishtha Kesswani. Analysis of KDD-Cup'99, NSL-KDD and UNSW-NB15 Datasets using Deep Learning in IoT, 2020. *Procedia Computer Science*, vol. 167, pp. 1561–1573 and <https://doi.org/10.1016/j.procs.2020.03.367>.
- Abdullah Emir Cil, Kazim Yildiz, and Ali Buldu. Detection of DDoS attacks with feed forward based deep neural network model, 2021. *Expert Systems with Applications*, vol. 169, pp. 114520 and <https://doi.org/10.1016/j.eswa.2020.114520>.
- Roger William Coêlho, Ronan Assumpção Silva, Luciana Andréia Fondazzi Martimiano, and Elvio João Leonardo. IoT and 5G Networks: A Discussion of SDN, NFV and Information Security, 2024. *Journal of the Brazilian Computer Society*, vol. 30, no. 1, pp. 212–227 and <https://doi.org/10.5753/jbcs.2024.3021>.
- Yunhe Cui, Lianshan Yan, Saifei Li, Huanlai Xing, Wei Pan, Jian Zhu, and Xiaoyang Zheng. SD-Anti-DDoS: Fast and efficient DDoS defense in software-defined networks, 2016. *Journal of Network and Computer Applications*, vol. 68, pp. 65–79 and <https://doi.org/10.1016/j.jnca.2016.04.005>.
- Sajjad Dadkhah, Euclides Carlos Pinto Neto, Raphael Ferreira, Reginald Chukwuka Molokwu, Somayeh Sadeghi, and Ali A. Ghorbani. CICIoMT2024: A benchmark dataset for multi-protocol security assessment in IoMT, 2024. *Internet of Things*, vol. 28, pp. 101351 and <https://doi.org/10.1016/j.iot.2024.101351>.
- Marcos V.O. de Assis, Luiz F. Carvalho, Joel J.P.C. Rodrigues, Jaime Lloret, and Mario L. Proença Jr. Near real-time security system applied to SDN environments in IoT networks using convolutional neural network, 2020. *Computers & Electrical Engineering*, vol. 86, pp. 106738 and <https://doi.org/10.1016/j.compeleceng.2020.106738>.
- Miguel Borges De Freitas, Pedro Quiterio, Luis Rosa, Tiago Cruz, and Paulo Simoes. SDN-assisted containerized security and monitoring components, 2020. *Proceedings of IEEE/IFIP Network Operations and Management Symposium 2020: Management in the Age of Softwarization and Artificial Intelligence, NOMS 2020* and <https://doi.org/10.1109/NOMS47738.2020.9110291>.
- Sedef Demirci and Seref Sagiroglu. Optimal placement of virtual network functions in software defined networks: A survey, 2019. *Journal of Network and Computer Applications*, vol. 147 and <https://doi.org/10.1016/j.jnca.2019.102424>.

- Lubna Luxmi Dhirani, Eddie Armstrong, and Thomas Newe. Industrial iot, cyber threats, and standards landscape: Evaluation and roadmap, 2021. *Sensors*, vol. 21, no. 11 and <https://doi.org/10.3390/s21113901>.
- José Dias, Pedro Pinto, Ricardo Santos, and Silvestre Malta. 5G Network Slicing: Security Challenges, Attack Vectors, and Mitigation Approaches, 2025. *Sensors*, vol. 25, no. 13 and <https://doi.org/10.3390/s25133940>.
- Abebe Abeshu Diro and Naveen Chilamkurti. Distributed attack detection scheme using deep learning approach for Internet of Things, 2018. *Future Generation Computer Systems*, vol. 82, pp. 761–768 and <https://doi.org/10.1016/j.future.2017.08.043>.
- Ionut Catalin Donca, Ovidiu Petru Stan, Marius Misaros, Anca Stan, and Liviu Miclea. Comprehensive Security for IoT Devices with Kubernetes and Raspberry Pi Cluster, 2024. *Electronics (Switzerland)*, vol. 13, no. 9 and <https://doi.org/10.3390/electronics13091613>.
- Ratnesh Kumar Dubey, Nidhi Dandotiya, Aravendra Sharma, Shubha Mishra, and Shashi Kant Gupta. Cyber attack Detection Using Machine Learning Techniques, 2023. 3rd IEEE International Conference on ICT in Business Industry and Government, ICTBIG 2023 and <https://doi.org/10.1109/ICTBIG59752.2023.10456080>.
- Rajendra Kumar Dwivedi, Sonali Pandey, and Rakesh Kumar. A study on Machine Learning Approaches for Outlier Detection in Wireless Sensor Network, 2018. 2018 8th International Conference on Cloud Computing, Data Science & Engineering (Confluence) : 11-12 Jan. 2018.
- Ahmad Hani El Fawal, Ali Mansour, Mohammad Ammad Uddin, and Abbass Nasser. Securing IoT Networks from DDoS Attacks Using a Temporary Dynamic IP Strategy, 2024. *Sensors*, vol. 24, no. 13 and <https://doi.org/10.3390/s24134287>.
- Mohammed El-Hajj, Ahmad Fadlallah, Maroun Chamoun, and Ahmed Serhrouchni. A survey of internet of things (IoT) authentication schemes, 2019. *Sensors (Switzerland)*, vol. 19, no. 5 and <https://doi.org/10.3390/s19051141>.
- ETSI. Network Functions Virtualisation (NFV), 2013. GS NFV 002 (v. 1.1.1).
- Rasoul Farahi. A comprehensive overview of load balancing methods in software-defined networks, 2025. *Discover Internet of Things*, vol. 5, no. 1, pp. 6 and <https://link.springer.com/10.1007/s43926-025-00098-5>.
- Cheng Feng and Pengwei Tian. Time Series Anomaly Detection for Cyber-physical Systems via Neural System Identification and Bayesian Filtering, 2021. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 2858–2867 and <https://doi.org/10.1145/3447548.3467137>.

- Wendi Feng, Zehua Guo, Chuanchang Liu, Yueming Zheng, Meng Wang, Bo Cheng, and Junliang Chen. BAGUETTE: Towards a Secure and Cost-effective Switch Upgrade in Hybrid Software-Defined Networks, 2020. 2020 IEEE International Conference on Communications (ICC), vol. 30, no. 01, pp. 1–6.
- Mohamed Amine Ferrag, Othmane Friha, Djallel Hamouda, Leandros Maglaras, and Helge Janicke. Edge-IIoTset: A New Comprehensive Realistic Cyber Security Dataset of IoT and IIoT Applications for Centralized and Federated Learning, 2022. IEEE Access, vol. 10, pp. 40281–40306 and <https://doi.org/10.1109/ACCESS.2022.3165809>.
- Giovanni Fioravanti, Mattia Giovanni Spina, and Floriano De Rango. Entropy based DDoS Detection in Software Defined Networks, 2023. 2023 IEEE 20th Consumer Communications & Networking Conference (CCNC), pp. 636–639 and <https://doi.org/10.1109/CCNC51644.2023.10059949>.
- Shu Fu, Qilin Fan, Yujie Tang, Haijun Zhang, Xin Jian, and Xiaoping Zeng. Cooperative Computing in Integrated Blockchain-Based Internet of Things, 2020. IEEE Internet of Things Journal, vol. 7, no. 3, pp. 1603–1612 and <https://doi.org/10.1109/JIOT.2019.2948144>.
- Xiuwen Fu, Dingyi Zheng, Xiangwei Liu, Liudong Xing, and Rui Peng. Systematic review and future perspectives on cascading failures in Internet of Things: Modeling and optimization, 2025. Reliability Engineering and System Safety, vol. 254 and <https://doi.org/10.1016/j.ress.2024.110582>.
- Angelo Furfaro, Pasquale Pace, and Andrea Parise. Facing DDoS bandwidth flooding attacks, 2020. Simulation Modelling Practice and Theory, vol. 98 and <https://doi.org/10.1016/j.simpat.2019.101984>.
- Mattias Gabriel, Luigi Iannone, and Benoit Donnet. Lisp Mapping System as DoS Amplification Vector, 2021. IEEE Networking Letters, vol. 3, no. 1, pp. 36–39 and <https://doi.org/10.1109/LNET.2021.3050814>.
- Metehan Gelgi, Yueting Guan, Sanjay Arunachala, Maddi Samba Siva Rao, and Nicola Dragoni. Systematic Literature Review of IoT Botnet DDOS Attacks and Evaluation of Detection Techniques, 2024. Sensors, vol. 24, no. 11 and <https://doi.org/10.3390/s24113571>.
- Sukhpal Singh Gill, Shreshth Tuli, Minxian Xu, Inderpreet Singh, Karan Vijay Singh, Dominic Lindsay, Shikhar Tuli, Daria Smirnova, Manmeet Singh, Udit Jain, Haris Pervaiz, Bhanu Sehgal, Sukhwinder Singh Kaila, Sanjay Misra, Mohammad Sadegh Aslanpour, Harshit Mehta, Vlado Stankovski, and Peter Garraghan. Transformative effects of IoT, Blockchain and Artificial Intelligence on cloud computing: Evolution, vision, trends and open challenges, 2019. Internet of Things (Netherlands), vol. 8 and <https://doi.org/10.1016/j.iot.2019.100118>.

- Lewis Golightly, Paolo Modesti, Rémi Garcia, and Victor Chang. Securing distributed systems: A survey on access control techniques for cloud, blockchain, IoT and SDN, 2023. *Cyber Security and Applications*, vol. 1 and <https://doi.org/10.1016/j.csa.2023.100015>.
- Shui Qing Gong, Jing Chen, Cong Hui Huang, and Qing Chao Zhu. Trust-aware secure virtual network embedding algorithm, 2015. *Tongxin Xuebao/Journal on Communications*, vol. 36, no. 11, pp. 1–10 and <https://doi.org/10.11959/j.issn.1000-436x.2015272>.
- Paul Göransson and Chuck Black. Software defined networks: a comprehensive approach, 2016. Morgan Kaufmann, pp. 1–353.
- Afrah Gueriani, Hamza Kheddar, and Ahmed Cherif Mazari. Adaptive Cyber-Attack Detection in IIoT Using Attention-Based LSTM-CNN Models, 2025. <https://doi.org/10.1109/ICTIS62692.2024.10894509>.
- Brij B Gupta, Amrita Dahiya, and Elhadj Benkhelifa. SDN and NFV: A New Dimension To Virtualization, 2024. World Scientific Publishing Co Pte Ltd, vol. 02, pp. 12–292 and <https://doi.org/10.1142/12794>.
- Sachin Gupta. IoT based Smart Home Automation using NodeMCU, 2021. *International Journal for Research in Applied Science and Engineering Technology*, vol. 9, no. 9, pp. 1179–1188 and <https://doi.org/10.22214/ijraset.2021.38165>.
- Pooyan Habibi, Mohammad Farhoudi, Sepehr Kazemian, Siavash Khorsandi, and Alberto Leon-Garcia. Fog Computing: A Comprehensive Architectural Survey, 2020. *IEEE Access*, vol. 8, pp. 69105–69133 and <https://doi.org/10.1109/ACCESS.2020.2983253>.
- Khan Mohammad Habibullah, Gregory Gay, and Jennifer Horkoff. Non-functional requirements for machine learning: understanding current use and challenges among practitioners, 2023. *Requirements Engineering*, vol. 28, no. 2, pp. 283–316 and <https://doi.org/10.1007/s00766-022-00395-3>.
- Qasim Hama Khurshid Hamamurad, Normal Mat Jusoh, and Uznir Ujang. Smart City Domains System Framework: Review of the Current Status, 2021. *Proceeding - 2021 IEEE 9th Conference on System, Process and Control, ICSPC 2021*, pp. 19–23 and <https://doi.org/10.1109/ICSPC53359.2021.9689187>.
- Ayyoob Hamza, Hassan Habibi Gharakheili, Theophilus A. Benson, and Vijay Sivaraman. Detecting Volumetric Attacks on IoT Devices via SDN-Based Monitoring of MUD Activity, 2019. *Proceedings of the 2019 ACM Symposium on SDN Research*, pp. 36–48 and <https://doi.org/10.1145/3314148.3314352>.

- Ayyoob Hamza, Dinesha Ranathunga, Hassan Habibi Gharakheili, Theophilus A. Benson, Matthew Roughan, and Vijay Sivaraman. Verifying and Monitoring IoTs Network Behavior Using MUD Profiles, 2022. *IEEE Transactions on Dependable and Secure Computing*, vol. 19, no. 1, pp. 1–18 and <https://doi.org/10.1109/TDSC.2020.2997898>.
- Bo Han, Vijay Gopalakrishnan, Lusheng Ji, and Seungjoon Lee. *Network_function_virtualization_Challenges_and_opportunities_for_innovations*, 2015.
- Zheng Hao, Zhaowen Lin, and Ran Li. A SDN/NFv security protection architecture with a function composition algorithm based on trie, 2018. *ACM International Conference Proceeding Series* and <https://doi.org/10.1145/3207677.3277992>.
- Mahmudul Hasan, Md. Milon Islam, Md Ishrak Islam Zarif, and M.M.A. Hashem. Attack and anomaly detection in IoT sensors in IoT sites using machine learning approaches, 2019. *Internet of Things*, vol. 7, pp. 100059 and <https://doi.org/10.1016/j.iot.2019.100059>.
- Hassan Hawilo, Manar Jammal, and Abdallah Shami. Network Function Virtualization-Aware Orchestrator for Service Function Chaining Placement in the Cloud, 2019. *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 3, pp. 643–655 and <https://doi.org/10.1109/JSAC.2019.2895226>.
- Jiawei He, Yejin Tan, Wangshu Guo, and Ming Xian. A Small Sample DDoS Attack Detection Method Based on Deep Transfer Learning, 2020. *2020 International Conference on Computer Communication and Network Security (CCNS)*, pp. 47–50 and <https://doi.org/10.1109/CCNS50731.2020.00019>.
- Arvin Hekmati, Jiahe Zhang, Tamoghna Sarkar, Nishant Jethwa, Eugenio Grippo, and Bhaskar Krishnamachari. Correlation-Aware Neural Networks for DDoS Attack Detection in IoT Systems, 2024. *IEEE/ACM Transactions on Networking*, vol. 32, no. 5, pp. 3929–3944 and <https://doi.org/10.1109/TNET.2024.3408675>.
- Almulla Hesham, Fragkiskos Sardis, Stan Wong, Toktam Mahmoodi, and Mallikarjun Tatipamula. A Simplified Network Access Control Design and Implementation for M2M Communication Using SDN, 2017. *2017 IEEE Wireless Communications and Networking Conference Workshops, WCNCW 2017* and <https://doi.org/10.1109/WCNCW.2017.7919082>.
- Alan Hevner and Samir Chatterjee. *Design Science Research in Information Systems*, 2010. pp. 9–22 and https://doi.org/10.1007/978-1-4419-5653-8_{ }2.
- Abdinasir Hirsi, Lukman Audah, and Adeb Salh. *SDN-DDoS Traffic Dataset*, 2024. *Universiti Tun Hussein Onn Malaysia*, vol. 1, no. 1.

- Vanlalruata Hnamte, Ashfaq Ahmad Najar, Hong Nhung-Nguyen, Jamal Hussain, and Manohar Naik Sugali. DDoS attack detection and mitigation using deep neural network in SDN environment, 2024. *Computers & Security*, vol. 138, pp. 103661 and <https://doi.org/10.1016/j.cose.2023.103661>.
- Md. Alamgir Hossain and Md. Saiful Islam. Ensuring network security with a robust intrusion detection system using ensemble-based machine learning, 2023. *Array*, vol. 19, pp. 100306 and <https://doi.org/10.1016/j.array.2023.100306>.
- Injung Hwang and Dongkun Shin. Application level network virtualization using selective connection, 2018. 2018 IEEE International Conference on Consumer Electronics, ICCE 2018, vol. 2018-January, pp. 1–2 and <https://doi.org/10.1109/ICCE.2018.8326277>.
- Muhammad Ishaq, Muhammad Hammad Afzal, Shehla Tahir, and Kifayat Ullah. A Compact Study of Recent Trends of Challenges and Opportunities in Integrating Internet of Things (IoT) and Cloud Computing, 2021. 2021 International Conference on Computing, Electronic and Electrical Engineering, ICE Cube 2021 - Proceedings and <https://doi.org/10.1109/ICECube53880.2021.9628191>.
- Uzair Javaid, Ang Kiang Siang, Muhammad Naveed Aman, and Biplab Sikdar. Mitigating IoT Device based DDoS Attacks using Blockchain, 2018. *Proceedings of the 1st Workshop on Cryptocurrencies and Blockchains for Distributed Systems*, pp. 71–76 and <https://doi.org/10.1145/3211933.3211946>.
- Abdul Rehman Javed, Muhammad Usman Sarwar, Mirza Omer Beg, Muhammad Asim, Thar Baker, and Hissam Tawfik. A collaborative healthcare framework for shared healthcare plan with ambient intelligence, 2020. *Human-centric Computing and Information Sciences*, vol. 10, no. 1 and <https://doi.org/10.1186/s13673-020-00245-7>.
- James A. Jerkins. Motivating a market or regulatory solution to IoT insecurity with the Mirai botnet code, 2017. 2017 IEEE 7th Annual Computing and Communication Workshop and Conference (CCWC), pp. 1–5 and <https://doi.org/10.1109/CCWC.2017.7868464>.
- Dingde Jiang, Zhihao Wang, Liuwei Huo, and Shaowei Xie. A Performance Measurement and Analysis Method for Software-Defined Networking of IoV, 2021. *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 6, pp. 3707–3719 and <https://doi.org/10.1109/TITS.2020.3029076>.
- Ebunayo Rachael Jimoh, Tinuke Omolewa Oladele, Oluwabunmi Titilope Oladele, Azeez Olawale Akinlolu, Mulikat Bola Akanbi, and Olasehinde Williams Olabanke. CNN-LSTM Model for Mitigation of DDoS Attacks in Software-Defined Networks, 2024. *International Conference on Science, Engineering and Business for Driving*

- Sustainable Development Goals, SEB4SDG 2024 and <https://doi.org/10.1109/SEB4SDG60871.2024.10630294>.
- Dhiraj Joshi, Nirmal Desai, Shyama Prosad Chowdhury, Wei-Han Lee, Luis Bathen, Shiqiang Wang, and Dinesh Verma. AI at the Edge: Challenges, Applications, and Directions, 2022. IoT for Defense and National Security, pp. 133–160 and <https://doi.org/10.1002/9781119892199.ch9>.
- Dimitrios Kafetzis, Spyridon Vassilaras, Georgios Vardoulas, and Iordanis Koutsopoulos. Software-Defined Networking Meets Software-Defined Radio in Mobile ad hoc Networks: State of the Art and Future Directions, 2022. IEEE Access, vol. 10, pp. 9989–10014 and <https://doi.org/10.1109/ACCESS.2022.3144072>.
- Kubra Kalkan, Levent Altay, Gurkan Gur, and Fatih Alagoz. JESS: Joint Entropy-Based DDoS Defense Scheme in SDN, 2018. IEEE Journal on Selected Areas in Communications, vol. 36, no. 10, pp. 2358–2372 and <https://doi.org/10.1109/JSAC.2018.2869997>.
- Amandeep Kaur, C. Rama Krishna, and Nilesh Vishwasrao Patil. A comprehensive review on Software-Defined Networking (SDN) and DDoS attacks: Ecosystem, taxonomy, traffic engineering, challenges and research directions, 2025. Computer Science Review, vol. 55 and <https://doi.org/10.1016/j.cosrev.2024.100692>.
- Sukhveer Kaur, Krishan Kumar, and Naveen Aggarwal. A review on P4-Programmable data planes: Architecture, research efforts, and future directions, 2021. Computer Communications, vol. 170, pp. 109–129 and <https://doi.org/10.1016/j.comcom.2021.01.027>.
- Shaden A. M. Khalifa, Nizar Elias, Mohamed A. Farag, Lei Chen, Aamer Saeed, Mohamed-Elamir F. Hegazy, Moustafa S. Moustafa, Aida Abd El-Wahed, Saleh M. Al-Mousawi, Syed G. Musharraf, Fang-Rong Chang, Arihiro Iwasaki, Kiyotake Suenaga, Muaaz Alajlani, Ulf Göransson, and Hesham R. El-Seedi. Marine Natural Products: A Source of Novel Anticancer Drugs, 2019. Marine Drugs, vol. 17, no. 9, pp. 491 and <https://doi.org/10.3390/md17090491>.
- Mohd Zain Khan and Mohammad Ubaidullah Bokhari. Comparative Study of Detection of IoT Attacks using Machine Learning Techniques, 2023. 10th International Conference on Computing for Sustainable Global Development (INDIACom).
- Sang Il Kim and Hwa Sung Kim. A VNF Placement Method based on VNF Characteristics, 2021. International Conference on Information Networking, vol. 2021-January, pp. 864–869 and <https://doi.org/10.1109/ICIN50884.2021.9334022>.
- Ye-Eun Kim, Yea-Sul Kim, and Hwankuk Kim. Effective Feature Selection Methods to Detect IoT DDoS Attack in 5G Core Network, 2022. Sensors, vol. 22, no. 10, pp. 3819 and <https://doi.org/10.3390/s22103819>.

- Nickolaos Koroniotis, Nour Moustafa, Elena Sitnikova, and Benjamin Turnbull. Towards the development of realistic botnet dataset in the Internet of Things for network forensic analytics: Bot-IoT dataset, 2019. *Future Generation Computer Systems*, vol. 100, pp. 779–796 and <https://doi.org/10.1016/j.future.2019.05.041>.
- Jakub Kufel, Katarzyna Bargiel-Łączek, Szymon Kocot, Maciej Koźlik, Wiktoria Bartnikowska, Michał Janik, Łukasz Czogalik, Piotr Dudek, Mikołaj Magiera, Anna Lis, Iga Paszkiewicz, Zbigniew Nawrat, Maciej Cebula, and Katarzyna Gruszczyńska. What Is Machine Learning, Artificial Neural Networks and Deep Learning?—Examples of Practical Applications in Medicine, 2023. *Diagnostics*, vol. 13, no. 15 and <https://doi.org/10.3390/diagnostics13152582>.
- Prashant Kumar, Chitra Kushwaha, Dimple Sethi, Debjani Ghosh, Punit Gupta, and Ankit Vidyarthi. Investigating the performance of multivariate LSTM models to predict the occurrence of Distributed Denial of Service (DDoS) attack, 2025. *PLOS ONE*, vol. 20, no. 1, pp. e0313930 and <https://doi.org/10.1371/journal.pone.0313930>.
- Sanjeev Kumar, Adrian Guerrero, and Christina Navarro. Cyber Security Flood Attacks and Risk Assessment for Internet of Things (IoT) Distributed Systems, 2023. 2023 IEEE World AI IoT Congress, AIIoT 2023, pp. 392–397 and <https://doi.org/10.1109/AIIoT58121.2023.10174553>.
- Pooja Kumari and Ankit Kumar Jain. A comprehensive study of DDoS attacks over IoT network and their countermeasures, 2023a. *Computers and Security*, vol. 127 and <https://doi.org/10.1016/j.cose.2023.103096>.
- Pooja Kumari and Ankit Kumar Jain. A comprehensive study of DDoS attacks over IoT network and their countermeasures, 2023b. *Computers & Security*, vol. 127, pp. 103096 and <https://doi.org/10.1016/j.cose.2023.103096>.
- Kuruva Lakshmana, Rajesh Kaluri, Nagaraja Gundluru, Zamil S. Alzamil, Dharmendra Singh Rajput, Arfat Ahmad Khan, Mohd Anul Haq, and Ahmed Alhussen. A Review on Deep Learning Techniques for IoT Data, 2022. *Electronics (Switzerland)*, vol. 11, no. 10 and <https://doi.org/10.3390/electronics11101604>.
- Jake Lever, Martin Krzywinski, and Naomi Altman. Model selection and overfitting, 2016. *Nature Methods*, vol. 13, no. 9, pp. 703–704 and <https://doi.org/10.1038/nmeth.3968>.
- Cheng Li, Zhengrui Qin, Ed Novak, and Qun Li. Securing SDN Infrastructure of IoT-Fog Networks from MitM Attacks, 2017. *IEEE Internet of Things Journal*, vol. 4, no. 5, pp. 1156–1164 and <https://doi.org/10.1109/JIOT.2017.2685596>.
- Chunyan Li, Jiaji Wang, Shuihua Wang, and Yudong Zhang. A review of IoT applications in healthcare, 2024. *Neurocomputing*, vol. 565 and <https://doi.org/10.1016/j.neucom.2023.127017>.

- Huiling Li, Jun Wu, Hansong Xu, Gaolei Li, and Mohsen Guizani. Explainable Intelligence-Driven Defense Mechanism against Advanced Persistent Threats: A Joint Edge Game and AI Approach, 2022a. *IEEE Transactions on Dependable and Secure Computing*, pp. 1–1 and <https://doi.org/10.1109/TDSC.2021.3130944>.
- Junhong Li. Detection of DDoS Attacks Based on Dense Neural Networks, Autoencoders and Pearson Correlation Coefficient, 2020. Dalhousie University.
- Yuzhen Li, Renjie Li, Zhou Zhou, Jiang Guo, Wei Yang, Meijie Du, and Qingyun Liu. GraphDDoS: Effective DDoS Attack Detection Using Graph Neural Networks, 2022b. 2022 IEEE 25th International Conference on Computer Supported Cooperative Work in Design (CSCWD), pp. 1275–1280 and <https://doi.org/10.1109/CSCWD54268.2022.9776097>.
- Xiaoyu Liang and Taieb Znati. A Long Short-Term Memory Enabled Framework for DDoS Detection, 2019. 2019 IEEE Global Communications Conference (GLOBECOM), pp. 1–6 and <https://doi.org/10.1109/GLOBECOM38437.2019.9013450>.
- Jenn Wei Lin, Pethuru Raj Chelliah, Meng Chieh Hsu, and Jia Xin Hou. Efficient Fault-Tolerant Routing in IoT Wireless Sensor Networks Based on Bipartite-Flow Graph Modeling, 2019. *IEEE Access*, vol. 7, pp. 14022–14034 and <https://doi.org/10.1109/ACCESS.2019.2894002>.
- Weiwei Lin, Chennian Xiong, Wentai Wu, Fang Shi, Keqin Li, and Minxian Xu. Performance Interference of Virtual Machines: A Survey, 2023. *ACM Computing Surveys*, vol. 55, no. 12 and <https://doi.org/10.1145/3573009>.
- Siquan Liu, Chuangchuang Zhang, Hongyong Yang, Guanghai Cui, Fuliang Li, and Xingwei Wang. Delay guaranteed SFC placement with VNF parallelization in multidomain IoT networks, 2025. *Scientific Reports*, vol. 15, no. 1 and <https://doi.org/10.1038/s41598-025-23407-y>.
- Xiangyu Liu, Yi Han, and Yanhui Du. IoT Device Identification Using Directional Packet Length Sequences and 1D-CNN, 2022. *Sensors*, vol. 22, no. 21, pp. 8337 and <https://doi.org/10.3390/s22218337>.
- Manuel López Martín. Novel applications of Machine Learning to Network Traffic Analysis and Prediction, 2019. Universidad de Valladolid and <http://uvadoc.uva.es/handle/10324/39457>.
- Yang Lu and Li Da Xu. Internet of things (IoT) cybersecurity research: A review of current research topics, 2019. *IEEE Internet of Things Journal*, vol. 6, no. 2, pp. 2103–2115 and <https://doi.org/10.1109/JIOT.2018.2869847>.
- Subhi R M Zeebaree, Rizgar R Zebari, Karwan Jacksi, and Dathar Abas Hasan. Security Approaches For Integrated Enterprise Systems Performance: A Review, 2019.

- INTERNATIONAL JOURNAL OF SCIENTIFIC & TECHNOLOGY RESEARCH, vol. 8 and www.ijstr.org.
- Alex Mavromatis, Carlos Colman-Meixner, Aloizio P. Silva, Xenofon Vasilakos, Reza Nejabati, and Dimitra Simeonidou. A Software-Defined IoT Device Management Framework for Edge and Cloud Computing, 2020. IEEE Internet of Things Journal, vol. 7, no. 3, pp. 1718–1735 and <https://doi.org/10.1109/JIOT.2019.2949629>.
- Bertha Mazon-Olivo and Alberto Pan. Internet of Things: State-of-the-art, Computing Paradigms and Reference Architectures; Internet of Things: State-of-the-art, Computing Paradigms and Reference Architectures, 2022. IEEE Latin America Transactions, vol. 20.
- Sri Harsha Mekala, Zubair Baig, Adnan Anwar, and Sherali Zeadally. Cybersecurity for Industrial IoT (IIoT): Threats, countermeasures, challenges and future directions, 2023. Computer Communications, vol. 208, pp. 294–320 and <https://doi.org/10.1016/j.comcom.2023.06.020>.
- Samuel Mergendahl, Devkishen Sisodia, Jun Li, and Hasan Cam. Source-End DDoS Defense in IoT Environments, 2017. Proceedings of the 2017 Workshop on Internet of Things Security and Privacy, pp. 63–64 and <https://doi.org/10.1145/3139937.3139954>.
- Mheni Merzouki, Charif Mahmoudi, Robert Bohn, and Cihan Tunc. Security Automation for Cloud-Based IoT Platforms, 2019. 2019 International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCoM) and IEEE Smart Data (SmartData), pp. 1185–1191 and <https://ieeexplore.ieee.org/document/8875342/>.
- Rashid Mijumbi, Joan Serrat, Juan Luis Gorricho, Niels Bouten, Filip De Turck, and Raouf Boutaba. Network function virtualization: State-of-the-art and research challenges, 2016. IEEE Communications Surveys and Tutorials, vol. 18, no. 1, pp. 236–262 and <https://doi.org/10.1109/COMST.2015.2477041>.
- Ashish Mishra, Saurabh Sharma, and Archana Pandey. A Review on DDOS Attack, TCP Flood Attack in Cloud Environment, 2020. SSRN Electronic Journal and <https://doi.org/10.2139/ssrn.3565043>.
- Nivedita Mishra and Sharnil Pandya. Internet of Things Applications, Security Challenges, Attacks, Intrusion Detection, and Future Visions: A Systematic Review, 2021. IEEE Access, vol. 9, pp. 59353–59377 and <https://doi.org/10.1109/ACCESS.2021.3073408>.
- Sandipan Rakeshkumar Mishra, Bharanidharan Shanmugam, Kheng Cher Yeo, and Suresh Thennadil. SDN-Enabled IoT Security Frameworks—A Review of Existing

- Challenges, 2025. Technologies, vol. 13, no. 3, pp. 121 and <https://doi.org/10.3390/technologies13030121>.
- Sudip Misra, P. Venkata Krishna, Harshit Agarwal, Antriksh Saxena, and Mohammad S. Obaidat. A Learning Automata Based Solution for Preventing Distributed Denial of Service in Internet of Things, 2011. 2011 International Conference on Internet of Things and 4th International Conference on Cyber, Physical and Social Computing, pp. 114–122 and <https://doi.org/10.1109/iThings/CPSCoM.2011.84>.
- Meenakshi Mittal, Krishan Kumar, and Sunny Behal. Deep learning approaches for detecting DDoS attacks: a systematic review, 2023. Soft Computing, vol. 27, no. 18, pp. 13039–13075 and <https://doi.org/10.1007/s00500-021-06608-1>.
- Yanmeng Mo, Huige Li, Dongsheng Wang, and Gaqiong Liu. An intrusion detection system based on convolution neural network, 2024. PeerJ Computer Science, vol. 10, pp. e2152 and <https://doi.org/10.7717/peerj-cs.2152>.
- K. Venkata Murali Mohan, Sarangam Kodati, and V. Krishna. Securing SDN Enabled IoT Scenario Infrastructure of Fog Networks From Attacks, 2022. Proceedings of the 2nd International Conference on Artificial Intelligence and Smart Energy, ICAIS 2022, pp. 1239–1243 and <https://doi.org/10.1109/ICAIS53314.2022.9742727>.
- Seyed Mohammad Mousavi and Marc St-Hilaire. Early detection of DDoS attacks against SDN controllers, 2015. 2015 International Conference on Computing, Networking and Communications (ICNC), pp. 77–81 and <https://doi.org/10.1109/ICNC.2015.7069319>.
- Xiaoshao Mu and Maxwell Fordjour Antwi-Afari. The applications of Internet of Things (IoT) in industrial management: a science mapping review, 2024. International Journal of Production Research, vol. 62, no. 5, pp. 1928–1952 and <https://doi.org/10.1080/00207543.2023.2290229>.
- Manoj Muniswamaiah, Tilak Agerwala, and Charles C. Tappert. Fog Computing and the Internet of Things (IoT): A Review, 2021. Proceedings - 2021 8th IEEE International Conference on Cyber Security and Cloud Computing and 2021 7th IEEE International Conference on Edge Computing and Scalable Cloud, CSCloud-EdgeCom 2021, pp. 10–12 and <https://doi.org/10.1109/CSCloud-EdgeCom52276.2021.00012>.
- Ashfaq Ahmad Najar and S. Manohar Naik. Cyber-Secure SDN: A CNN-Based Approach for Efficient Detection and Mitigation of DDoS attacks, 2024. Computers & Security, vol. 139, pp. 103716 and <https://doi.org/10.1016/j.cose.2024.103716>.
- KG Raghavendra Narayan, Srijanee Mookherji, Vanga Odelu, Rajendra Prasath, Anish Chand Turlapaty, and Ashok Kumar Das. IIDS: Design of Intelligent Intrusion Detection System for Internet-of-Things Applications, 2023.

- Mohd Nazeer, Areej Alasiry, Mohammed Qayyum, Vemana Karunakar Madhan, Gouri Patil, and Pulipati Srilatha. Enhancing Cyber Security in Autonomous Vehicles: A Hybrid XG Boost-Deep Learning Approach for Intrusion Detection in the CAN Bus, 2024. *Journal European des Systemes Automatisés*, vol. 57, no. 5, pp. 1295–1304 and <https://doi.org/10.18280/jesa.570505>.
- Ahsan Nazir, Jingsha He, Nafei Zhu, Saima Siraj Qureshi, Siraj Uddin Qureshi, Faheem Ullah, Ahsan Wajahat, and Muhammad Salman Pathan. A deep learning-based novel hybrid CNN-LSTM architecture for efficient detection of threats in the IoT ecosystem, 2024. *Ain Shams Engineering Journal*, vol. 15, no. 7, pp. 102777 and <https://doi.org/10.1016/j.asej.2024.102777>.
- Euclides Carlos Pinto Neto, Sajjad Dadkhah, Raphael Ferreira, Alireza Zohourian, Rongxing Lu, and Ali A. Ghorbani. CICIOT2023: A Real-Time Dataset and Benchmark for Large-Scale Attacks in IoT Environment, 2023. *Sensors*, vol. 23, no. 13, pp. 5941 and <https://doi.org/10.3390/s23135941>.
- Juniper Networks. Learn About Network Functions Virtualization, 2022. <http://www.docker.com>.
- Stefania Loredana Nita and Marius Iulian Mihailescu. On Artificial Neural Networks used in Cloud Computing Security - A Survey, 2018. *ECAI 2018 - International Conference – 10th Edition*, vol. 31, pp. 91–97.
- Beny Nugraha and Rathan Narasimha Murthy. Deep Learning-based Slow DDoS Attack Detection in SDN-based Networks, 2020. *2020 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN)*, pp. 51–56 and <https://doi.org/10.1109/NFV-SDN50289.2020.9289894>.
- Bruno Astuto A. Nunes, Marc Mendonca, Xuan Nam Nguyen, Katia Obraczka, and Thierry Turetletti. A survey of software-defined networking: Past, present, and future of programmable networks, 2014. *IEEE Communications Surveys and Tutorials*, vol. 16, no. 3, pp. 1617–1634 and <https://doi.org/10.1109/SURV.2014.012214.00180>.
- Hamed Haddad Pajouh, Gholam Hossein Dastghaibiyfard, and Sattar Hashemi. Two-tier network anomaly detection model: a machine learning approach, 2017. *Journal of Intelligent Information Systems*, vol. 48, no. 1, pp. 61–74 and <https://doi.org/10.1007/s10844-015-0388-x>.
- Ricardo Parizotto, Bruno Loureiro Coelho, Diego Cardoso Nunes, Israat Haque, and Alberto Schaeffer-Filho. Offloading Machine Learning to Programmable Data Planes: A Systematic Survey, 2023. *ACM Computing Surveys*, vol. 56, no. 1 and <https://doi.org/10.1145/3605153>.
- Manthan Patel, P.P. Amritha, Vinay B Sudheer, and M. Sethumadhavan. DDoS Attack Detection Model using Machine Learning Algorithm in Next Generation Firewall,

2024. *Procedia Computer Science*, vol. 233, pp. 175–183 and <https://doi.org/10.1016/j.procs.2024.03.207>.
- Vinay Tila Patil¹ and Shailesh Shivaji Deore². Deep Learning-Driven IoT Defence: Comparative Analysis of CNN and LSTM for DDoS Detection and Mitigation ARTICLE INFO ABSTRACT, 2024. *Journal of Information Systems Engineering and Management*, vol. 2025, no. 8s, pp. 2468–4376 and <https://www.jisem-journal.com/>.
- Ken Peffers, Tuure Tuunanen, Marcus A. Rothenberger, and Samir Chatterjee. A Design Science Research Methodology for Information Systems Research, 2007. *Journal of Management Information Systems*, vol. 24, no. 3, pp. 45–77 and <https://doi.org/10.2753/MIS0742-1222240302>.
- Mateus Pelloso, Andressa Vergutz, Aldri Santos, and Michele Nogueira. A Self-Adaptable System for DDoS Attack Prediction Based on the Metastability Theory, 2018. 2018 IEEE Global Communications Conference (GLOBECOM), pp. 1–6 and <https://doi.org/10.1109/GLOCOM.2018.8647934>.
- S. Prabavathy, K. Sundarakantham, and S. Mercy Shalinie. Design of cognitive fog computing for intrusion detection in Internet of Things, 2018. *Journal of Communications and Networks*, vol. 20, no. 3, pp. 291–298 and <https://doi.org/10.1109/JCN.2018.000041>.
- M. Premkumar and T. V.P. Sundararajan. Defense Countermeasures for DoS Attacks in WSNs Using Deep Radial Basis Networks, 2021. *Wireless Personal Communications*, vol. 120, no. 4, pp. 2545–2560 and <https://doi.org/10.1007/s11277-021-08545-6>.
- Lakshmi Priya Rachakonda, Madhuri Siddula, and Vanlin Sathya. A comprehensive study on IoT privacy and security challenges with focus on spectrum sharing in Next-Generation networks (5G/6G/beyond), 2024. *High-Confidence Computing*, vol. 4, no. 2 and <https://doi.org/10.1016/j.hcc.2024.100220>.
- Md Mahbubur Rahman, Shaharia Al Shakil, and Mizanur Rahman Mustakim. A survey on intrusion detection system in IoT networks, 2025. *Cyber Security and Applications*, vol. 3, pp. 100082 and <https://doi.org/10.1016/j.csa.2024.100082>.
- Mohd Rahul, Hesham Alhumiany, Mohd Muntjir, and Hesham A Alhumyani. An Analysis of Internet of Things(IoT): Novel Architectures, Modern Applications, Security Aspects and Future Scope with Latest Case Studies Task Scheduling in Cloud Computing View project Internet of Things View project An Analysis of Internet of Things (IoT): Novel Architectures, Modern Applications, Security Aspects and Future Scope with Latest Case Studies IoT in Latest Trends, 2017. www.ijert.org.
- Shailendra Rathore and Jong Hyuk Park. Semi-supervised learning based distributed attack detection framework for IoT, 2018. *Applied Soft Computing*, vol. 72, pp. 79–89 and <https://doi.org/10.1016/j.asoc.2018.05.049>.

- Hemanth Kumar Ravuri, Maria Torres Vega, Jeroen Van Der Hooft, Tim Wauters, Bin Da, and Filip De Turck. On Routing Scalability in Flat SDN Architectures, 2020. Proceedings of the 11th International Conference on Network of the Future, NoF 2020, pp. 23–27 and <https://doi.org/10.1109/NoF50125.2020.9249170>.
- Danda B. Rawat and Swetha R. Reddy. Software Defined Networking Architecture, Security and Energy Efficiency: A Survey, 2017. IEEE Communications Surveys and Tutorials, vol. 19, no. 1, pp. 325–346 and <https://doi.org/10.1109/COMST.2016.2618874>.
- Partha Pratim Ray and Neeraj Kumar. SDN/NFV architectures for edge-cloud oriented IoT: A systematic review, 2021. Computer Communications, vol. 169, pp. 129–153 and <https://doi.org/10.1016/j.comcom.2021.01.018>.
- Khaled Riad, Teng Huang, and Lishan Ke. A dynamic and hierarchical access control for IoT in multi-authority cloud storage, 2020. Journal of Network and Computer Applications, vol. 160 and <https://doi.org/10.1016/j.jnca.2020.102633>.
- Adriana Riofrio and Gustavo Caiza. Cloud Computing Architecture with Dockers for an Industrial Process, 2021. 7th International Conference on Engineering and Emerging Technologies, ICEET 2021 and <https://doi.org/10.1109/ICEET53442.2021.9659760>.
- Monika Roopak, Gui Yun Tian, and Jonathon Chambers. An Intrusion Detection System Against DDoS Attacks in IoT Networks, 2020. 2020 10th Annual Computing and Communication Workshop and Conference (CCWC), pp. 0562–0567 and <https://doi.org/10.1109/CCWC47524.2020.9031206>.
- Syaifuddin Saif, Ridi Ferdiana, and Widyawan Widyawan. IoT-DH Dataset, 2024. IoT-DH Dataset and <https://doi.org/10.17632/8dns3xbckv.1>.
- T. Saipriya and M. Anand. To Secure IoT sensor nodes through Fog computing, 2021. Proceedings of the 2nd International Conference on Electronics and Sustainable Communication Systems, ICESC 2021, pp. 836–844 and <https://doi.org/10.1109/ICESC51422.2021.9532693>.
- Rishi Sairam, Suman Sankar Bhunia, Vijayanand Thangavelu, and Mohan Gurusamy. NETRA: Enhancing IoT Security Using NFV-Based Edge Traffic Analysis, 2019. IEEE Sensors Journal, vol. 19, no. 12, pp. 4660–4671 and <https://doi.org/10.1109/JSEN.2019.2900097>.
- Azar Abid Salih and Maiwan Bahjat Abdulrazaq. Cybernet Model: A New Deep Learning Model for Cyber DDoS Attacks Detection and Recognition, 2024. Computers, Materials & Continua, vol. 78, no. 1, pp. 1275–1295 and <https://doi.org/10.32604/cmc.2023.046101>.

- Mayra Samaniego, Cristian Espana, and Ralph Deters. Smart virtualization for IoT, 2018. Proceedings - 3rd IEEE International Conference on Smart Cloud, SmartCloud 2018, pp. 125–128 and <https://doi.org/10.1109/SmartCloud.2018.00028>.
- Hunor Sándor, Béla Genge, and Gheorghe Sebestyén-Pál. Resilience in the Internet of Things: The Software Defined Networking approach, 2015. Proceedings - 2015 IEEE 11th International Conference on Intelligent Computer Communication and Processing, ICCP 2015, pp. 545–552 and <https://doi.org/10.1109/ICCP.2015.7312717>.
- Abimbola Sangodoyin, Babagana Modu, Irfan Awan, and Jules Pagna Disso. An Approach to Detecting Distributed Denial of Service Attacks in Software Defined Networks, 2018a. 2018 IEEE 6th International Conference on Future Internet of Things and Cloud (FiCloud), pp. 436–443 and <https://doi.org/10.1109/FiCloud.2018.00069>.
- Abimbola Sangodoyin, Tshiamo Sigwele, Prashant Pillai, Yim Fun Hu, Irfan Awan, and Jules Disso. DoS Attack Impact Assessment on Software Defined Networks, 2018b. Lecture Notes of the Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering, LNICST, vol. 231, pp. 11–22 and https://doi.org/10.1007/978-3-319-76571-6_{ }2.
- Reneilson Santos, Danilo Souza, Walter Santo, Admilson Ribeiro, and Edward Moreno. Machine learning algorithms to detect DDoS attacks in SDN, 2020. Concurrency and Computation: Practice and Experience, vol. 32, no. 16 and <https://doi.org/10.1002/cpe.5402>.
- Tinshu Sasi, Arash Habibi Lashkari, Rongxing Lu, Pulei Xiong, and Shahrear Iqbal. A comprehensive survey on IoT attacks: Taxonomy, detection mechanisms and challenges, 2024. Journal of Information and Intelligence, vol. 2, no. 6, pp. 455–513 and <https://doi.org/10.1016/j.jiixd.2023.12.001>.
- Mohamed S. Sawah, Hela Elmannai, Alaa A. El-Bary, Kh. Lotfy, and Osama E. Sheta. Distributed denial of service (DDoS) classification based on random forest model with backward elimination algorithm and grid search algorithm, 2025. Scientific Reports, vol. 15, no. 1, pp. 19063 and <https://doi.org/10.1038/s41598-025-03868-x>.
- Oussama Sbai and Mohamed El boukhari. Data Flooding Intrusion Detection System for MANETs Using Deep Learning Approach, 2020. Proceedings of the 13th International Conference on Intelligent Systems: Theories and Applications, pp. 1–5 and <https://doi.org/10.1145/3419604.3419777>.
- Sebastian Garcia, Agustin Parmisano, and Maria Jose Erquiaga. IoT-23: A labeled dataset with malicious and benign IoT network traffic, 2020. (1.0.0) [Data set]. Zenodo.

- Ahmed Ramzy Shaaban, Essam Abdelwanees, and Mohamed Hussein. Distributed Denial of Service Attacks Analysis, Detection, and Mitigation for the Space Control Ground Network, 2020. Proceedings of the Pakistan Academy of Sciences: Pakistan Academy of Sciences A. Physical and Computational Sciences, vol. 57, no. 2, pp. 97–108.
- Devarshi Shah, Jin Wang, and Q. Peter He. Feature engineering in big data analytics for IoT-enabled smart manufacturing – Comparison between deep learning and statistical learning, 2020. Computers and Chemical Engineering, vol. 141 and <https://doi.org/10.1016/j.compchemeng.2020.106970>.
- Iman Sharafaldin, Arash Habibi Lashkari, Saqib Hakak, and Ali A Ghorbani. Developing Realistic Distributed Denial of Service (DDoS) Attack Dataset and Taxonomy, 2019. IEEE.
- Sabrina Sicari, Alessandra Rizzardi, Daniele Miorandi, and Alberto Coen-Porisini. REATO: REActing TO Denial of Service attacks in the Internet of Things, 2018. Computer Networks, vol. 137, pp. 37–48 and <https://doi.org/10.1016/j.comnet.2018.03.020>.
- Chandrapal Singh and Ankit Kumar Jain. A comprehensive survey on DDoS attacks detection & mitigation in SDN-IoT network, 2024. e-Prime - Advances in Electrical Engineering, Electronics and Energy, vol. 8 and <https://doi.org/10.1016/j.prime.2024.100543>.
- Jagdeep Singh and Sunny Behal. Detection and mitigation of DDoS attacks in SDN: A comprehensive review, research challenges and future directions, 2020. Computer Science Review, vol. 37, pp. 100279 and <https://doi.org/10.1016/j.cosrev.2020.100279>.
- Mitali Sinha. A comprehensive survey of DDoS attack defense systems for different SDN architectures, 2025. Computer Networks, vol. 272 and <https://doi.org/10.1016/j.comnet.2025.111711>.
- Arunan Sivanathan. IoT Behavioral Monitoring via Network Traffic Analysis, 2020. <http://arxiv.org/abs/2001.10632>.
- David Smekal, Jan Hajny, and Zdenek Martinasek. Packet Generators on Field Programmable Gate Array Platform, 2017. vol. 31, pp. 1–4.
- Aviral Srivastava, Dhyan Thakkar, Sharda Valiveti, Pooja Shah, and Gaurang Raval. Anticipated Network Surveillance – An extrapolated study to predict cyber-attacks using Machine Learning and Data Analytics, 2023. <http://arxiv.org/abs/2312.17270>.

- Kalupahana Liyanage Kushan Sudheera, Dinil Mon Divakaran, Rhishi Pratap Singh, and Mohan Gurusamy. ADEPT: Detection and Identification of Correlated Attack Stages in IoT Networks, 2021. *IEEE Internet of Things Journal*, vol. 8, no. 8, pp. 6591–6607 and <https://doi.org/10.1109/JIOT.2021.3055937>.
- Bambang Susilo and Riri Fitri Sari. Intrusion Detection in Software Defined Network Using Deep Learning Approach, 2021. 2021 IEEE 11th Annual Computing and Communication Workshop and Conference, CCWC 2021, pp. 807–812 and <https://doi.org/10.1109/CCWC51732.2021.9375951>.
- K. Swathi and G. Narsimha. Robust Deep Learning Based Framework for Detecting Cyber Attacks from Abnormal Network Traffic, 2023. *International Journal on Recent and Innovation Trends in Computing and Communication*, vol. 11, no. 7, pp. 341–353 and <https://doi.org/10.17762/ijritcc.v11i7.7958>.
- Lo'ai Tawalbeh, Fadi Muheidat, Mais Tawalbeh, and Muhannad Quwaider. IoT privacy and security: Challenges and solutions, 2020. *Applied Sciences (Switzerland)*, vol. 10, no. 12 and <https://doi.org/10.3390/APP10124102>.
- Sahrish Khan Tayyaba, Hasan Ali Khattak, Ahmad Almogren, Munam Ali Shah, Ikram Ud Din, Ibrahim Alkhalifa, and Mohsen Guizani. 5G vehicular network resource management for improving radio access through machine learning, 2020. *IEEE Access*, vol. 8, pp. 6792–6800 and <https://doi.org/10.1109/ACCESS.2020.2964697>.
- Pankaj Thorat and Niraj Kumar Dubey. SDN-based Machine Learning Powered Alarm Manager for Mitigating the Traffic Spikes at the IoT Gateways, 2020. *Proceedings of CONECCT 2020 - 6th IEEE International Conference on Electronics, Computing and Communication Technologies*, no. M1 and <https://doi.org/10.1109/CONECCT50063.2020.9198356>.
- Sergio Trilles, Sahibzada Saadoon Hammad, and Ditsuhi Iskandaryan. Anomaly detection based on Artificial Intelligence of Things: A Systematic Literature Mapping, 2024. *Internet of Things*, vol. 25, pp. 101063 and <https://doi.org/10.1016/j.iot.2024.101063>.
- Bhagyashri Tushir, Hetesh Sehgal, Rohan Nair, Behnam Dezfouli, and Yuhong Liu. The Impact of DoS Attacks on Resource-constrained IoT Devices: A Study on the Mirai Attack, 2021. *arXiv* and <http://arxiv.org/abs/2104.09041>.
- Stephen Ugwuanyi, Rameez Asif, and James Irvine. Network Virtualization: Proof of Concept for Remote Management of Multi-Tenant Infrastructure, 2020. *Proceedings - 2020 IEEE 6th International Conference on Dependability in Sensor, Cloud and Big Data Systems and Application, DependSys 2020*, pp. 98–105 and <https://doi.org/10.1109/DependSys51298.2020.00023>.

- Imtiaz Ullah and Qusay H. Mahmoud. Design and Development of RNN Anomaly Detection Model for IoT Networks, 2022. *IEEE Access*, vol. 10, pp. 62722–62750 and <https://doi.org/10.1109/ACCESS.2022.3176317>.
- United Nations. World Urbanization Prospects: The 2018 Revision, 2018. United Nations, pp. 1–21.
- Ismael Amezcua Valdovinos, Jesús Arturo Pérez-Díaz, Kim Kwang Raymond Choo, and Juan Felipe Botero. Emerging DDoS attack detection and mitigation strategies in software-defined networks: Taxonomy, challenges and future directions, 2021. *Journal of Network and Computer Applications*, vol. 187 and <https://doi.org/10.1016/j.jnca.2021.103093>.
- Karan B. Virupakshar, Manjunath Asundi, Kishor Channal, Pooja Shettar, Somashekar Patil, and D.G. Narayan. Distributed Denial of Service (DDoS) Attacks Detection System for OpenStack-based Private Cloud, 2020. *Procedia Computer Science*, vol. 167, pp. 2297–2307 and <https://doi.org/10.1016/j.procs.2020.03.282>.
- Ruchi Vishwakarma and Ankit Kumar Jain. A survey of DDoS attacking techniques and defence mechanisms in the IoT network, 2020. *Telecommunication Systems*, vol. 73, no. 1, pp. 3–25 and <https://doi.org/10.1007/s11235-019-00599-z>.
- Jin Wang and Liping Wang. SDN-Defend: A Lightweight Online Attack Detection and Mitigation System for DDoS Attacks in SDN, 2022. *Sensors*, vol. 22, no. 21, pp. 8287 and <https://doi.org/10.3390/s22218287>.
- Mingzhe Wang and Qiuliang Zhang. Optimized data storage algorithm of IoT based on cloud computing in distributed system, 2020. *Computer Communications*, vol. 157, pp. 124–131 and <https://doi.org/10.1016/j.comcom.2020.04.023>.
- Tao Wang, Hang Zhu, Fabian Ruffy, Jin Xin, Sivaraman Anirudh, and Dan R K Ports. Multitenancy for Fast and Programmable Networks in the Cloud, 2020.
- Yue Wang, Yiming Jiang, and Julong Lan. FCNN: An Efficient Intrusion Detection Method Based on Raw Network Traffic, 2021. *Security and Communication Networks*, vol. 2021 and <https://doi.org/10.1155/2021/5533269>.
- David Wetherall. Active network vision and reality: lessons from a capsule-based system, 2002.
- Yirui Wu, Dabao Wei, and Jun Feng. Network attacks detection methods based on deep learning techniques: A survey, 2020. *Security and Communication Networks*, vol. 2020 and <https://doi.org/10.1155/2020/8872923>.
- Tan Xinyu, Kujala Pekka, and Hakala Ismo. A design and modeling approach for resource-constrained internet of things devices, 2025. *Design Automation for Embedded Systems*, vol. 29, no. 3 and <https://doi.org/10.1007/s10617-025-09294-6>.

- Tong Xu, Deyun Gao, Ping Dong, Hongke Zhang, Chuan Heng Foh, and Han Chieh Chao. Defending Against New-Flow Attack in SDN-Based Internet of Things, 2017. IEEE Access, vol. 5, pp. 3431–3443 and <https://doi.org/10.1109/ACCESS.2017.2666270>.
- Xiaolong Xu, Qingxiang Liu, Yun Luo, Kai Peng, Xuyun Zhang, Shunmei Meng, and Lianyong Qi. A computation offloading method over big data for IoT-enabled cloud-edge computing, 2019. Future Generation Computer Systems, vol. 95, pp. 522–533 and <https://doi.org/10.1016/j.future.2018.12.055>.
- Kun Yang, Junjie Zhang, Yang Xu, and Jonathan Chao. DDoS Attacks Detection with AutoEncoder, 2020. NOMS 2020 - 2020 IEEE/IFIP Network Operations and Management Symposium, pp. 1–9 and <https://doi.org/10.1109/NOMS47738.2020.9110372>.
- Lingfeng Yang and Hui Zhao. DDoS Attack Identification and Defense Using SDN Based on Machine Learning Method, 2018. 2018 15th International Symposium on Pervasive Systems, Algorithms and Networks (I-SPAN), pp. 174–178 and <https://doi.org/10.1109/I-SPAN.2018.00036>.
- Aiting Yao, Frank Jiang, Xuejun Li, Chengzu Dong, Jia Xu, Yi Xu, Gang Li, and Xiao Liu. A Novel Security Framework for Edge Computing based UAV Delivery System, 2021. 2021 IEEE 20th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom), pp. 1031–1038 and <https://doi.org/10.1109/TrustCom53373.2021.00142>.
- Ibrar Yaqoob, Ibrahim Abaker Targio Hashem, Yasir Mehmood, Abdullah Gani, Salimah Mokhtar, and Sghaier Guizani. Enabling communication technologies for smart cities, 2017. IEEE Communications Magazine, vol. 55, no. 1, pp. 112–120 and <https://doi.org/10.1109/MCOM.2017.1600232CM>.
- Bilal Zahran, Naveed Ahmed, Abdel Rahman Alzoubaidi, and Asri Ngadi. Security and Privacy Issues in Network Function Virtualization: A Review from Architectural Perspective, 2024. IJACSA) International Journal of Advanced Computer Science and Applications, vol. 15, no. 6 and www.ijacsa.thesai.org.
- Ahmad Zainudin, Love Allen Chijioke Ahakonye, Rubina Akter, Dong-Seong Kim, and Jae-Min Lee. An Efficient Hybrid-DNN for DDoS Detection and Classification in Software-Defined IIoT Networks, 2023. IEEE Internet of Things Journal, vol. 10, no. 10, pp. 8491–8504 and <https://doi.org/10.1109/JIOT.2022.3196942>.
- Alejandro Molina Zarca, Jorge Bernal Bernabe, Antonio Skarmeta, and Jose M. Alcaraz Calero. Virtual IoT HoneyNets to mitigate cyberattacks in SDN/NFV-Enabled IoT networks, 2020. IEEE Journal on Selected Areas in Communications, vol. 38, no. 6, pp. 1262–1277 and <https://doi.org/10.1109/JSAC.2020.2986621>.

- Muhammad Zeeshan, Qaiser Riaz, Muhammad Ahmad Bilal, Muhammad K. Shahzad, Hajira Jabeen, Syed Ali Haider, and Azizur Rahim. Protocol-Based Deep Intrusion Detection for DoS and DDoS Attacks Using UNSW-NB15 and Bot-IoT Data-Sets, 2022. IEEE Access, vol. 10, pp. 2269–2283 and <https://doi.org/10.1109/ACCESS.2021.3137201>.
- Peiying Zhang, Chunxiao Jiang, Xue Pang, and Yi Qian. STEC-IoT: A Security Tactic by Virtualizing Edge Computing on IoT, 2021a. IEEE Internet of Things Journal, vol. 8, no. 4, pp. 2459–2467 and <https://doi.org/10.1109/JIOT.2020.3017742>.
- Peiying Zhang, Chao Wang, Chunxiao Jiang, and Abderrahim Benslimane. Security-Aware Virtual Network Embedding Algorithm Based on Reinforcement Learning, 2021b. IEEE Transactions on Network Science and Engineering, vol. 8, no. 2, pp. 1095–1105 and <https://doi.org/10.1109/TNSE.2020.2995863>.
- Peiying Zhang, Ning Chen, Shibao Li, Kim-Kwang Raymond Choo, and Chunxiao Jiang. Multi-Domain Virtual Network Embedding Algorithm based on Horizontal Federated Learning, 2022. pp. 1–12 and <http://arxiv.org/abs/2205.14665>.
- Wei Zhe Zhang, Ibrahim A. Elgendy, Mohamed Hammad, Abdullah M. Iliyasu, Xiaojiang Du, Mohsen Guizani, and Ahmed A. Abd El-Latif. Secure and Optimized Load Balancing for Multitier IoT and Edge-Cloud Computing Systems, 2021c. IEEE Internet of Things Journal, vol. 8, no. 10, pp. 8119–8132 and <https://doi.org/10.1109/JIOT.2020.3042433>.
- Wenxiu Zhang, Shan Jing, and Chuan Zhao. A Survey of SDN Data Plane Attacks and Defense Strategies, 2023. ACM International Conference Proceeding Series, pp. 59–65 and <https://doi.org/10.1145/3605801.3605814>.
- Ziming Zhao, Zhaoxuan Li, Zhihao Zhou, Jiongchi Yu, Zhuoxue Song, Xiaofei Xie, Fan Zhang, and Rui Zhang. DDoS family: A novel perspective for massive types of DDoS attacks, 2024. Computers & Security, vol. 138, pp. 103663 and <https://doi.org/10.1016/j.cose.2023.103663>.
- Lu Zhou, Ye Zhu, Tianrui Zong, and Yong Xiang. A feature selection-based method for DDoS attack flow classification, 2022. Future Generation Computer Systems, vol. 132, pp. 67–79 and <https://doi.org/10.1016/j.future.2022.02.006>.
- Bo Zong and Silan Yi. Distributed Security Framework-based Internet of Things Network Slicing for 5G Networks, 2025. Web Intelligence, vol. 23, no. 4, pp. 613–630 and <https://doi.org/10.1177/24056456251385731>.
- Xinbin Zuo, Xue Pang, Pengping Zhang, Junsan Zhang, Tao Dong, and Peiying Zhang. A security-aware software-defined IoT network architecture, 2020. 2020 IEEE Computing, Communications and IoT Applications, ComComAp 2020 and <https://doi.org/10.1109/ComComAp51192.2020.9398887>.