

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: Carl R. Richardson (2026) "Analysis of Lurie Systems for Learning and Control", University of Southampton, Faculty of Engineering and Physical Sciences, School of Electronics and Computer Science, PhD Thesis, 149pp.

UNIVERSITY OF SOUTHAMPTON

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

Analysis of Lurie Systems for Learning and Control

DOI: <https://doi.org/10.5258/SOTON/PG/T152>

by

Carl Robert Richardson

BEng, MSc

ORCID: [0000-0001-9799-896X](https://orcid.org/0000-0001-9799-896X)

*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

Analysis of Lurie Systems for Learning and Control

by Carl Robert Richardson

Due to the remarkable success of deep learning in a wide variety of applications, neural networks (NNs) are increasingly being adopted to model and control dynamical systems. When controlling dynamical systems, it is key to provide robust guarantees of the system's behaviour, especially in safety-critical applications. However, efficiently verifying such safety and robustness properties for systems involving NNs is challenging, due to the presence and large number of activation functions inside the NN. Furthermore, it is vital to learn robust and generalisable models of dynamical systems, to deal with noise and the, potentially, large space of initial conditions. This is particularly difficult since learning generic, high-dimensional functions is a cursed estimation problem.

This thesis focuses on the forced Lurie system as a general modelling framework which, amongst others, captures many systems involving NNs as special cases. Examples include recurrent neural networks, neural oscillators, and the interconnection of a linear time invariant system with a feed-forward NN. The analysis of Lurie systems has been well-studied in the control theory literature, with the absolute stability problem being a pertinent example. This thesis builds upon this rich literature to develop less conservative stability criteria for Lurie systems involving NNs and to exploit known properties for constructing robust and generalisable models of convergent dynamical systems.

The first contribution develops less conservative stability criteria for Lurie systems, involving NNs with only ReLU activations. Properties of the ReLU function are leveraged to construct tailored quadratic constraints which are incorporated in the stability criteria, posed as semidefinite programs, via the S-procedure. Both continuous and discrete-time cases are examined. The second contribution repeats the same steps for Lurie systems with magnitude nonlinearities. This is a generalisation of the first contribution since a loop transformation between the magnitude and (leaky) ReLU is shown to exist, meaning the stability of Lurie systems involving ReLU or leaky ReLU activations can be verified. Numerical examples highlight the reduced conservatism, particularly for high-dimensional systems.

The final contribution proposes the k -contracting Lurie network (LN) as a robust and generalisable model of convergent dynamical systems. The absolute stability framework is used to establish conditions which guarantee the LN is k -contracting for all slope-restricted nonlinearities. Unconstrained parametrisations of these conditions are then established to restrict training to only search over LNs satisfying the k -contraction property. When tested on dynamical systems datasets involving multiple equilibrium points and limit cycles, the k -contracting LN is an order of magnitude more accurate than existing models, when initial conditions are sampled outside the training distribution and subjected to additive noise.

Contents

List of Figures	ix
List of Tables	xi
Declaration of Authorship	xiii
Acknowledgements	xv
Notation	xix
1 Introduction	1
1.1 Motivation: Control	1
1.2 Motivation: Learning	3
1.3 Thesis Outline	3
2 Background	7
2.1 Mathematical Preliminaries	7
2.1.1 Signals and Systems	7
2.1.1.1 Signals	7
2.1.1.2 Systems	8
2.1.2 Lyapunov Stability Analysis	11
2.1.3 Semidefinite Programming	15
2.1.4 Stochastic Gradient Descent	16
2.1.5 Matrix Decompositions	18
2.1.5.1 Eigenvalue Decomposition	18
2.1.5.2 Singular Value Decomposition	18
2.1.6 Compound Matrices	19
2.1.7 Volume of k -sets	22
2.1.8 k -contraction Analysis	23
2.2 Forced Lurie Systems and Relevant Special Cases	26
2.2.1 Forced Lurie Systems	26
2.2.1.1 Lurie Systems	26
2.2.1.2 Lurie Networks	27
2.2.2 Well-posedness	27
2.2.3 Lurie Systems Involving Feed-forward Neural Networks	28
2.2.4 Relationship Between Lurie Networks and Neural ODEs	29
2.3 Absolute Stability	31
2.3.1 Common Quadratic Constraints	31

2.3.2	S-procedure	34
2.3.3	Literature on Absolute Stability	35
3	Stability Analysis of Lurie Systems with ReLU Nonlinearities	43
3.1	Introduction	43
3.2	Preliminaries	45
3.2.1	Problem Setup	45
3.2.2	Properties of the ReLU Function	46
3.3	Quadratic Constraints	47
3.4	Global Stability Analysis	49
3.5	Convex Relaxations for Theorem 3.2	53
3.5.1	Specific Choice of Λ and \mathbf{W}	53
3.5.2	Specific Choice of \mathbf{H}	54
3.6	Numerical Examples	55
3.6.1	Numerical Setup	55
3.6.2	Discussion	56
3.7	Conclusion	58
4	Stability Analysis of Discrete Lurie Systems with ReLU Nonlinearities	61
4.1	Introduction	61
4.2	Preliminaries	63
4.2.1	Mean Value Theorem	63
4.2.2	Properties of the ReLU Function	63
4.2.3	Positively Homogenous Functions	64
4.2.4	Quadratic Constraints Satisfied by the Repeated ReLU	65
4.2.5	Problem Setup	65
4.3	Global Stability Analysis	66
4.4	Equivalence of Local and Global Stability	69
4.5	Numerical Examples	71
4.5.1	Numerical Setup	71
4.5.2	Discussion	72
4.6	Conclusion	73
5	Stability Analysis of Lurie Systems with Magnitude Nonlinearities	75
5.1	Introduction	75
5.2	Preliminaries	76
5.2.1	Problem Setup	76
5.2.2	Loop Transformations Between Magnitude and ReLU	77
5.2.3	Well-posedness	80
5.3	Quadratic Constraints	81
5.3.1	Simplifying the Positivity QC	83
5.3.2	Sector-like Lemma for Magnitude Nonlinearities	84
5.4	Global Stability Analysis	85
5.4.1	Quadratic Lyapunov Function	85
5.4.2	Lurie-type Lyapunov Function	86
5.5	Convex Relaxations for Theorem 5.2	89
5.5.1	Specific Choice of Λ and \mathbf{M}	89

5.5.2	Specific Choice of \mathbf{H}	89
5.6	Equivalence of Local and Global Stability	90
5.7	Numerical Examples	92
5.7.1	Discussion	93
5.8	Conclusion	95
6	Lurie Networks with Robust Convergent Dynamics	97
6.1	Introduction	97
6.2	Preliminaries: k -contraction Analysis	99
6.3	Lurie Networks	100
6.3.1	Example Lurie Networks	101
6.3.2	Graph Lurie Networks	102
6.4	k -contracting Lurie Networks	103
6.4.1	k -contracting Lurie Networks	103
6.4.2	k -contracting Graph Lurie Networks	107
6.5	Parametrisation of k -contracting Lurie Networks	108
6.5.1	Parametrisation of k -contracting Lurie Networks	109
6.5.2	Parametrisation of k -contracting Graph Lurie Networks	113
6.5.3	Complexity Comparison	114
6.6	Empirical Evaluation: Dynamical Systems	114
6.6.1	Data	114
6.6.1.1	Opinion Dynamics	115
6.6.1.2	Hopfield Network	116
6.6.1.3	Simple Attractor	116
6.6.2	Training	116
6.6.3	k -contracting Lurie Networks	117
6.6.4	k -contracting Graph Lurie Networks	124
6.7	Empirical Evaluation: Fashion MNIST	126
6.7.1	Data	126
6.7.2	Models and Training	126
6.7.3	Results	128
6.8	Related Work	130
6.9	Conclusion	131
7	Conclusion	133
7.1	Thesis Summary	133
7.2	Future Work	134
	References	137
	Appendix A	149

List of Figures

1.1	Contrasting approaches for tokamak plasma control [1].	2
2.1	Trajectories overlaying 2D vector fields with a stable (left), unstable (centre) & asymptotically stable (right) equilibrium point at the origin.	12
2.2	$V(x) = x^\top x$ overlaying the 2D vector fields, from Fig. 2.1, with a stable (left), unstable (centre) & asymptotically stable (right) equilibrium point at the origin.	13
2.3	The 3-parallelotope with vertices $x_1, x_2, x_3 \in \mathbb{R}^3$	23
2.4	Trajectories from 1-contracting (left) and 2-contracting systems (right).	25
2.5	Block diagram of Forced Lurie System.	27
2.6	Interconnection of a feed-forward neural network and an LTI system.	29
2.7	Nonlinearity satisfying the Sector $[0, 1]$ conditions over $y \in [-\frac{\pi}{2}, \frac{\pi}{2}]$	32
2.8	Nonlinearity satisfying the Slope $[-2, 2]$ conditions over $y \in [-1, 2]$	34
2.9	Equivalent Lurie system augmented by the multiplier $M(s)$	39
3.1	Lurie system with static nonlinearity.	44
3.2	ReLU function.	47
3.3	Time taken to solve LMIs for Hopfield network's in Examples 9-12.	58
4.1	Discrete-time Lurie system with static nonlinearity.	62
4.2	Time taken to solve LMIs for Hopfield network's in Examples 9-12.	74
5.1	Loop transformations from (a) ReLU to magnitude (Proposition 5.1) and (b) magnitude to ReLU (Proposition 5.2).	80
5.2	Sector $[-1, 1]$ & Sector $[0, 1]$ alongside the magnitude & (leaky) ReLU.	81
5.3	Time taken to solve LMIs for Hopfield network's in Examples 9-12.	95
6.1	Trajectories from k -contracting dynamical systems. For $k \in \{1, 2\}$, time-invariant k -contracting systems exponentially converge to equilibrium points. 3-contracting systems exponentially converge to limit cycles.	99
6.2	Ground truth trajectories and the associated predictions of each model for the opinion dynamics test set.	120
6.3	Ground truth trajectories and the associated predictions of each model for the OOD opinion dynamics test set.	120
6.4	Ground truth trajectories and the associated predictions of each model for the noisy OOD opinion dynamics test set.	121
6.5	Ground truth trajectories and the associated predictions of each model for the Hopfield network test set.	121
6.6	Ground truth trajectories and the associated predictions of each model for the OOD Hopfield network test set.	122

6.7	Ground truth trajectories and the associated predictions of each model for the noisy OOD Hopfield network test set.	122
6.8	Ground truth trajectories and the associated predictions of each model for the simple attractor test set.	123
6.9	Ground truth trajectories and the associated predictions of each model for the OOD simple attractor test set.	123
6.10	Ground truth trajectories and the associated predictions of each model for the noisy OOD simple attractor test set.	124
6.11	Random sample of images from each class of FMNIST.	127
6.12	t-SNE plots of FMNIST test set (top) and output of k -contracting Lurie network with test set as initial conditions (bottom).	129

List of Tables

2.1	Sector-bounded and slope-restricted activation functions over \mathbb{R}	33
2.2	Criteria comparison.	41
3.1	Properties satisfied by the ReLU function.	48
3.2	Example state space models (A, B, C, D)	55
3.3	Comparison of the maximum series gain.	56
3.4	Comparison of the decision variable count.	57
4.1	Comparison of the maximum series gain.	72
4.2	Comparison of the decision variable count.	73
5.1	Comparison of the maximum series gain.	93
5.2	Comparison of the decision variable count.	94
6.1	Default training settings for the isolated and GC datasets.	117
6.2	Deviations from the default settings for the isolated and GC datasets. . .	118
6.3	Average and best MSE on isolated test sets.	119
6.4	MSE on OOD and noisy isolated test sets.	119
6.5	Average and best MSE on GC test sets.	125
6.6	MSE on OOD and noisy GC test sets.	125
6.7	k -Lurie network and Lurie network settings.	127
6.8	CNN settings.	127
6.9	Training settings used by the Lurie network and k -contracting Lurie network with and without the CNN.	127
6.10	Classification accuracy on FMNIST test set.	128

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:

C. R. Richardson, M. C. Turner, and S. R. Gunn, "Strengthened Circle and Popov Criteria for the stability analysis of feedback systems with ReLU neural networks," *IEEE Control Systems Letters*, 2023

C. R. Richardson, M. C. Turner, and S. R. Gunn, "Strengthened Circle and Popov Criteria and the analysis of ReLU neural networks," in *2024 UKACC 14th International Conference on Control (CONTROL)*, pp. 127–128, IEEE, 2024

C. R. Richardson, M. C. Turner, S. R. Gunn, and R. Drummond, "Strengthened stability analysis of discrete-time Lurie systems involving ReLU neural networks," in *6th Annual Learning for Dynamics and Control Conference*, pp. 209–221, PMLR, 2024

C. R. Richardson, M. C. Turner, and S. R. Gunn, "Lurie networks with k-contracting dynamics," in *New Frontiers in Associative Memory Workshop at ICLR*, 2025

C. R. Richardson, M. C. Turner, and S. R. Gunn, "Lurie networks with robust convergent dynamics," *TMLR: Transactions on Machine Learning Research*, 2025

C. R. Richardson, M. C. Turner, and S. R. Gunn, "Analysis of Lurie systems with magnitude nonlinearities and connections to neural network stability analysis," *IEEE Transactions on Automatic Control*, 2026

Signed:.....

Date:.....

Acknowledgements

Above all, I would like to express my deepest thanks to my supervisor, Professor Matthew Turner, for his unwavering support throughout my PhD. I am truly appreciative for his patience, insightful guidance, and the considerable time and effort he has devoted to supporting me. He consistently offered help when I needed it, while also pushing me to pursue the research directions that interested me most. I feel incredibly fortunate to have had him as a mentor, and I will always look back fondly on our many discussions over coffee and around the whiteboard.

Next, I would like to sincerely thank my examiners, Professor Bing Chu and Professor William Heath, for their careful reading of this thesis and for their insightful and constructive feedback. Their comments and suggestions have been invaluable in refining and strengthening this work.

I am also grateful to my internship advisors, Dr Ján Drgoňa and Dr Ethan King, for giving me the chance to visit the DSMI Group at PNNL. I learnt a lot from our discussions and thoroughly enjoyed the visit. My sincere thanks also go to Emma Miles and Estelle Garcia from the MINDS CDT for their invaluable help throughout my studies, this visit wouldn't have happened without you.

A special thank you to Dr David Llewellyn-Jones from the Alan Turing Institute (ATI) for his brilliant support related to the use of the Baskerville HPC. His timely assistance was instrumental in enabling the completion and publication of Chapter 6.

More generally, I am thankful to the EPSRC and DSTL for funding my PhD. Without their backing I would not have had the privilege to pursue this research.

I would also like to thank my housemates and friends from home, the MINDS CDT, and the ATI enrichment scheme for the escapes to the pub, weekends away, and shared experiences that brought fun and camaraderie to this journey.

Most importantly, thank you to Kate, for always reminding me of what is important in life. Your love and support has been a blessing throughout my PhD - you are the kindest person I know and this journey would not have been the same without you.

Last but not least, thank you to all of my family for your support through the ups and downs. A special thank you to my grandma, June, for always taking an interest in my endeavours. Thank you to my brother, Elliott, for the funny stories and always checking in on me. Finally, thank you to my parents, to whom I dedicate this thesis. You have always believed in me and gave me the best chance to succeed. You are both an inspiration to me and I will always be thankful for having you there.

To my parents,

Notation

Sets

\mathbb{N}	The natural numbers
$[i, j]_{\mathbb{N}}$	The set $\{i, i + 1, \dots, j \mid i < j \text{ and } i, j \in \mathbb{N}\}$
$\mathbb{R}_{\geq 0}$	Set of non-negative real numbers
$\mathbb{R}_{\geq 0}^m$	Set of m -dimensional vectors with non-negative real elements
$\mathbb{R}_{\geq 0}^{m \times n}$	Set of $m \times n$ matrices with non-negative real elements
\mathbb{S}^m	Set of $m \times m$ symmetric matrices
$\mathbb{S}_{\geq 0}^m$	Set of $m \times m$ symmetric matrices with non-negative elements
\mathbb{S}_+^m	Set of $m \times m$ positive definite matrices
\mathbb{D}^{mn}	Set of $m \times n$ diagonal matrices: for $M = [m_{ij}]$, non-zero terms can only occur at elements $i = j$
\mathbb{D}^m	Set of $m \times m$ diagonal matrices
\mathbb{D}_+^{mn}	Set of $m \times n$ diagonal matrices with positive elements
\mathbb{D}_+^m	Set of $m \times m$ diagonal matrices with positive elements
$\mathcal{O}(n)$	Set of $n \times n$ orthogonal matrices
$\text{Skew}(n)$	Set of $n \times n$ skew-symmetric matrices
\mathcal{Z}^m	Set of $m \times m$ Z matrices (non-positive off-diagonal elements)
\mathbb{M}^m	Set of $m \times m$ Metzler matrices (non-negative off-diagonal elements)
\mathcal{L}_2	All functions which are finite square integrable
\mathcal{L}_∞	All functions which are bounded and analytic on the imaginary axis or unit disk
\mathcal{H}_∞	All functions which are analytic in the closed right half complex plane or open unit disk
\mathcal{RH}_∞	Space of real rational transfer function matrices, analytic in the closed right half complex plane or open unit disk
\mathcal{RL}_∞	Space of real rational transfer function matrices, bounded and analytic on the imaginary axis or unit disk

Logic

$f : \mathcal{D}_1 \rightarrow \mathcal{D}_2$

The function f is a map from set \mathcal{D}_1 to set \mathcal{D}_2

$A \Rightarrow B$

Statement A implies statement B

Definitions

$M_1 := M_2$

 M_1 is defined by M_2

$M = [m_{ij}]$

A matrix M with elements m_{ij}

$I \ (I_m)$

The $(m \times m)$ identity matrix

$\binom{n}{k} = \frac{n!}{k!(n-k)!}$

Binomial coefficient where $n, k \in \mathbb{R}_{\geq 0}$ and $n - k \geq 0$ **Limits**

$A \rightarrow B$

 A tends to B

$\epsilon \rightarrow 0_+$

 ϵ tends to 0 from above**Inequalities**

$M \succ 0$

A positive definite matrix M

$M \succeq 0$

A positive semi-definite matrix M

$M \prec 0$

A negative definite matrix M

$M \preceq 0$

A negative semi-definite matrix M

$\sigma_1(M) \geq \dots \geq \sigma_{\min(m,n)}(M)$

Ordered singular values of $m \times n$ matrix M

$\lambda_1(M) \geq \dots \geq \lambda_m(M)$

Ordered eigenvalues of $m \times m$ matrix M **Operators**

$M^{(k)}$

 k -multiplicative compound of matrix M

$M^{[k]}$

 k -additive compound matrix of matrix M

$He(M) := M + M^\top$

Symmetric component of a matrix $M \in \mathbb{R}^{n \times n}$

$\text{trace}(M)$

Trace of matrix M

$\text{diag}(m_1, \dots, m_n)$

Diagonal matrix with elements m_1, \dots, m_n

$\text{blockdiag}(M_1, \dots, M_m)$

Block diagonal matrix with diagonal blocks

M_1, \dots, M_m

$J_f(t, x) := \frac{\partial f}{\partial x}(t, x)$

Jacobian matrix of a function $f : \mathbb{R}_{\geq 0} \times \mathbb{R}^n \rightarrow \mathbb{R}^n$

$\text{vol}(\mathcal{D})$

Volume of the set $\mathcal{D} \subset \mathbb{R}^m$

$\det(M)$

Determinant of matrix M

$\exp(\cdot)$

The exponential function

Norms and Matrix Measures

$\|x\|$

Vector norm of $x \in \mathbb{R}^n$

$\|x\|_{\mathcal{L}_2}$

 \mathcal{L}_2 signal norm of either $x : \mathbb{R} \rightarrow \mathbb{R}^n$ or $x : \mathbb{N} \rightarrow \mathbb{R}^n$

$\|G\|_{\mathcal{L}_\infty}$

 \mathcal{L}_∞ norm of the matrix G

$\|x\|_2 := \sqrt{x^\top x}$

Vector 2-norm of $x \in \mathbb{R}^n$

$$\|x\|_{2,\Theta} := \|\Theta x\|_2$$

Scaled vector 2-norm of $x \in \mathbb{R}^n$ with respect to an invertible scaling matrix $\Theta \in \mathbb{R}^{n \times n}$

$$\mu_2(M) := \lambda_1\left(\frac{M+M^\top}{2}\right)$$

Matrix measure of M induced by the 2-norm

$$\mu_{2,\Theta}(M) := \mu_2(\Theta M \Theta^{-1})$$

Matrix measure of M induced by the scaled 2-norm with respect to an invertible scaling matrix $\Theta \in \mathbb{R}^{n \times n}$

Chapter 1

Introduction

1.1 Motivation: Control

Traditionally, model-driven (conventional) approaches are used to design and control safety-critical systems such as those within an aircraft [8], chemical plant [9], or medical device [10]. Such an approach has been extremely effective and even contributed towards the first moon landing [11; 12]. A key principle of this approach is to construct a simplified model of the system, which accurately describes its behaviour over a local (safe) region of the state space; engineering measures are then put in place to ensure the initial state of the system is restricted to this region [13]. Based on this model, a simple feedback controller is then designed such that the closed loop system, within that region, satisfies the required safety certificates. The linearisation method from Lyapunov analysis [14] is one example of this approach.

In contrast, machine learning (ML) approaches for control are gaining momentum due to their success in a number of challenging control applications [15; 16; 17; 1]. These approaches require vast amounts of data, rather than a simplified model of the system [18]. This data is then leveraged to train a flexible feedback controller, typically in a simulated environment, tuned to optimise some performance metric. Unlike the model-driven approaches, a typical ML approach is not accompanied by any safety certificates [19].

To illustrate some key differences between model-driven and ML approaches, consider the time-varying, nonlinear, and multivariate tokamak plasma control problem [1]. Figures 1.1 (a)-(c) show the key components of the ML approach; (e)-(f) compare the real-time control system architecture for the ML (“our architecture”) and model-driven (“conventional control”) approaches; (g)-(h) depict the tokamak setup. Some of the notable differences include:

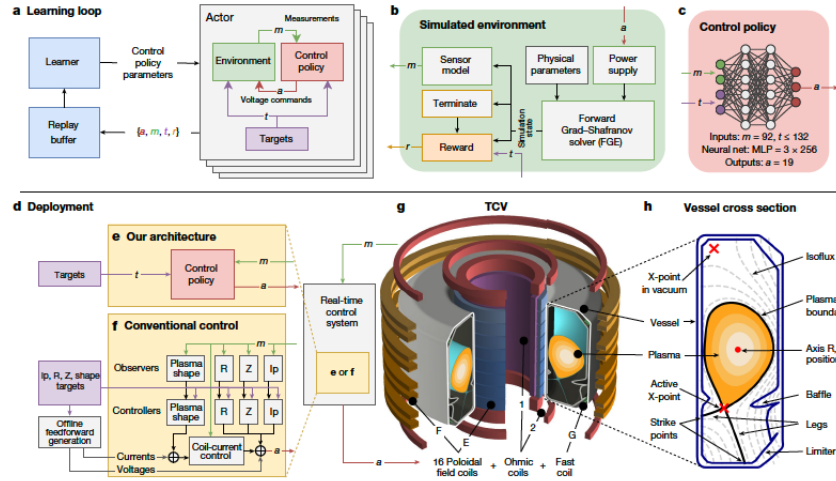


FIGURE 1.1: Contrasting approaches for tokamak plasma control [1].

- The ML approach optimises a single neural network (NN) controller, with multiple inputs and outputs, to directly act on the nonlinear system. In contrast, the model-driven approach uses a set of single input and output proportional-integral-derivative (PID) controllers, designed to not mutually interfere. Furthermore, these PID controllers are designed based on linearised dynamics.
- Whilst the ML controller design process requires a lot of data and compute, it just requires the choice of a policy architecture and the careful design of a performance metric. On the other hand, the model-driven approach requires significant engineering to ensure the sub-systems do not interfere and the linearised dynamics models, which the controllers are determined by, faithfully represent the nonlinear system over the local region of interest.

These differences highlight the benefit of the ML approach in simplifying controller design and for designing controllers to solve difficult problems. However, without any safety certificates, deploying such a controller runs a much greater risk of the plasma's shape becoming unstable leading to cooling, damage to the tokamak, or energy loss.

One notion of safety in this context is stability of the closed loop system. Some attempts have been made to obtain stability certificates, where the stability verification is performed post training [20] and built into the training [21]. Whilst undoubtedly progress has been made, the number of neurons and the highly nonlinear nature of the NNs makes them difficult to analyse. The number of neurons poses the challenge of developing computationally tractable stability criteria for high-dimensional systems. Simultaneously, the criteria must scale well in terms of conservatism, made challenging by the nonlinearities, to avoid misclassification.

1.2 Motivation: Learning

Despite the success of deep learning in a plethora of domains such as sequential processing [22; 23; 24], computer vision [25], language modelling [26] and computational chemistry [27], there is typically an absence of formal guarantees or certificates regarding their robustness and stability. Even in physics-informed ML, the physical constraints are encouraged through the design of an augmented loss function, rather than being imposed on the model [28].

In biological neural systems, stability is a requirement [29]; hence, its use as an inductive bias is another motivation. Some works have approached this problem by parametrising the network weights to guarantee stability [30], exponential stability [31], and contraction [32]. Imposing these constraints presents a fundamental trade-off. The stronger the inductive bias, the more likely the model is to be robust and generalise well; however, the bias will be prohibitive if the assumption isn't true. This trade-off pervades all of ML, with [33] presenting a unified approach for exploiting geometrical properties of a problem.

When learning to model time-invariant dynamical systems from trajectory data, existing stability-constrained models, such as [30; 31; 32], can only learn representations of systems with very particular dynamics. An inductive bias which could capture a larger class of dynamics, such as convergence to multiple equilibrium points or limit cycles, would be more widely applicable, whilst still being robust and more likely to generalise.

1.3 Thesis Outline

The structure of this thesis is described below, outlining the contributions and contents of each chapter:

Chapter 2 begins by presenting the mathematical concepts used throughout this thesis. The fundamental model, and important special cases, studied throughout this thesis are then motivated and discussed. Finally, the absolute stability framework is introduced, along with a review of the literature.

Chapter 3 develops new stability criteria for Lurie systems with a repeated ReLU non-linearity, as this is a popular NN activation function. New quadratic constraints which characterised the ReLU function are presented and leveraged to strengthen classical absolute stability criteria, with low complexity. Numerical examples illustrate the reduced conservatism of the new criteria across a range of example systems, including high-dimensional Hopfield networks. These results have been

published in

C. R. Richardson, M. C. Turner, and S. R. Gunn, “Strengthened Circle and Popov Criteria for the stability analysis of feedback systems with ReLU neural networks,” *IEEE Control Systems Letters*, 2023

C. R. Richardson, M. C. Turner, and S. R. Gunn, “Strengthened Circle and Popov Criteria and the analysis of ReLU neural networks,” in *2024 UKACC 14th International Conference on Control (CONTROL)*, pp. 127–128, IEEE, 2024

Chapter 4 extends the results of Chapter 3 to discrete-time Lurie systems with a repeated ReLU nonlinearity. This is a more natural setting since NN controllers would typically be implemented digitally. Additionally, under mild conditions it is shown that if such a Lurie system has a unique equilibrium point, then it must be globally stable or unstable. The global stability criteria are tested on several example systems, illustrating a desirable trade-off between conservatism and complexity. Publication of these results can be found in

C. R. Richardson, M. C. Turner, S. R. Gunn, and R. Drummond, “Strengthened stability analysis of discrete-time Lurie systems involving ReLU neural networks,” in *6th Annual Learning for Dynamics and Control Conference*, pp. 209–221, PMLR, 2024

Chapter 5 derives innovative stability conditions for Lurie systems with a repeated magnitude nonlinearity. Novel quadratic constraints, capturing the behaviour of the magnitude function, are presented and applied to strengthen classical absolute stability criteria, with low complexity. A loop transformation is also shown to exist between Lurie systems with (leaky) ReLU nonlinearities and magnitude nonlinearities; hence these results can be applied to a wider set of NN activation functions. Finally, under mild conditions, it is shown that if a Lurie system with a positively homogenous nonlinearity has a unique equilibrium point, then it must be globally stable or unstable. This is a parallel result to that in Chapter 4, but for continuous-time Lurie systems. These findings have been reported in

C. R. Richardson, M. C. Turner, and S. R. Gunn, “Analysis of Lurie systems with magnitude nonlinearities and connections to neural network stability analysis,” *IEEE Transactions on Automatic Control*, 2026

Chapter 6 leans on k -contraction analysis to propose absolute stability-like results for studying Lurie networks which converge to multiple equilibrium points or multiple limit cycles. Parametrisations of Lurie networks which satisfy these results are also proposed. As illustrated by some numerical examples, this parametrised Lurie network can be trained to identify systems with convergent dynamics,

whilst ensuring the identified model theoretically satisfies the convergence property, regardless of the time domain which the trajectory data was recorded over. This was shown to improve robustness and out of distribution generalisation. Furthermore, convergent dynamics are thought to be present in a myriad of biological neural functions. Motivated by this, the model is deployed on the MNIST classification task, exhibiting interpretable separation of classes in the latent space. This research has been documented in

C. R. Richardson, M. C. Turner, and S. R. Gunn, "Lurie networks with k-contracting dynamics," in *New Frontiers in Associative Memory Workshop at ICLR*, 2025

C. R. Richardson, M. C. Turner, and S. R. Gunn, "Lurie networks with robust convergent dynamics," *TMLR: Transactions on Machine Learning Research*, 2025

Chapter 7 concludes the thesis and suggests directions for future work.

Chapter 2

Background

2.1 Mathematical Preliminaries

2.1.1 Signals and Systems

This section introduces some basic definitions of signals and systems used throughout the thesis. This background material is based upon [34, Chapter 3], [35], and [14, Chapter 5].

2.1.1.1 Signals

A signal is a (Lebesgue) measurable function that either maps non-negative real numbers to real vectors, in the case of continuous signals, or natural numbers to real vectors for discrete signals. These sets of signals are defined as

$$\mathcal{S}^c := \left\{ u : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^n \right\} \quad \mathcal{S}^d := \left\{ u : \mathbb{N} \rightarrow \mathbb{R}^n \right\} \quad (2.1)$$

where a discrete-time signal, such as $u[k]$, is denoted throughout this thesis by u_k . Readers unfamiliar with measure theory may consider these as sets containing all continuous or discrete signals which could possibly occur in an engineering system, as well as many signals which could not conceivably occur.

Signals form a natural vector space under addition and scalar multiplication, which are defined below for continuous and discrete signals u, v and $\alpha \in \mathbb{R}$

$$(u + v)(t) = u(t) + v(t) \qquad (\alpha u)(t) = \alpha u(t) \qquad (2.2)$$

$$(u + v)_k = u_k + v_k \qquad (\alpha u)_k = \alpha u_k \qquad (2.3)$$

In order to address stability issues, the behaviour of continuous and discrete signals over infinite time intervals must be considered. The infinite-horizon \mathcal{L}_2 space of continuous and discrete signals are defined by

$$\mathcal{L}_2 := \left\{ u \in \mathcal{S}^c \mid \|u\|_{\mathcal{L}_2} < \infty \right\} \qquad \|u\|_{\mathcal{L}_2} := \left(\int_0^\infty \|u(t)\|_2^2 dt \right)^{\frac{1}{2}} \qquad (2.4)$$

$$\mathcal{L}_2 := \left\{ u \in \mathcal{S}^d \mid \|u\|_{\mathcal{L}_2} < \infty \right\} \qquad \|u\|_{\mathcal{L}_2} := \left(\sum_{k=0}^\infty \|u_k\|_2^2 \right)^{\frac{1}{2}} \qquad (2.5)$$

The notation is used interchangeably since it should be clear which is in use depending on whether the signal is continuous or discrete. The \mathcal{L}_2 -norm, $\|\cdot\|_{\mathcal{L}_2}$, is a signal norm and provides a way to measure the size of a signal. As with all norms, the three standard properties hold

- $\|u\|_{\mathcal{L}_2} \geq 0 \quad \forall u \in \mathcal{S}^c$ and $\|u\|_{\mathcal{L}_2} = 0$ if and only if $u(t) \equiv 0$
- $\|\alpha u\|_{\mathcal{L}_2} = \alpha \|u\|_{\mathcal{L}_2} \quad \forall \alpha \in \mathbb{R}_{\geq 0}$
- $\|u + v\|_{\mathcal{L}_2} \leq \|u\|_{\mathcal{L}_2} + \|v\|_{\mathcal{L}_2}$

The first property was presented for continuous signals, but holds analogously for discrete signals too.

2.1.1.2 Systems

A system, G , is a map from an input signal space, \mathcal{S}_1 , to an output signal space, \mathcal{S}_2 . That is, $G : \mathcal{S}_1 \rightarrow \mathcal{S}_2$. Mapping an input signal, $u \in \mathcal{S}_1$, to an output signal, $v \in \mathcal{S}_2$, is expressed using the operator notation $v = Gu$.

Systems form a linear space under addition and multiplication by a scalar ($\alpha \in \mathbb{R}$), which are defined by

$$(G_1 + G_2)u = G_1u + G_2u \qquad (2.6)$$

$$(\alpha G)u = \alpha(Gu) \qquad (2.7)$$

A number of properties about systems are presented next:

- A system is *causal* if the output up to time T depends only on the input up to time T , for every T . That is, G is causal if $P_T G u = P_T G P_T u$, in which P_T is the truncation operator defined below for the continuous and discrete cases

$$(P_T u) = \begin{cases} u(t) & t \leq T \\ 0 & t > T \end{cases} \quad (P_T u) = \begin{cases} u_k & k \leq T \\ 0 & k > T \end{cases} \quad (2.8)$$

where $t \in \mathbb{R}_{\geq 0}$, $k \in \mathbb{N}$ and T is defined appropriately.

- A system is *time-invariant* if the response to an input time-shifted by T is the output also time-shifted by T , for every T . That is, G is time-invariant if $G S_T u = S_T G u$, in which S_T is the time-shift operator defined below for the continuous and discrete cases

$$(S_T u)(t) = u(t - T) \quad (S_T u)_k = u_{k-T} \quad (2.9)$$

where $t - T \in \mathbb{R}_{\geq 0}$ and $k - T \in \mathbb{N}$.

- A system is *input-output stable* if $u \in \mathcal{L}_2 \Rightarrow v \in \mathcal{L}_2$ when $v = G u$.
- A system is *linear* if $G(\alpha_1 u_1 + \alpha_2 u_2) = \alpha_1 G u_1 + \alpha_2 G u_2$ for $\alpha_1, \alpha_2 \in \mathbb{R}$.

Continuous and discrete *state space models* are systems of the form

$$\dot{x} = f(t, x, u) \quad x_{k+1} = f(k, x_k, u_k) \quad (2.10a)$$

$$y = h(t, x, u) \quad y_k = h(k, x_k, u_k) \quad (2.10b)$$

where u is the input signal, x is the system's internal state with $x(0) = x_0$ and y is the output signal. For each fixed $x_0 \in \mathcal{D} \subseteq \mathbb{R}^n$ the state space models in (2.10) define an operator which maps u to y .

Linear Systems

Any linear time-invariant (LTI) system may be represented by the convolution integral or sum

$$v(t) = v_o(t) + \int_0^t G(t - \tau)u(\tau) d\tau \quad v[k] = v_o[k] + \sum_{\tau=0}^k G[k - \tau]u[\tau] \quad (2.11)$$

where $v_o \in \mathbb{R}^n$ denotes the unforced response, $u \in \mathbb{R}^m$ is the forcing function, and $G \in \mathbb{R}^{n \times m}$ is a matrix-valued function called the *impulse response*.

Taking the Laplace transform (for continuous systems) or the Z-transform (for discrete systems) of (2.11) results in

$$v(s) = G(s)u(s) \quad v[z] = G[z]u[z] \quad (2.12)$$

$$G(s) = \int_0^\infty G(t)e^{-st} dt \quad G[z] = \sum_{k=0}^\infty G[k]z^{-k} \quad (2.13)$$

where $s, z \in \mathbb{C}$, and $G(s), G[z]$ are Hermitian and typically referred to as *real transfer function matrices*.

The transfer function matrix describes an input-output stable system if and only if the continuous-time (discrete-time) system G is analytic in the right-half complex plane (analytic in the open unit disk) and the \mathcal{L}_∞ norm of the system is finite. Such a class of systems is known as \mathcal{H}_∞

$$\begin{aligned} \mathcal{H}_\infty &= \left\{ G \mid G(s) \text{ is analytic in } \operatorname{Re}(s) > 0 \text{ and } \|G\|_{\mathcal{L}_\infty} < \infty \right\} \\ \mathcal{H}_\infty &= \left\{ G \mid G[z] \text{ is analytic in } \|z\| > 1 \text{ and } \|G\|_{\mathcal{L}_\infty} < \infty \right\} \end{aligned} \quad (2.14)$$

where the definition is used interchangeably for continuous and discrete systems. In the case that G is *rational*, $G \in \mathcal{H}_\infty$ is true if and only if G has no poles in the right half complex plane (open unit disk). The rational subset of \mathcal{H}_∞ is denoted by $\mathcal{RH}_\infty \subset \mathcal{H}_\infty$.

Linear time invariant (continuous and discrete) state space models have the form

$$\dot{x}(t) = Ax(t) + Bu(t) \quad x_{k+1} = Ax_k + Bu_k \quad (2.15a)$$

$$y(t) = Cx(t) + Du(t) \quad y_k = Cx_k + Du_k \quad (2.15b)$$

One can convert the state space representation of a system to its transfer function representation by taking the Laplace transform (Z-transform). Assuming $x_0 = 0$, this equates to

$$G(s) = C(sI - A)^{-1}B + D \quad G[z] = C(zI - A)^{-1}B + D \quad (2.16)$$

This allows the system to be analysed in the time-domain, via the state space representation, or the frequency domain, via the transfer function representation. In this thesis, the term linear system, and sometimes denoted by $G(s)$ or $G[z]$, is used to refer to both the transfer function (2.16) and state-space realisation (2.15).

2.1.2 Lyapunov Stability Analysis

The stability of dynamical systems is a foundational concern in control theory and applied mathematics. In the context of autonomous systems, systems whose behaviour is governed by time-invariant differential equations, stability analysis is crucial for understanding the long-term evolution of signals and ensuring the reliability of system performance. One of the most powerful and widely used tools for analysing stability of such systems is Lyapunov stability theory, first introduced by the Russian mathematician Aleksandr Lyapunov in the late 19th century.

This section introduces Lyapunov stability theory but only covers the critical results needed for this thesis. The results are presented for continuous-time systems; however, each result has an analogous discrete-time result, with the key difference being that the derivative of the Lyapunov function (\dot{V}) is replaced by a difference equation ($\Delta V = V_{k+1} - V_k$). The results are presented without proof; however, the interested reader can find a more comprehensive introduction in [14, Chapter 4], including proofs.

An autonomous dynamical system can generally be described by a set of ordinary differential equations of the form

$$\begin{aligned} \dot{x} &= f(x) \\ y &= h(x) \end{aligned} \quad (2.17)$$

where $\mathcal{D} \subseteq \mathbb{R}^n$, the vector field $f : \mathcal{D} \rightarrow \mathbb{R}^n$ is locally Lipschitz, and $h : \mathcal{D} \rightarrow \mathbb{R}^m$ maps the state to the output. The central objective is to assess the behaviour of trajectories near equilibrium points. That is, whether solutions that start near an equilibrium remain close to it (stability), converge to it over time (asymptotic stability), or diverge (instability). Figure 2.1 illustrates examples of these three behaviours for vector fields with $n = 2$. The following definition states this more formally, where an equilibrium point is any state $\bar{x} \in \mathcal{D}$ which satisfies $f(\bar{x}) = 0$.




FIGURE 2.1: Trajectories overlaying 2D vector fields with a stable (left), unstable (centre) & asymptotically stable (right) equilibrium point at the origin.

Definition 2.1. *The equilibrium point $x = 0$ of (2.17) is*

- stable if, for each $\epsilon > 0$, there exists $\delta = \delta(\epsilon) > 0$ such that

$$\|x(0)\| < \delta \Rightarrow \|x(t)\| < \epsilon \quad \forall t \geq 0 \quad (2.18)$$

- unstable if, it is not stable.
- asymptotically stable if, it is stable and δ can be chosen such that

$$\|x(0)\| < \delta \Rightarrow \lim_{t \rightarrow \infty} x(t) = 0 \quad (2.19)$$

It should be noted that the various definitions in Definition 2.1 are all with respect to an equilibrium point at the origin. Although no assumptions should be made about the location of equilibrium points, the stability of any equilibrium point can be studied as if it was at the origin, without loss of generality. This can easily be seen by defining the following change of variable $\hat{x} := x - \bar{x}$ which shifts the equilibrium point, \bar{x} , to the origin. Suppose $\bar{x} \neq 0$, then the derivative of \hat{x} is

$$\hat{\dot{x}} = \dot{x} = f(x) = f(\hat{x} + \bar{x}) := g(\hat{x}) \quad (2.20)$$

where $g(0) = 0$. Hence, in the new variable, \hat{x} , the system, $g(\cdot)$, has an equilibrium point at the origin.

Lyapunov's direct method, also known as the second method of Lyapunov, offers an approach to stability analysis without requiring knowledge of the explicit solution, $x(t)$, to (2.17). The method involves the construction of a scalar function $V : \mathbb{R}^n \rightarrow \mathbb{R}$, called a Lyapunov function, which serves as an energy-like measure of the system's state. By analysing the time derivative of $V(x)$ along the trajectories of the system, one can infer the stability properties of an equilibrium point. Lyapunov's direct method is presented next.




FIGURE 2.2: $V(x) = x^\top x$ overlaying the 2D vector fields, from Fig. 2.1, with a stable (left), unstable (centre) & asymptotically stable (right) equilibrium point at the origin.

Theorem 2.1. *Let $x = 0$ be an equilibrium point of (2.17) and $\mathcal{D} \subset \mathbb{R}^n$ be a domain containing $x = 0$. Let $V : \mathcal{D} \rightarrow \mathbb{R}$ be a continuously differentiable function such that*

$$V(0) = 0 \quad V(x) > 0 \quad \forall x \neq 0 \in \mathcal{D} \quad (2.21)$$

$$\dot{V}(x) \leq 0 \quad \forall x \in \mathcal{D} \quad (2.22)$$

then, $x = 0$ is locally stable over the domain \mathcal{D} . Moreover, if

$$\dot{V}(x) < 0 \quad \forall x \neq 0 \in \mathcal{D} \quad (2.23)$$

then, $x = 0$ is locally asymptotically stable over the domain \mathcal{D} .

Note that conditions (2.21) are independent of the vector field, so should be satisfied by design when choosing a Lyapunov function. Conditions (2.22) and (2.23) on the other hand, are dependent on the vector field since

$$\dot{V}(x) = \frac{\partial V^\top}{\partial x} f(x) \quad (2.24)$$

Hence, determining which property of Definition 2.1 the equilibrium point satisfies, via Theorem 2.1, is entirely dependent on the choice of $V(\cdot)$ and the vector field, $f(\cdot)$, governing the dynamical system.

Figure 2.2 illustrates how a trajectory, $x(t)$, of the system (2.17) has a corresponding scalar trajectory $V(x(t))$. This can be seen by following the arrows from an arbitrary initial condition $x(0) \neq 0 \in \mathcal{D}$. In the stable example, the trajectory rotates in a circle without the radius from the origin decaying or growing. For this particular Lyapunov function, this corresponds to $V(x(t))$ staying at a constant value. For the unstable

example, the trajectory is diverging from the origin which corresponds to $V(x(t))$ increasing. Finally, for the asymptotically stable example, the trajectory rotates and converges towards the origin, which corresponds to $V(x(t))$ asymptotically decreasing until it converges to its lower bound at the origin, $V(0) = 0$.

The next result is known as the *Barbashin-Krasovskii Theorem* and provides Lyapunov conditions which ensure the origin is *globally asymptotically stable* [14, Theorem 4.2]. With respect to Definition 2.1, this implies the equilibrium point is asymptotically stable for $\delta \rightarrow \infty$. It can also be considered as an analogous result to Theorem 2.1 for the case $\mathcal{D} = \mathbb{R}^n$. This result is fundamental to most of the later chapters in this thesis.

Theorem 2.2. *Let $x = 0$ be an equilibrium point of (2.17) and let $V : \mathbb{R}^n \rightarrow \mathbb{R}$ be a continuously differentiable function such that*

$$V(0) = 0 \quad \|x\| \rightarrow \infty \Rightarrow V(x) \rightarrow \infty \quad V(x) > 0 \quad \forall x \neq 0 \quad (2.25)$$

$$\dot{V}(x) < 0 \quad \forall x \neq 0 \quad (2.26)$$

then, $x = 0$ is globally asymptotically stable.

Akin to Theorem 2.1, the conditions (2.25) are independent of the vector field, so should be satisfied by design when constructing a Lyapunov function. When studying global asymptotic stability, this requires the Lyapunov function to additionally be radially unbounded. Condition (2.26) on the other hand is dependent on the vector field.

In practice, the challenge often lies in constructing an appropriate Lyapunov function. Approaches for doing this, involving semi-definite programming, will be discussed later in this chapter for certain classes of Lyapunov functions (Section 2.3.2). For now it is just important to highlight that the failure to verify stability using a chosen Lyapunov function does not provide conclusive evidence the origin is unstable for a given vector field; instead, it just implies stability cannot be verified with that choice of Lyapunov function. It may be possible to verify stability with a different choice. On the other hand, verifying stability using a chosen Lyapunov function does provide conclusive evidence the origin is stable. Nevertheless, Lyapunov stability analysis provides a systematic and rigorous approach to stability analysis and is widely used in both theoretical studies and engineering applications such as robotics, aerospace systems, and power networks [36; 37].

2.1.3 Semidefinite Programming

Linear matrix inequalities (LMIs) play a central role in many problems in control theory [38], where they are frequently used to express convex constraints. An LMI constraint has the following canonical form

$$F_0 + \sum_{i=1}^m F_i x_i \succeq 0 \quad (2.27)$$

where $x \in \mathbb{R}^m$ is the decision variable, and the matrices $F_0, \dots, F_m \in \mathbb{S}^n$ are provided as part of the problem data. Many LMIs are not presented in this canonical form, but can easily be manipulated into it. To illustrate this, an example is presented next.

Example 2.1. Consider the following LMI

$$\mathbf{P}A + A^\top \mathbf{P} \prec 0 \quad (2.28)$$

used for verifying global asymptotic stability of the linear system $\dot{z} = Az$, with a quadratic Lyapunov function $V(z) = z^\top \mathbf{P}z$. The matrix $A \in \mathbb{R}^{n \times n}$ is the problem data and the matrix $\mathbf{P} \in \mathbb{S}^n$ contains the decision variables.

Firstly, this negative definite inequality can be expressed as a positive semi-definite inequality by simply negating the left hand side and subtracting the term ϵI where $\epsilon \rightarrow 0_+$. This results in

$$-(\mathbf{P}A + A^\top \mathbf{P} + \epsilon I) \succeq 0 \quad (2.29)$$

For the case $n = 2$, this is equivalent to (2.27) with $x = [p_{11}, p_{12}, p_{22}]^\top$, $F_0 = -\epsilon I$ and

$$F_1 = - \begin{bmatrix} 2a_{11} & a_{12} \\ a_{12} & 0 \end{bmatrix} \quad F_2 = - \begin{bmatrix} 2a_{21} & a_{11} + a_{22} \\ a_{11} + a_{22} & 2a_{12} \end{bmatrix} \quad F_3 = - \begin{bmatrix} 0 & a_{21} \\ a_{21} & 2a_{22} \end{bmatrix} \quad (2.30)$$

A class of optimisation problems that naturally incorporates LMIs is semi-definite programming (SDP). In SDP, the objective is a linear function of the decision variables, and the constraints are expressed as one or more LMIs. Given $F_0, \dots, F_m \in \mathbb{S}^n$ and a vector $b \in \mathbb{R}^m$, the canonical form of a SDP problem is

$$\begin{aligned} & \max_{x \in \mathbb{R}^m} \quad b^\top x \\ & \text{subject to} \quad F_0 + \sum_{i=1}^m F_i x_i \succeq 0 \end{aligned} \quad (2.31)$$

Throughout this thesis, several SDP problems are formulated; however, like Example 2.1, they do not involve an objective function. These problems are referred to as *feasibility problems* since the user is only interested in the existence of a solution to the LMI constraints.

SDP problems are a subclass of convex optimisation problems. Importantly, LMI constraints define convex sets, and the intersection of multiple LMIs also yields a convex set. Since both the objective function and the feasible set defined by LMIs are convex, SDPs can be solved efficiently using well-established algorithms. In particular, interior-point methods are capable of solving SDP problems in polynomial-time, making them tractable for problems of moderate size [39].

The development of general-purpose SDP solvers, such as SeDuMi [40], SDPT3 [41], and MOSEK [42] has enabled the widespread application of SDP across control and machine learning. Moreover, modelling tools like YALMIP [43] provide user-friendly interfaces for specifying and solving SDP problems, significantly simplifying their implementation. For example, these tools can automatically handle the conversion of an LMI, such as (2.28), to the canonical form (2.27).

The main limitation of SDP is caused by the polynomial-time complexity where the actual time taken to solve an LMI is dependent on several factors including the number of decision variables, number of LMI constraints and dimensions of the LMIs. For large problems, such as those involving large neural networks, solving SDP problems can become intractable. In recent years, considerable research has been devoted to enhancing the scalability of SDP solvers; for example, exploiting sparsity in SDP problems to improve efficiency [44]. Sparse SDP problems arise in various ways, including problems involving neural networks.

2.1.4 Stochastic Gradient Descent

Nonlinear optimisation plays a central role in machine learning. At its core, training a neural network is a high-dimensional nonlinear optimisation problem, often involving millions of parameters. In many cases, the goal is to minimise a loss function that quantifies the discrepancy between the model's predictions and the observed data. Gradient methods are a suitable class of algorithms for this type of problem where the process entails iteratively traversing the parameter space in the search of an optimal or near-optimal point where the loss function is minimised [45]. The stochastic gradient descent (SGD) algorithm is one such method and was employed in this thesis. This section provides the essential background to understand the mechanics behind this algorithm, based on [46], and some of the related design choices made in Chapter 6. Many alternative algorithms exist, but a discussion on this is beyond the scope of this section.

Historically, nonlinear optimisation theory evolved from calculus and numerical analysis, with methods such as gradient descent, Newton’s method, and quasi-Newton methods (e.g., BFGS) forming the foundation of unconstrained nonlinear optimisation [45]. These techniques rely on gradient (and in some cases, Hessian) information to iteratively refine parameter estimates. While effective in many low to moderate dimensional settings, their computational cost and memory requirements often scale poorly with the size of modern deep learning models and datasets. As an example, the gradient descent algorithm takes the form

$$w_{k+1} = w_k - \eta \nabla_{w_k} L(y, \hat{y}) \quad \text{where} \quad \hat{y} = \bar{y}(w_k, x) \quad (2.32)$$

where the observed output data is denoted by y . The model’s predictions are represented by \hat{y} and are dependent on the model parameters at iteration k , w_k , along with the chosen model, \bar{y} . Finally, the observed input data is denoted by x , the learning rate is η , and the loss function is $L(\cdot, \cdot)$. In this setting, each iteration k corresponds to an epoch (pass over the entire dataset). The mean squared error (MSE) is one example of a loss function and is used later in this thesis.

Unlike traditional gradient descent, which computes gradients over the entire training dataset, stochastic gradient descent (SGD) operates on smaller randomly sampled batches of data. This not only addresses issues with memory capacity, encountered when loading in a full dataset for gradient descent, but the introduction of randomly sampled data batches has been observed to help escape saddle points and shallow local minima, thus facilitating better generalisation in neural networks. One epoch of the SGD algorithm takes the form

Algorithm 1: Stochastic Gradient Descent

Input:

Data split into N batches $\{(x_1, y_1), \dots, (x_N, y_N)\}$, with batch size Q for each (x_i, y_i)

Initial model parameters w_1

Learning rate η

for $k = 1, \dots, N$ **do**

$$\left[\begin{array}{l} w_{k+1} = w_k - \eta \nabla_{w_k} L(y_k, \hat{y}_k) \quad \text{where} \quad \hat{y}_k = \bar{y}(w_k, x_k) \end{array} \right. \quad (2.33)$$

where the final model parameters become the initial model parameters for the next epoch.

Despite its simplicity and effectiveness, SGD suffers from several limitations including sensitivity to the choice of learning rate and slow convergence. These limitations have motivated the development of adaptive gradient methods (beyond the scope of this section) and practical tools such as learning rate scheduling which reduces the learning

rate by a specified factor when a certain condition is met (for example, the training loss has plateaued) or at a specified epoch. Refer to [46] for a discussion on more advanced algorithms and practical approaches. If used, these would likely further improve the empirical results in Chapter 6; however, such methods are not the focus of this thesis so are omitted here.

2.1.5 Matrix Decompositions

Matrix decompositions are fundamental tools in linear algebra and lie at the heart of a broad range of applications across control, optimisation, and machine learning. These decompositions provide structured ways to factor matrices into simpler, interpretable components, enabling more efficient algorithms for solving linear systems, performing dimensionality reduction, and encoding properties of dynamical systems (Chapter 6).

Two of the most widely used decompositions are the eigenvalue decomposition and the singular value decomposition, both of which offer deep insight into the structure of a matrix and have well-established theoretical foundations. This section briefly reviews these two decompositions, focusing on their mathematical properties and practical significance.

2.1.5.1 Eigenvalue Decomposition

The eigenvalue decomposition (EVD) is applicable to square matrices and is based on the principle that, under certain conditions, a matrix can be represented in terms of its eigenvalues and eigenvectors. For a diagonalisable and symmetric matrix $A \in \mathbb{S}^n$, there exists an orthogonal matrix of eigenvectors, $V \in \mathcal{O}(n)$, and a real diagonal matrix of corresponding eigenvalues, $\Lambda \in \mathbb{D}^n$, such that

$$A = V\Lambda V^\top \tag{2.34}$$

where $V = [v_1, \dots, v_n]$ and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$ satisfy the eigenvalue problem $Av_i = \lambda_i v_i$.

Generalisations of the EVD exist for when A is not symmetric or diagonalisable [14, Section 4.3]; however, neither of these decompositions are needed in this thesis.

2.1.5.2 Singular Value Decomposition

The singular value decomposition (SVD) states that for any matrix $A \in \mathbb{R}^{m \times n}$ there exists orthogonal matrices $U \in \mathcal{O}(m)$, $V \in \mathcal{O}(n)$ and a diagonal matrix $\Sigma \in \mathbb{D}_+^{mn}$ such that

$$A = U\Sigma V^\top \quad (2.35)$$

The diagonal entries of Σ , known as the singular values of A , reveal important information about the rank and conditioning of the matrix. The vectors of U and V are respectively referred to as the left and right singular vectors.

For a matrix $A \in \mathbb{R}^{m \times n}$, the SVD of A is related to the EVD of $A^\top A$ and AA^\top . Using the SVD of A and the fact that orthogonal matrices satisfy the property $U^\top = U^{-1}$, the following relationships are clear

$$A^\top A = V\Sigma^2 V^{-1} \quad AA^\top = U\Sigma^2 U^{-1} \quad (2.36)$$

and are in the form of the EVD.

2.1.6 Compound Matrices

In this section, several known definitions and algebraic results related to compound matrices are documented. Most of this thesis can be understood without these concepts, but they are vital to the development of Chapter 6. The results are included without proof; the interested reader should refer to [47] for a more detailed tutorial on the topic.

Let n be a positive integer and fix $k \in [1, n]_{\mathbb{N}}$. The ordered set of *increasing* sequences of k integers from $[1, n]_{\mathbb{N}}$ is denoted by $Q(k, n)$. For example

$$Q(3, 4) = \{(1, 2, 3), (1, 2, 4), (1, 3, 4), (2, 3, 4)\} \quad (2.37)$$

In this example, the first argument of Q , $k = 3$, indicates the length of the sequence and the second argument, $n = 4$, indicates the upper bound in the set of sequences.

Now consider a matrix $W \in \mathbb{R}^{n \times m}$. For $\alpha \in Q(k, n)$ and $\beta \in Q(k, m)$, the matrix $W[\alpha \parallel \beta]$ denotes the $k \times k$ sub-matrix obtained by taking the entries of W along the rows indexed by α and columns indexed by β . As an example, if $k = 2$ and $n = m = 4$, then

$$Q(2, 4) = \{(1, 2), (1, 3), (1, 4), (2, 3), (2, 4), (3, 4)\} \quad (2.38)$$

The sub-matrix $W[\alpha \parallel \beta]$, for elements $\alpha = (1, 2)$ and $\beta = (3, 4)$, would then be given by

$$W[(1,2) \parallel (3,4)] = \begin{bmatrix} w_{13} & w_{14} \\ w_{23} & w_{24} \end{bmatrix} \quad (2.39)$$

The k -minors of the matrix W are defined as $W(\alpha \parallel \beta) := \det(W[\alpha \parallel \beta])$.

Definition 2.2 (k -multiplicative compound). Let $W \in \mathbb{R}^{n \times m}$ and fix $k \in [1, \min(n, m)]_{\mathbb{N}}$. The k -multiplicative compound of W , denoted $W^{(k)}$, is the $\binom{n}{k} \times \binom{m}{k}$ matrix containing all the k -minors of W ordered lexicographically.

For example, if $n = m = 3$ and $k = 2$ then $\alpha, \beta \in Q(2, 3) = \{(1, 2), (1, 3), (2, 3)\}$ and

$$W^{(2)} = \begin{bmatrix} W((1,2) \parallel (1,2)) & W((1,2) \parallel (1,3)) & W((1,2) \parallel (2,3)) \\ W((1,3) \parallel (1,2)) & W((1,3) \parallel (1,3)) & W((1,3) \parallel (2,3)) \\ W((2,3) \parallel (1,2)) & W((2,3) \parallel (1,3)) & W((2,3) \parallel (2,3)) \end{bmatrix} \quad (2.40)$$

Some important special cases include

$$W^{(1)} = W \quad W^{(n)} = \det(W) \quad (pI_n)^{(k)} = p^k I_s \quad W \in \mathbb{D}^n \rightarrow W^{(k)} \in \mathbb{D}^s \quad (2.41)$$

with $s := \binom{n}{k}$. Next, a series of algebraic results concerned with the k -multiplicative compound are presented.

Fact 2.1 (Cauchy-Binet Formula). If $U \in \mathbb{R}^{n \times m}$, $V \in \mathbb{R}^{m \times p}$ and $k \in [1, \min(n, m, p)]_{\mathbb{N}}$, then

$$(UV)^{(k)} = U^{(k)} V^{(k)} \quad (2.42)$$

Fact 2.2. Fix $k \in [1, \min(n, m)]_{\mathbb{N}}$. As a consequence of Definition 2.2, if $W \in \mathbb{R}^{n \times m}$ then

$$(W^\top)^{(k)} = (W^{(k)})^\top \quad (2.43)$$

Fact 2.3. Fix $k \in [1, n]_{\mathbb{N}}$. If $W \in \mathbb{R}^{n \times n}$ is non-singular, then by Fact 2.1

$$(W^{-1})^{(k)} = (W^{(k)})^{-1} \quad (2.44)$$

Fact 2.4. Fix $k \in [1, \min(n, m, p)]_{\mathbb{N}}$. If $W \in \mathbb{R}^{n \times n}$, $U \in \mathbb{R}^{p \times n}$ and $V \in \mathbb{R}^{n \times p}$, then by Fact 2.1

$$(UWV)^{(k)} = U^{(k)} W^{(k)} V^{(k)} \quad (2.45)$$

Fact 2.5. Fix $k \in [1, n]_{\mathbb{N}}$. An implication of Fact 2.1 is that if $W \in \mathbb{R}^{n \times n}$ with eigenvalues $\lambda_1, \dots, \lambda_n$, then the eigenvalues of $W^{(k)}$ are the $\binom{n}{k}$ products

$$\left\{ \prod_{l=1}^k \lambda_{i_l} : 1 \leq i_1 < \dots < i_k \leq n \right\} \quad (2.46)$$

The definition of a second compound matrix, the k -additive compound, and a set of related algebraic results are introduced next.

Definition 2.3 (k -additive compound). Let $W \in \mathbb{R}^{n \times n}$ and $k \in [1, n]_{\mathbb{N}}$. The k -additive compound of W is the $\binom{n}{k} \times \binom{n}{k}$ matrix defined by

$$W^{[k]} := \frac{d}{d\epsilon} (I_n + \epsilon W)^{(k)} \Big|_{\epsilon=0} \quad (2.47)$$

Special cases include

$$W^{[1]} = W \quad W^{[n]} = \text{trace}(W) \quad (pI_n)^{[k]} = kpI_s \quad W \in \mathbb{D}^n \rightarrow W^{[k]} \in \mathbb{D}^s \quad (2.48)$$

with $s := \binom{n}{k}$. Some useful algebraic results related to the k -additive compound are presented next.

Fact 2.6. If $W \in \mathbb{R}^{n \times n}$ and $k \in [1, n]_{\mathbb{N}}$, then as a consequence of Definition 2.3

$$(W^\top)^{[k]} = (W^{[k]})^\top \quad (2.49)$$

Fact 2.7. Fix $k \in [1, n]_{\mathbb{N}}$. For $W \in \mathbb{R}^{n \times n}$ with eigenvalues $\lambda_1, \dots, \lambda_n$, the eigenvalues of $W^{[k]}$ are the $\binom{n}{k}$ sums

$$\left\{ \sum_{l=1}^k \lambda_{i_l} : 1 \leq i_1 < \dots < i_k \leq n \right\} \quad (2.50)$$

An important consequence of Fact 2.7 is that if W is positive definite (semi-definite), then this property is upheld by $W^{[k]}$. Opposite conclusions can be drawn if W is negative definite (semi-definite).

Fact 2.8. Fix $k \in [1, n]_{\mathbb{N}}$. If $U, V \in \mathbb{R}^{n \times n}$, then

$$(U + V)^{[k]} = U^{[k]} + V^{[k]} \quad (2.51)$$

Fact 2.9. Fix $k \in [1, \min(n, p)]_{\mathbb{N}}$. If $W \in \mathbb{R}^{n \times n}$, $U \in \mathbb{R}^{p \times n}$, $V \in \mathbb{R}^{n \times p}$ and $UV = I_p$, then

$$(UWV)^{[k]} = U^{(k)} W^{[k]} (U^{(k)})^{-1} \quad (2.52)$$

2.1.7 Volume of k -sets

This section aims to provide a clear geometric interpretation of the k -multiplicative compound. The section begins by defining a k -set (the codomain of a function dependent on k variables) before presenting Theorem 2.3, the key result which exposes the relationship between the volume of a k -set and the k -multiplicative compound of the Jacobian. A k -parallelotope is then shown as an example. Like before, these are existing results, so are presented without proof. Refer to [48] for more information.

Definition 2.4 (k -sets). Consider a compact set $\mathcal{D} \subset \mathbb{R}^k$ and a continuous differentiable map $\Psi : \mathcal{D} \rightarrow \mathbb{R}^n$, with $k \in [1, n]_{\mathbb{N}}$. The codomain of Ψ is given by the parametrised set

$$\Psi(\mathcal{D}) := \left\{ \Psi(r) : r \in \mathcal{D} \right\} \subseteq \mathbb{R}^n \quad (2.53)$$

Since \mathcal{D} is compact and $\Psi(\cdot)$ is continuous, $\Psi(\mathcal{D})$ is a closed set.

Theorem 2.3 (Volume of k -sets). Fix $k \in [1, n]_{\mathbb{N}}$. Consider a compact set $\mathcal{D} \subset \mathbb{R}^k$ and a continuously differentiable map $\Psi : \mathcal{D} \rightarrow \mathbb{R}^n$. The volume of the parametrised set (2.53) is given by

$$\text{vol}(\Psi(\mathcal{D})) = \int_{\mathcal{D}} \left\| J_{\Psi}^{(k)}(r) \right\|_2 dr \quad (2.54)$$


where $J_{\Psi}(r) = \left[\frac{\partial \Psi(r)}{\partial r_1} \quad \dots \quad \frac{\partial \Psi(r)}{\partial r_k} \right]$ is the Jacobian of Ψ . Note that as $J_{\Psi} : \mathcal{D} \rightarrow \mathbb{R}^{n \times k}$, it implies $J_{\Psi}^{(k)} : \mathcal{D} \rightarrow \mathbb{R}^s$ with $s = \binom{n}{k}$.

Theorem 2.3 states that the volume of a k -set is governed by the k -multiplicative compound of its Jacobian. An important geometrical feature, is that for $k \in \{1, 2, 3\}$ the volume of the k -set is equivalent to the standard notions of length, area and volume. The k -parallelotope is now presented as an example of a k -set.

Definition 2.5 (k -parallelotope). Fix $k \in [1, n]_{\mathbb{N}}$ and let vectors $x_1, \dots, x_k \in \mathbb{R}^n$. The parallelotope generated by these vectors (and the zero vertex) is the set given by

$$P(x_1, \dots, x_k) := \left\{ \sum_{i=1}^k r_i x_i : r_i \in [0, 1] \forall i \right\} \quad (2.55)$$

Based on Definition 2.5, the k -parallelotope is a k -set with the following compact domain \mathcal{D} and continuous differentiable function $\Psi(r; x_1, \dots, x_k)$, as illustrated for $k = n = 3$ in Figure 2.3.

FIGURE 2.3: The 3-parallelotope with vertices $x_1, x_2, x_3 \in \mathbb{R}^3$.

$$\mathcal{D} := \left\{ r \in \mathbb{R}^k : r_i \in [0, 1] \forall i \in [1, k]_{\mathbb{N}} \right\} \quad \Psi(r; x_1, \dots, x_k) := \sum_{i=1}^k r_i x_i \quad (2.56)$$

The Jacobian of the k -parallelotope is $J_{\Psi}(r) = X := [x_1, \dots, x_k]$ and from Theorem 2.3, the volume of the k -parallelotope is given by $\|X^{(k)}\|_2$.

2.1.8 k -contraction Analysis

This section leverages the algebraic results from Section 2.1.6 and the geometrical intuition from Section 2.1.7 to define k -contraction and provide intuition for its geometrical interpretation. Refer to [48; 47] for more in depth tutorials on the topic. Consider the time-varying nonlinear system

$$\dot{x} = f(t, x) \quad (2.57)$$

where $f : \mathbb{R}_{\geq 0} \times \mathbb{R}^n \rightarrow \mathbb{R}^n$. It is assumed throughout that f is continuously differentiable with respect to x . Fix $k \in [1, n]_{\mathbb{N}}$ and let \mathcal{S}^k denote the unit simplex.

$$\mathcal{S}^k := \left\{ r \in \mathbb{R}^k : r_i \geq 0 \text{ and } r_1 + \dots + r_k \leq 1 \right\} \quad (2.58)$$

The convex combination of a set of initial conditions $x_1, \dots, x_{k+1} \in \mathbb{R}^n$ is defined by $h : \mathcal{S}^k \rightarrow \mathbb{R}^n$.

$$h(r; x_1, \dots, x_{k+1}) := \sum_{i=1}^k r_i x_i + \left(1 - \sum_{i=1}^k r_i \right) x_{k+1} \quad (2.59)$$

The set $h(\mathcal{S}^k)$ is a k -set (Definition 2.4) and can be thought of as a k -dimensional body of states representing initial conditions of (2.57). Next, $w_i(t, r)$ is defined as a measure of the sensitivity of a solution to (2.57) at time t , to a change in the initial condition $h(r)$, caused by a change in r_i .

$$w_i(t, r) := \frac{\partial x(t, h(r))}{\partial r_i} \quad \text{where } w_i(0, r) = \frac{\partial h(r)}{\partial r_i} = x_i - x_{k+1} \quad \text{for all } i \in [1, k]_{\mathbb{N}} \quad (2.60)$$

Now the necessary background has been detailed, k -contraction can be defined along with an intuitive geometric interpretation.

Definition 2.6 (k -contraction). Fix $k \in [1, n]_{\mathbb{N}}$. The nonlinear system (2.57) is k -contracting if there exists an $\eta > 0$ and a vector norm $\|\cdot\|_2$ such that for any $x_1, \dots, x_{k+1} \in \mathbb{R}^n$ and any $r \in \mathcal{S}^k$, the mapping $W : \mathbb{R}_{\geq 0} \times \mathcal{S}^k \rightarrow \mathbb{R}^{n \times k}$ defined by $W(t, r) := [w_1(t, r), \dots, w_k(t, r)]$ satisfies

$$\|W^{(k)}(t, r)\|_2 \leq \exp(-\eta t) \|W^{(k)}(0, r)\|_2 \quad \forall t \in \mathbb{R}_{\geq 0} \quad (2.61)$$

To explain the geometric meaning of this definition, pick a domain $\mathcal{D} \subseteq \mathcal{S}^k$ and recall that $h(\mathcal{D})$ is a k -set representing k -dimensional bodies of initial conditions for (2.57); thus, $x(t, h(\mathcal{D})) := \{x(t, h(r)) : r \in \mathcal{D}\}$ is a k -set describing how k -dimensional bodies evolve over time. Theorem 2.3 is now leveraged to show how the volume of these bodies evolves over time, when governed by (2.57).

$$\begin{aligned} \text{vol}(x(t, h(\mathcal{D}))) &= \int_{\mathcal{D}} \|J_x(t, h(r))^{(k)}\|_2 dr \\ &= \int_{\mathcal{D}} \left\| \begin{bmatrix} \frac{\partial x(t, h(r))}{\partial r_1} & \dots & \frac{\partial x(t, h(r))}{\partial r_k} \end{bmatrix}^{(k)} \right\|_2 dr \\ &= \int_{\mathcal{D}} \|W^{(k)}(t, r)\|_2 dr \end{aligned} \quad (2.62)$$

If (2.57) is k -contracting, the volume of these bodies is upper bounded by the initial volume scaled by an exponentially decaying term i.e., (2.61).

$$\begin{aligned} \text{vol}(x(t, h(\mathcal{D}))) &\leq \exp(-\eta t) \int_{\mathcal{D}} \|W^{(k)}(0, r)\|_2 dr \\ &= \exp(-\eta t) \left\| \begin{bmatrix} (x_1 - x_{k+1}) & \dots & (x_k - x_{k+1}) \end{bmatrix}^{(k)} \right\|_2 \int_{\mathcal{D}} dr \end{aligned} \quad (2.63)$$




FIGURE 2.4: Trajectories from 1-contracting (left) and 2-contracting systems (right).

Therefore, k -contraction of (2.57) implies the volume of k -dimensional bodies $x(t, h(\mathcal{D}))$ converges to zero at an exponential rate. This can also be interpreted as *the volume of k -dimensional bodies is contracting or converging to a $(k - 1)$ -dimensional subspace*. Figure 2.4 provides an illustration of a 1-contracting and 2-contracting system.

Theorem 2.4. Fix $k \in [1, n]_{\mathbb{N}}$ and consider the nonlinear system (2.57) assumed to be continuously differentiable. If there exists $\eta > 0$ and an invertible matrix $\Theta \in \mathbb{R}^{n \times n}$ such that

$$\mu_{2, \Theta^{(k)}}(J_f^{[k]}(t, x)) \leq -\eta \quad \forall x \in \mathbb{R}^n \text{ and } t \in \mathbb{R}_{\geq 0} \quad (2.64)$$

then (2.57) is k -contracting in the 2-norm with respect to the metric $P := \Theta^\top \Theta$.

Many existing k -contraction results, including Theorem 2.4, are expressed in terms of matrix measures. An overview of their definitions and properties can be found in [35, Section 2.2]. Theorem 2.4 provides a sufficient condition for verifying k -contraction in the 2-norm with respect to a metric P . The 2-norm was chosen for this work due to its relationship with the eigenvalues of its argument, but other norms could be chosen. Furthermore, like with Lyapunov analysis, one may apply an invertible linear transformation Θ to $w_i(t, r)$ and the k -contraction analysis may be performed in this new domain whilst implying the same property holds in the original domain. This idea is made clear in [49] for 1-contraction and translates analogously to the k -contraction case. When using such an invertible transformation, the system is said to be k -contracting with respect to the metric $P = \Theta^\top \Theta$.

2.2 Forced Lurie Systems and Relevant Special Cases

This section presents the fundamental continuous-time model studied throughout this thesis. The discrete-time counterpart is defined analogously. The most general form of the model, the *Forced Lurie System* [50], is presented first and then the special cases studied throughout the thesis are highlighted after. Conditions which ensure the system is well-posed are then presented, followed by some relevant examples.

2.2.1 Forced Lurie Systems

Consider the interconnection in Figure 2.5, where $P(s) \in \mathcal{RH}_\infty$ is an LTI system with state space realisation (A, B, C, D) and $\Phi : \mathbb{R}^m \rightarrow \mathbb{R}^m$ is a static nonlinearity. The system is modelled by

$$\dot{x} = Ax + B_u u + v_1 \quad y = Cx + D_u u + v_2 \quad u = \Phi(y) \quad (2.65)$$

where $A \in \mathbb{R}^{n \times n}$; $B := [B_u, I_n, 0] \in \mathbb{R}^{n \times (n+2m)}$ with $B_u \in \mathbb{R}^{n \times m}$; $C \in \mathbb{R}^{m \times n}$; and $D := [D_u, 0, I_m] \in \mathbb{R}^{m \times (n+2m)}$ with $D_u \in \mathbb{R}^{m \times m}$. Such an interconnection is a type of *Forced Lurie System*¹. The structured form of the B and D matrices accommodates the three explicitly defined inputs to the LTI system, where the two external inputs, v_1 and v_2 , are respectively applied to the state equation and the output equation.

2.2.1.1 Lurie Systems

An important special case of (2.65) is the unforced system, when $v_1(t) \equiv v_2(t) \equiv 0$, and is referred to as a *Lurie system* throughout the thesis. As the external inputs have been removed, the system becomes time-invariant. In this case, the B and D matrices reduce to $B = B_u \in \mathbb{R}^{n \times m}$ and $D = D_u \in \mathbb{R}^{m \times m}$. For the system to be *well-posed*, a unique solution $x(t)$ must exist for every $x(0)$. This requires some mild assumptions about $\Phi(\cdot)$ and for the implicit equation $y = Cx + D\Phi(y)$ to have a unique solution, y , for every x . These implicit equations are sometimes called *algebraic loops* which frequently arise in control problems; more general results on these can be found in [51; 52; 53]. Section 2.2.2 discusses the conditions which ensure well-posedness is guaranteed.

Lurie systems are the focus of Chapter 3 - Chapter 5. As discussed in these chapters, such systems are prevalent in nature and engineering. Those involving recurrent neural networks (RNNs) and feed-forward neural networks (NNs) are of particular interest in this thesis, with the feed-forward example discussed in Section 2.2.3.

¹Named after Anatolii Isakovitch Lurie and sometimes spelt Lur'e or Lurye.




FIGURE 2.5: Block diagram of Forced Lurie System.

2.2.1.2 Lurie Networks

Another important special case of (2.65) is referred to as the *Lurie network* and occurs when $D_u = 0$ and the external inputs are replaced by constant bias terms, $v_1(t) \equiv b_1$ and $v_2(t) \equiv b_2$. This special case is also time invariant and since $D_u = 0$, the output equation is not implicit; hence only mild conditions on $\Phi(\cdot)$ are required to ensure the Lurie network is well-posed.

Lurie networks are the focus of Chapter 6. As discussed in that chapter, many artificial and biological learning systems are instances of this structure and hence, are of interest in this thesis.

2.2.2 Well-posedness

Well-posedness of the Lurie system and Lurie network are equivalent to the existence of a unique solution to the state space equations (2.65) for the respective special cases. Assuming $\Phi(\cdot)$ is globally Lipschitz (and differentiable almost everywhere), the Lurie network is immediately well-posed, since $D_u = 0$ [14, Chapter 3]. When $D \neq 0$, the Lurie system additionally requires the existence of a unique solution, $y(\eta)$, to the equation $F(y) := y - D\Phi(y) = \eta$. A sufficient condition for this is given by Lemma 2.1 and adapted from, for example, [54, Section II].

Lemma 2.1. *Assuming $\Phi(\cdot)$ is globally Lipschitz, a unique solution, $x(t)$, exists to (2.65) with $v_1(t) \equiv v_2(t) \equiv 0$ and $D = D_u$ for all $x(0) \in \mathbb{R}^n$ and $t \geq 0$ if there exists $\mathbf{U} \in \mathbb{D}_+^m$ such that*

$$2\mathbf{U} - \mathbf{U}D - D^\top \mathbf{U} \succ 0 \quad (2.66)$$

An analogous result is required to ensure well-posedness of discrete-time Lurie systems, where the condition (2.66) is identical. As will be shown in the later chapters, for many Lurie systems of interest, such as those involving NNs, this condition is naturally satisfied.

2.2.3 Lurie Systems Involving Feed-forward Neural Networks

This section presents the interconnection of a feed-forward NN and an LTI dynamical system in the Lurie system framework. The derivation is adapted from [55].

The continuous-time LTI system, $G(s) \in \mathcal{RH}_\infty$, with state space realisation $(A, B, C, 0)$; and, the L -layer feed-forward NN with zero biases, are expressed by (2.67) and depicted in Figure 2.6. The NN has zero biases to ensure the origin is an equilibrium point of the system. The nonlinearity $\Phi_i(\cdot) : \mathbb{R}^K \rightarrow \mathbb{R}^K$ applies the activation function $\phi : \mathbb{R} \rightarrow \mathbb{R}$ element-wise to the K neurons of layer i . That is $\Phi_i(\cdot) = [\phi(z_1), \dots, \phi(z_K)]^\top$.

$$G : \begin{cases} \dot{x} &= Ax + Bu \\ y &= Cx \end{cases} \quad \text{NN} : \begin{cases} w_0 &= y \\ v_i &= W_{i-1}w_{i-1} \\ w_i &= \Phi_i(v_i) \\ u &= W_L w_L \end{cases} \quad (2.67)$$

The next step is to express (2.67) as a dynamical system interconnected with a static nonlinearity. To this end, the following definitions are made

$$v := \begin{bmatrix} v_1 \\ \vdots \\ v_L \end{bmatrix} \quad w := \begin{bmatrix} w_1 \\ \vdots \\ w_L \end{bmatrix} \quad \Phi(v) := \begin{bmatrix} \Phi_1(v_1) \\ \vdots \\ \Phi_L(v_L) \end{bmatrix} \quad (2.68)$$

Using these definitions

$$\begin{bmatrix} v_1 \\ v_2 \\ v_3 \\ \vdots \\ v_{L-1} \\ v_L \end{bmatrix} = \underbrace{\begin{bmatrix} 0 & 0 & 0 & \dots & 0 & 0 \\ W_1 & 0 & 0 & \dots & 0 & 0 \\ 0 & W_2 & 0 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \ddots & 0 & 0 \\ 0 & 0 & 0 & \dots & W_{L-1} & 0 \end{bmatrix}}_{=:N} \begin{bmatrix} w_1 \\ w_2 \\ w_3 \\ \vdots \\ w_{L-1} \\ w_L \end{bmatrix} + \underbrace{\begin{bmatrix} W_0 \\ 0 \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix}}_{=:N_0} y \quad (2.69)$$

$$u = \underbrace{\begin{bmatrix} 0 & 0 & 0 & \dots & 0 & W_L \end{bmatrix}}_{=:N_L} w$$

Thus, the NN can be expressed more compactly by




FIGURE 2.6: Interconnection of a feed-forward neural network and an LTI system.

$$\text{NN} : \begin{cases} v &= N_0 y + N w \\ w &= \Phi(v) \\ u &= N_L w \end{cases} \quad (2.70)$$

In Lurie form, v must be the output and w must be the input of the dynamical system as they are, respectively, the input and output of $\Phi(\cdot)$. This is achieved by substituting u of (2.69) into G of (2.67) and substituting y of (2.67) into v of (2.70) to get

$$P : \begin{cases} \dot{x} &= Ax + BN_L w \\ v &= N_0 Cx + N w \end{cases} \quad (2.71)$$

This is a Lurie system with a state space realisation $(A, BN_L, N_0 C, N)$. It has the same form as (2.65) with $B_u = BN_L$, $D_u = N$, $w = \Phi(v)$ and $v_1 \equiv v_2 \equiv 0$.

2.2.4 Relationship Between Lurie Networks and Neural ODEs

Consider the vector field $\dot{z} = f(z)$ represented by the following L -layer feed-forward network (also referred to as a neural ODE)

$$\begin{aligned} \dot{z} &= W_L \Phi(u_{L-1}) + b_L \\ u_{L-1} &= W_{L-1} \Phi(u_{L-2}) + b_{L-1} \\ u_{L-2} &= W_{L-2} \Phi(u_{L-3}) + b_{L-2} \\ &\vdots \\ u_3 &= W_3 \Phi(u_2) + b_3 \\ u_2 &= W_2 \Phi(u_1) + b_2 \\ u_1 &= W_1 z + b_1 \end{aligned} \quad (2.72)$$

This representation was first proposed in [56]. To illustrate the superior expressivity of the Lurie network, it is now shown how a special case of (2.65) can approximate the neural ODE (2.72). An alternative expression for (2.72) is

$$\begin{aligned}
\dot{z} &= 0z + W_L\Phi(u_{L-1}) + b_L \\
\epsilon\dot{u}_{L-1} &= -u_{L-1} + W_{L-1}\Phi(u_{L-2}) + b_{L-1} \\
\epsilon\dot{u}_{L-2} &= -u_{L-2} + W_{L-2}\Phi(u_{L-3}) + b_{L-2} \\
&\vdots \\
\epsilon\dot{u}_3 &= -u_3 + W_3\Phi(u_2) + b_3 \\
\epsilon\dot{u}_2 &= -u_2 + W_2\Phi(u_1) + b_2 \\
u_1 &= W_1z + b_1
\end{aligned} \tag{2.73}$$

where $\epsilon \rightarrow 0$. Defining a new state $x := [z, u_{L-1}, u_{L-2}, \dots, u_3, u_2]^\top$ and an output vector $y := [u_{L-1}, u_{L-2}, \dots, u_3, u_2, u_1]$ it is clear that (2.73) is a special case of (2.65) with $D_u = 0$, $v_1 = b_x$, $v_2 = b_y$ and the other state space matrices defined by the sparse structures below.

$$A = \epsilon^{-1} \begin{bmatrix} 0 & 0 & 0 & \dots & 0 & 0 \\ 0 & -I & 0 & \dots & 0 & 0 \\ 0 & 0 & -I & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & -I & 0 \\ 0 & 0 & 0 & \dots & 0 & -I \end{bmatrix} \quad B_u = \epsilon^{-1} \begin{bmatrix} \epsilon W_L & 0 & 0 & \dots & 0 & 0 \\ 0 & W_{L-1} & 0 & \dots & 0 & 0 \\ 0 & 0 & W_{L-2} & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & W_3 & 0 \\ 0 & 0 & 0 & \dots & 0 & W_2 \end{bmatrix} \tag{2.74}$$

$$C = \begin{bmatrix} 0 & I & 0 & \dots & 0 & 0 \\ 0 & 0 & I & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & I & 0 \\ 0 & 0 & 0 & \dots & 0 & I \\ W_1 & 0 & 0 & \dots & 0 & 0 \end{bmatrix} \quad b_y = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 0 \\ b_1 \end{bmatrix} \quad b_x = \epsilon^{-1} \begin{bmatrix} \epsilon b_L \\ b_{L-1} \\ b_{L-2} \\ \vdots \\ b_3 \\ b_2 \end{bmatrix} \tag{2.75}$$

This is just one realisation of a Lurie network which can approximate a neural ODE, when ϵ is appropriately set. Other permutations of the state would result in different realisations of the Lurie networks weights and biases. Finally, due to the division by ϵ , it would not be possible to train a Lurie network with the exact same form as (2.73);

however, this analysis shows that it is possible to approximate the structure of a neural ODE with a Lurie network.

2.3 Absolute Stability

The absolute stability problem is a foundational question in nonlinear control theory that is concerned with the stability of forced Lurie systems (Section 2.2.1). The problem arises in the analysis of a feedback interconnection between a linear time-invariant (LTI) system and a, potentially time-varying, memoryless nonlinearity that adheres to certain conditions. The objective is to determine whether the closed-loop system satisfies a chosen definition of stability (e.g., Section 2.1.1.2 or Section 2.1.2) for all nonlinearities within a class, independent of the specific functional form of the nonlinearity. References to a small selection of papers on the topic are listed here [57; 58; 59; 50; 14; 60; 61; 62; 63; 64; 65].

The absolute stability problem is motivated by practical control scenarios where the linear part of the system (often the plant or controller) is well understood, but the nonlinearity may only be known to satisfy broad constraints, such as sector or slope conditions. A canonical example is the Lurie system (Figure 2.5 with $v_1(t) \equiv v_2(t) \equiv 0$), which consists of an LTI system in feedback with a nonlinearity constrained to a sector, determined by the lines $\phi_1(y) = k_1y$ and $\phi_2(y) = k_2y$, and denoted by $\text{Sector}[k_1, k_2]$. The absolute stability problem then asks

Does the closed-loop system satisfy the chosen definition of stability for all nonlinearities within this sector?

The domain which such a definition of stability can be shown to hold depends upon the domain which the sector condition holds for. An example is displayed in Figure 2.7 which indicates the sector bound only holds over the local domain $y \in [-\frac{\pi}{2}, \frac{\pi}{2}]$.

Characterising the class of nonlinearities which the absolute stability criterion is required to hold for is a key step. One common way to do this is with the use of quadratic constraints (QCs), as discussed next.

2.3.1 Common Quadratic Constraints

Sector-bounded and slope-restricted nonlinearities represent two broad classes which commonly appear in absolute stability problems. A scalar function $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is, respectively, sector-bounded or slope-restricted over some domain, \mathcal{D} , if the following inequalities hold




FIGURE 2.7: Nonlinearity satisfying the Sector[0, 1] conditions over $y \in [-\frac{\pi}{2}, \frac{\pi}{2}]$.

$$\phi \in \text{Sector}[k_1, k_2] \Rightarrow k_1 y \leq \phi(y) \leq k_2 y \quad (2.76)$$

$$\phi \in \text{Slope}[k_1, k_2] \Rightarrow k_1 \leq \frac{\phi(y) - \phi(\tilde{y})}{y - \tilde{y}} \leq k_2 \quad (2.77)$$

for all $y \neq \tilde{y} \in \mathcal{D}$ where $k_1 \leq k_2$. Many neural network activation functions satisfy these properties globally, as highlighted in Table 2.1. Two illustrative examples are also presented in Figure 2.7 and Figure 2.8.

These scalar sector and slope inequalities can be used to construct sector-bounded and slope-restricted QCs for repeated nonlinearities $\Phi : \mathbb{R}^m \rightarrow \mathbb{R}^m$ with the form

$$\Phi(\cdot) = \begin{bmatrix} \phi(\cdot) \\ \vdots \\ \phi(\cdot) \end{bmatrix} \quad (2.78)$$

Fact 2.10. Let $\Phi : \mathbb{R}^m \rightarrow \mathbb{R}^m$ be a repeated nonlinearity and $K_1, K_2 \in \mathbb{D}$ with $K_2 - K_1 \in \mathbb{D}_+^m$. If there exists $\mathbf{V} \in \mathbb{D}_+^m$ such that

$$[\Phi(y) - K_1 y]^\top \mathbf{V} [K_2 y - \Phi(y)] \geq 0 \quad \forall y \in \mathcal{D} \quad (2.79)$$

then $\Phi(y) \in \text{Sector}[K_1, K_2]$ over the domain \mathcal{D} .

Proof: Assuming $\phi(\cdot) \in \text{Sector}[k_{1,i}, k_{2,i}]$, then the following two inequalities can be extracted from (2.76)

$$\phi(y_i) - k_{1,i} y_i \geq 0 \quad k_{2,i} y_i - \phi(y_i) \geq 0 \quad (2.80)$$

TABLE 2.1: Sector-bounded and slope-restricted activation functions over \mathbb{R} .

Activation Function	Sector-bound ($[k_1, k_2]$)	Slope-restriction ($[k_1, k_2]$)
ReLU	$[0, 1]$	$[0, 1]$
Leaky ReLU	$[0, 1]$	$[0, 1]$
Shifted Sigmoid	$[0, 1]$	$[0, 0.25]$
Tanh	$[0, 1]$	$[0, 1]$

Multiplying the two inequalities together and by $v_{ii} \geq 0$ results in

$$(\phi(y_i) - k_{1,i}y_i)v_{ii}(k_{2,i}y_i - \phi(y_i)) \geq 0 \quad (2.81)$$

Assuming $\Phi(\cdot)$ is a repeated nonlinearity constructed from $\phi(\cdot)$ then, summing over $i = 1, \dots, m$ and expressing in the equivalent quadratic form leads to Fact 2.10. Note that $K_j = \text{diag}(k_{j,1}, \dots, k_{j,m})$ for $j \in \{1, 2\}$ and $\mathbf{V} = \text{diag}(v_{11}, \dots, v_{mm})$. \square

A similar fact can be obtained for slope-restricted repeated nonlinearities.

Fact 2.11. Let $\Phi : \mathbb{R}^m \rightarrow \mathbb{R}^m$ be a repeated nonlinearity, $\Psi(y, \tilde{y}) := \Phi(y) - \Phi(\tilde{y})$ and $K_1, K_2 \in \mathbb{D}$ with $K_2 - K_1 \in \mathbb{D}_+^m$. If there exists $\mathbf{W} \in \mathbb{D}_+^m$ such that

$$[\Psi(y, \tilde{y}) - K_1(y - \tilde{y})]^\top \mathbf{W} [K_2(y - \tilde{y}) - \Psi(y, \tilde{y})] \geq 0 \quad \forall y \neq \tilde{y} \in \mathcal{D} \quad (2.82)$$

then $\Phi(y) \in \text{Slope}[K_1, K_2]$ over the domain \mathcal{D} .

Proof: Assuming $\phi(\cdot) \in \text{Slope}[k_{1,i}, k_{2,i}]$, then the following two inequalities can be extracted from (2.77)

$$\psi(y_i, \tilde{y}_i) - k_{1,i}(y_i - \tilde{y}_i) \geq 0 \quad k_{2,i}(y_i - \tilde{y}_i) - \psi(y_i, \tilde{y}_i) \geq 0 \quad (2.83)$$

Multiplying the two inequalities together and by $w_{ii} \geq 0$ results in

$$(\psi(y_i, \tilde{y}_i) - k_{1,i}(y_i - \tilde{y}_i))w_{ii}(k_{2,i}(y_i - \tilde{y}_i) - \psi(y_i, \tilde{y}_i)) \geq 0 \quad (2.84)$$

Assuming $\Phi(\cdot)$ is a repeated nonlinearity constructed from $\phi(\cdot)$ then, summing over $i = 1, \dots, m$ and expressing in the equivalent quadratic form leads to Fact 2.11. Note that $K_j = \text{diag}(k_{j,1}, \dots, k_{j,m})$ for $j \in \{1, 2\}$ and $\mathbf{W} = \text{diag}(w_{11}, \dots, w_{mm})$. \square

Both the sector-bounded and slope-restricted QCs include variables (\mathbf{V} and \mathbf{W}) which can later be determined. This freedom can be leveraged when posing stability criteria




FIGURE 2.8: Nonlinearity satisfying the Slope $[-2, 2]$ conditions over $y \in [-1, 2]$.

(Section 2.1.2) as a semi-definite programming (SDP) problem (Section 2.1.3), as will be illustrated in the next section.

2.3.2 S-procedure

The S-procedure is a classical result from control theory and convex analysis that provides sufficient conditions under which a quadratic inequality can be inferred from another [38, Section 2.6.3]. Specifically, it addresses the problem of determining when an implication of the form

$$g(x) \geq 0 \quad \Rightarrow \quad f(x) \geq 0 \quad (2.85)$$

holds for all $x \in \mathbb{R}^n$ when f, g are quadratic functions. The S-procedure introduces a scalar multiplier, often referred to as an S-procedure multiplier, to transform such an implication into a single inequality involving a linear combination of f and g . Formally, the S-procedure states that if there exists a scalar $\lambda \geq 0$ such that

$$f(x) - \lambda g(x) \geq 0 \quad \forall x \in \mathbb{R}^n \quad (2.86)$$

then this is sufficient for the implication $g(x) \geq 0 \Rightarrow f(x) \geq 0$ to hold.

The significance of the S-procedure lies in its ability to convert non-convex implication constraints into convex semidefinite constraints, enabling the use of powerful tools from semidefinite programming (Section 2.1.3). In this thesis, the S-procedure is regularly used to form a conservative approximation, $g(x)$, of the constraint of interest, $f(x)$. A classical result from absolute stability analysis, the Circle Criterion [14], is studied below to illustrate this.

Example 2.2. Consider a Lurie system of the form (2.65) with $v_1(t) \equiv v_2(t) \equiv 0$, $D = 0$ and $\Phi(y) \in \text{Sector}[0, I]$ for all $y \in \mathbb{R}^m$. The Circle Criterion provides a sufficient condition to test if such a Lurie system is globally asymptotically stable (GAS).

Following Theorem 2.2, the Circle Criterion proposes a quadratic Lyapunov function of the form $V(x) = x^\top \mathbf{P}x$ with $\mathbf{P} \in \mathbf{S}_+^n$. Calculating $\dot{V}(x)$ results in

$$\dot{V}(x) = \begin{bmatrix} x \\ \Phi \end{bmatrix}^\top \begin{bmatrix} \mathbf{P}A + A^\top \mathbf{P} & \mathbf{P}B \\ B^\top \mathbf{P} & 0 \end{bmatrix} \begin{bmatrix} x \\ \Phi \end{bmatrix} \quad (2.87)$$

Checking if $\dot{V}(x) < 0 \forall x \neq 0 \in \mathbb{R}^n$ is equivalent to verifying if there exists $\mathbf{P} \in \mathbf{S}_+^n$ which satisfies the LMI

$$\begin{bmatrix} \mathbf{P}A + A^\top \mathbf{P} & \mathbf{P}B \\ * & 0 \end{bmatrix} \prec 0 \quad (2.88)$$

This inequality can be efficiently solved (if a solution exists) using SDP (Section 2.1.3).

As highlighted by a Schur complements argument [66, Section 2.3], a solution $\mathbf{P} \in \mathbf{S}_+^n$ cannot be found to satisfy this LMI since the (2,2) block can never be negative definite. At the expense of some conservatism, the Circle Criterion leverages the S-procedure to handle this issue by adding Fact 2.10 to $\dot{V}(x)$ to get

$$g(x) = \dot{V}(x) + \underbrace{2\Phi(y)^\top \mathbf{V}[y - \Phi(y)]}_{\geq 0} \quad (2.89)$$

Since the second term is non-negative, $g(x) < 0 \Rightarrow \dot{V}(x) < 0 \forall x \neq 0 \in \mathbb{R}^n$. Checking if $g(x) < 0 \forall x \neq 0 \in \mathbb{R}^n$ is equivalent to verifying if there exists $\mathbf{P} \in \mathbf{S}_+^n$ and $\mathbf{V} \in \mathbf{D}_+^m$ which satisfy the LMI

$$\begin{bmatrix} \mathbf{P}A + A^\top \mathbf{P} & \mathbf{P}B + C^\top \mathbf{V} \\ * & -2\mathbf{V} \end{bmatrix} \prec 0 \quad (2.90)$$

Since the (2,2) block is now negative definite, a Schur complements argument no longer prevents a solution from existing. The existence of a solution will now depend on the Lurie system parameters (A, B, C) .

2.3.3 Literature on Absolute Stability

Over the decades, several analytical techniques have been developed to address the absolute stability problem. Notable contributions include the Circle Criterion, Popov Criterion, Zames-Falb multipliers and integral quadratic constraints (IQCs). These approaches blend frequency-domain and time-domain insights and have played a central

role in shaping modern robust and nonlinear control theory. In the following section, a number of approaches are discussed, followed by a comparison of their respective strengths and limitations. As commonly done throughout this chapter, the continuous-time results will be prioritised in the discussion; however, the discrete-time setting is conceptually the same.

Each approach is classified according to how the method is formulated and the definition of stability used. Approaches derived in the time-domain, via the S-procedure (Section 2.3.2), which result in Lyapunov-based stability criteria (Section 2.1.2) are referred to as *quadratic forms*. A popular frequency domain approach is to study the input-output stability properties (Section 2.1.1.2) of an equivalent system, augmented by a *multiplier*. Thanks to the Kalman-Yakubovich-Popov (KYP) Lemma [14, Lemma 6.3], it is possible to convert a criterion between the two domains. The final class of methods are those which frame the problem as an IQC. This is a powerful approach because it captures existing results from robust control and enables one to combine different stability results, including quadratic forms and multipliers, into a common analysis framework. A recent review of the field is presented in [65]; whereas [67; 68] detail accessible introductions to multipliers and IQCs.

Quadratic Forms

The Circle Criterion [69; 14] was one of the earliest and most intuitive results in absolute stability theory. It was originally developed to verify asymptotic stability for scalar Lurie systems, with a sector-bounded nonlinearity. The criterion had a graphical representation for testing stability, the Nyquist plot, which was the motivation behind the name. A generalisation for testing asymptotic stability of multivariate Lurie systems was later established; however, the intuitive graphical representation could no longer be applied. Instead the criterion was posed as a semi-definite programming (SDP) problem which could be verified computationally (Section 2.3.2). The stability result was posed as a linear matrix inequality (LMI) and was derived in several ways, such as appending the sector-bound quadratic constraint (Fact 2.10) via the S-procedure in the Lyapunov analysis (Section 2.1.2). This was demonstrated in Example 2.2 which highlights the resulting LMI in (2.90).

The Popov Criterion [69; 14] involves the use of the, more general, Lurie-type Lyapunov function. However, this choice requires the linear system to be *strictly proper*, which equates to assuming $D = 0$ in the linear state space representation. As with the Circle Criterion, the result can be derived in various ways, but is now most commonly presented as an LMI [38]. This included appending the sector-bounded quadratic constraint (QC) via the S-procedure in the Lyapunov analysis. This can easily be seen by following the same steps as Example 2.2 but starting with the Lurie-type Lyapunov function $V(x) = x^\top \mathbf{P}x + 2 \int_0^y \Lambda \Phi(\sigma) d\sigma$. The resulting LMI was given by

$$\begin{bmatrix} \mathbf{P}A + A^\top \mathbf{P} & \mathbf{P}B + C^\top \mathbf{V} + A^\top C^\top \boldsymbol{\Lambda} \\ \star & \boldsymbol{\Lambda}CB + B^\top C^\top \boldsymbol{\Lambda} - 2\mathbf{V} \end{bmatrix} \prec 0 \quad (2.91)$$

with $\mathbf{P} \in \mathbf{S}_+^n$, $\boldsymbol{\Lambda} \in \mathbf{D}_+^m$ and $\mathbf{V} \in \mathbf{D}_+^m$. The LMI associated with the Circle Criterion is a special case of this with $\boldsymbol{\Lambda} = 0$. In [70; 71], the Popov Criterion was also shown to hold for $\boldsymbol{\Lambda} \in \mathbf{D}$; thus increasing the freedom of the SDP decision variables.

An alternative perspective on the absolute stability problem is obtained through the use of the Small Gain Theorem (SGT) [69; 38] which was developed as an analysis tool for studying the input-output stability (Section 2.1.1.2) of two interconnected systems. The theorem states that if both individual systems have a finite L_2 gain, then the interconnected system is input-output stable providing the product of these gains is less than 1. For Lurie systems, a finite L_2 gain of the nonlinearity is equivalent to $\|\phi(y_i)\|^2 \leq \|y_i\|^2$. This is a QC and can be expressed with extra flexibility by multiplying through with the variable $v_{ii} \in \mathbb{R}_{\geq 0}$ [38, Section 5.1]. Summing over $i = 1, \dots, m$ results in the QC

$$\begin{bmatrix} x \\ \Phi(y) \end{bmatrix}^\top \begin{bmatrix} C^\top \mathbf{V}C & C^\top \mathbf{V}D \\ \star & D^\top \mathbf{V}D - \mathbf{V} \end{bmatrix} \begin{bmatrix} x \\ \Phi(y) \end{bmatrix} \geq 0 \quad \forall x \in \mathbb{R}^n \quad (2.92)$$

When appended to the time derivative of the quadratic Lyapunov function, the resulting LMI is given by

$$\begin{bmatrix} \mathbf{P}A + A^\top \mathbf{P} + C^\top \mathbf{V}C & \mathbf{P}B + C^\top \mathbf{V}D \\ \star & D^\top \mathbf{V}D - \mathbf{V} \end{bmatrix} \prec 0 \quad (2.93)$$

where $\mathbf{P} \in \mathbf{S}_+^n$ and $\mathbf{V} \in \mathbf{D}_+^m$.

In addition to the classical approaches discussed above, more recent works such as [72; 73; 74] also fall into this category. The approaches presented in [72; 73] assume the nonlinearity is sector and slope restricted. These properties were leveraged to construct several non-negative integral inequalities which were included in a novel Lyapunov function by adding them as a weighted sum to a quadratic term. Consequently, the associated LMI arising from the negative definite Lyapunov condition was much larger than those previously considered. A fairly similar approach was taken in [74]; however, a more sophisticated quadratic term was used in the Lyapunov function which resulted in an LMI where the \mathbf{P} matrix was relaxed to just being symmetric and the weighting terms of the integrals were no longer required to be positive. Each of these methods derived the corresponding stability criterion in the time-domain and assumed $D = 0$.

Passivity

Passivity is a broad and powerful concept in linear and nonlinear systems theory, with deep connections to stability analysis, control design, and feedback interconnections. Due to its breadth, only a brief overview is presented here, primarily drawing on the treatment in [14, Chapter 6]. This section also forms the basis of the upcoming section on *multipliers*.

A state space system is said to be *passive* if the energy supplied to the system is always greater than or equal to the increase in the energy stored by the system. Formally, for a system with input $u(t)$, state $x(t)$, output $y(t)$, and a non-negative storage function $V(x)$, passivity is defined via the dissipation inequality

$$y^\top(t)u(t) \geq \dot{V}(t) = \frac{\partial V}{\partial x} \cdot \dot{x} \quad (2.94)$$

The system is *strictly passive* if the condition is satisfied with a strict inequality.

In the linear time-invariant (LTI) case, the system is (strictly) passive if the corresponding transfer function is (*strictly*) *positive real*. Specifically, for a scalar system, the real part of the system's frequency response must be non-negative (positive) for all frequencies. This bounds the phase of the frequency response to the (open) interval between ± 90 degrees. Less intuitive generalisations of these conditions exist for multivariate systems.

Passivity-based stability results are typically derived using Lyapunov techniques, where the storage function $V(x)$ serves as a Lyapunov function candidate. Passivity of subsystems plays a central role in establishing the stability of interconnected systems: for instance, the feedback interconnection of two passive systems is guaranteed to be stable under mild technical conditions. An example is detailed in the discussion on multipliers. Whilst passivity provides a convenient framework for studying the stability of system interconnections, it is a strong assumption which limits its applicability. For example, any system with transfer function of relative degree two or more, such as many simple mechanical and electrical systems, do not satisfy this assumption, and neither do systems which exhibit a phase lead beyond 90 degrees.

Many advanced stability results can be interpreted through the lens of passivity via loop transformation techniques. These transformations recast stability problems into equivalent passive feedback interconnection problems, with relaxed assumptions. Notably, tools such as Zames-Falb multipliers, widely used in the analysis of systems with slope-restricted and/or sector-bounded nonlinearities, can often be interpreted in terms of passivity properties.




FIGURE 2.9: Equivalent Lurie system augmented by the multiplier $M(s)$.

Multipliers

To illustrate the multiplier framework, consider the Lurie system from Figure 2.5 with $v_1(t) \equiv v_2(t) \equiv 0$ as the original system of interest. It is well known that, under negative feedback, this system is input-output stable if the linear system, $P(s)$, is strictly passive and the nonlinearity, Φ , is passive [14, Theorem 6.4]. This places a very strong requirement on $P(s)$ which is rarely satisfied in practice. To relax this strong requirement, one can introduce a new transfer function, $M(s)$, known as a *multiplier*, into the loop. Using the same passivity arguments, the augmented but equivalent Lurie system (Figure 2.9) is input-output stable if $\hat{P}(s)$ is strictly passive and the nonlinearity, $\hat{\Phi}$, is passive. As the strictly passive requirement is now on $P(s)M(s)$, where there is freedom to choose $M(s)$, this should be easier to enforce. The challenge is to search for a multiplier which guarantees input-output stability for as large a class of nonlinearities as possible.

O'Shea proposed a class of multipliers suitable for monotone nonlinearities [75; 76] which was later formalised in the pioneering work by Zames and Falb [58]. The class of Zames-Falb (ZF) multipliers is defined below and has a number of appealing properties, as discussed in the review article [61].

Definition 2.7. *The class of Zames-Falb multipliers is given by all transfer functions, $M(s)$, bounded and analytic on the imaginary axis, whose inverse Laplace transform is given by*

$$m(t) = \delta(t) - \sum_{i=1}^{\infty} h_i \delta(t - t_i) - h(t) \quad \|h\|_{\mathcal{L}_1} + \sum_{i=1}^{\infty} \|h_i\| < 1 \quad (2.95)$$

where $\delta(\cdot)$ is the dirac delta function, $h_i \in \mathbb{R}$, and $\|h\|_{\mathcal{L}_1} = \int_{-\infty}^{\infty} \|h(t)\| dt$.

Zames-Falb multipliers can be applied to Lurie systems with monotone nonlinearities; if the class of nonlinearities isn't odd, only a subclass of multipliers can be applied. Most recent work has focused on how to search for multipliers which either maximise the slope size for which stability is guaranteed or minimise the \mathcal{L}_2 gain [61]. Such searches are typically restricted to the rational subclass $M \in \mathcal{RL}_{\infty}$, where $h_i = 0$ for all i [77; 78; 79; 80; 81; 82]. As discussed in [68, Section 5], the searches above, as well

as others, can be further classified into three main approaches. Those which choose rational multipliers where $H(s)$ has the following form

$$H(s) = \prod_{i=1}^n \frac{k_i}{s + a_i} \quad (2.96)$$

These multipliers require the order, n , and the poles, a_i , to be chosen such that the arising matrix inequalities are linear. Other approaches assume the order of the multiplier is at least the order of the plant but require an upper bound on the \mathcal{L}_1 norm (Definition 2.7) which is known to be extremely conservative in some cases. A final set of approaches are those which use irrational multipliers; whilst these are promising, they are currently quite fragile and are challenging to combine with synthesis techniques.

One limitation common to most multiplier approaches is that they are typically posed for scalar Lurie systems. Generalisations to multivariate Lurie systems are not trivial and require further assumptions about the nonlinearity to be made (e.g., it must be the derivative of a convex potential function [61]). Even more problematic for high-dimensional Lurie systems is that the searches become significantly more difficult and computationally demanding. Promising methods in this regard, such as [83], are still computationally demanding and require most parameters of (2.96) to be chosen without much guidance. Of course, this reduces the freedom to find a multiplier which satisfies the required passivity conditions needed to guarantee input-output stability of the Lurie system. Other attempts to address this limitation include [84; 85; 86; 87].

Integral Quadratic Constraints

For completeness, a brief discussion on integral quadratic constraints (IQC) is included, although they do not play a role in this thesis. Unlike Lyapunov analysis, which can be applied to any nonlinear system, the IQC approach typically restricts attention to the analysis of systems which can be expressed as a feedback interconnection of one known and “well-behaved” element, and a block of other elements which are less precisely known. A Lurie system is a good example. The result was first presented in [59], but was strongly influenced by the work of Yakubovich [57]. Similar to quadratic constraints, IQCs characterise a class of operator’s in terms of quadratic forms. Example operators include nonlinearities, delays, time-varying coefficients, and uncertain linear time invariant (LTI) dynamics.

IQCs are typically presented in the frequency domain (although they can be formulated in the time domain [88]) and can be directly applied to obtain an input-output stability condition, also in the frequency domain. With respect to Figure 2.5, the nonlinearity, Φ , is said to satisfy the IQC defined by $\Pi \in \mathcal{RL}_\infty$ if the following inequality holds

$$\int_{-\infty}^{\infty} \begin{bmatrix} Y(jw) \\ U(jw) \end{bmatrix}^* \Pi(jw) \begin{bmatrix} Y(jw) \\ U(jw) \end{bmatrix} dw \geq 0 \quad (2.97)$$

TABLE 2.2: Criteria comparison.

Criteria	Nonlinearity	$D \neq 0$	Variable Count	LMI Count
Circle	Sector	✓	$0.5n(n+1) + m$	3
Popov	Sector	✗	$0.5n(n+1) + 2m$	4
SGT	Norm-bounded	✓	$0.5n(n+1) + m$	3
Park [72]	Sector, Slope	✗	$0.5(n+m)(n+m+1) + 10m$	12
ZF ²	Slope, Sector, Repeated	✗	$2.5n(n+1) + 2n + m^2$	12

where $Y(jw), U(jw)$ are the respective Fourier transforms of $y(t), u(t)$ and $\Pi(jw) = \Pi^*(jw)$. One can easily obtain an IQC from a QC, such as (2.79), by integrating over the time domain and subsequently applying Parseval's identity. The benefit of establishing an IQC of the form (2.97) is that it can be directly applied to verify input-output stability in the frequency domain. Assuming the Lurie system in Figure 2.5 has $v_1 \equiv v_2 \equiv 0$ and $P(s) \in \mathcal{RH}_\infty$, then the Lurie system is input-output stable if

$$\begin{bmatrix} P(jw) \\ I \end{bmatrix}^* \Pi(jw) \begin{bmatrix} P(jw) \\ I \end{bmatrix} \leq -\epsilon I \quad \forall w \in \mathbb{R} \quad (2.98)$$

By application of the KYP Lemma, it follows that (2.98) can be expressed as a linear matrix inequality (LMI) in terms of $\mathbf{P} \in \mathcal{S}_+^n$ and the other matrix variables which arise from the IQC characterisation. Such problems can be posed as a semidefinite programming (SDP) problem and solved efficiently (Section 2.1.3).

The key step in IQC analysis is characterising the troublesome nonlinear, time-varying, or uncertain component. Several IQCs were established in [59] and, more recently, IQCs characterising neural network nonlinearities were presented in [89]. Much work has also focused on formulating existing absolute stability results in the IQC framework [90; 60; 62; 64]. Whilst this provides a unifying framework for comparing quadratic forms and multiplier stability criteria, a more practical benefit is gained when multiple troublesome components are present in a system. As shown in [67], a single IQC can be constructed from the individual IQCs which characterise each of the troublesome components.

Comparison

When studying asymptotic or input-output stability of Lurie systems involving neural networks (NNs), the resulting LMIs become much larger than those traditionally studied in the literature. For this reason, criteria with fewer variables and lower conservatism are desirable to minimise the computation required and to avoid wasting computation (e.g., the system is stable but the criterion cannot verify this). With this in mind, some relevant properties of the criteria presented earlier in this section are detailed in Table 2.2.

²[91] used as ZF choice.

First, the assumptions made about the nonlinearity were considered. With reference to Table 2.1, it is clear that many activation functions of interest are sector-bounded and slope-restricted. In artificial NNs the user chooses the nonlinearity, so the repeated property can be satisfied by design. This is also common in models of biological NNs and is leveraged in this thesis.

Secondly, assumptions made about the LTI system were considered. For example, only two criteria can deal with the case when $D \neq 0$, which is limiting when studying systems involving feed-forward NNs. It should be noted that whilst the considered ZF approach [91] assumes $D = 0$, this is not a general requirement for ZF methods. Additionally, whilst most of the criteria considered rely on passivity to motivate the choice of Lyapunov function, as will be shown later in the thesis, a more computational approach may also be taken (i.e., directly trying to reduce conservatism in the LMIs by introducing new variables and relaxing constraints on existing ones).

A brief comparison of the computational demands was made. It is understood that the dimension of the LMIs (dependent on the number and size of each LMI) provides the most significant contribution to slowing down the LMI solver [38]. It should be noted that the criteria in [72; 91] not only contain many more LMIs, but also some of the LMIs involved were significantly larger than those in the other approaches. However, it should be emphasised that there are a lot of choices to be made in ZF analysis, so the number of variables and LMIs are only accurate in the case of [91]. In terms of the number of variables, [72] and [91] are quadratic in m (number of activation functions), whereas the other approaches are linear; hence, [72; 91] are considerably more computationally demanding.

Finally, this section is concluded by highlighting some miscellaneous factors which should be considered when choosing a stability criterion. For example, quadratic form approaches tend to be more systematic than multiplier approaches, which often require intuition about the problem when choosing parameter values. All criteria focus on stability of equilibrium points or well-behaved signals; there are no global results for systems with multiple equilibrium points or limit cycles. Finally, the IQC approach is elegant and unrivalled when dealing with multiple uncertainties; however, deriving IQCs is perhaps easier once a time-domain quadratic constraint has been established.

Chapter 3

Stability Analysis of Lurie Systems with ReLU Nonlinearities

This chapter considers the stability analysis of a Lurie system with a static repeated ReLU (rectified linear unit) nonlinearity. Properties of the ReLU function are leveraged to derive new tailored quadratic constraints (QCs) which are satisfied by the repeated ReLU. These QCs are used to strengthen the Circle and Popov Criteria for this specialised Lurie system. It is shown that the criteria can be cast as a set of linear matrix inequalities (LMIs) with less restrictive conditions on the matrix variables. Many systems involving a neural network (NN) with ReLU activations are important instances of this specialised Lurie system; for example, a Hopfield network or the interconnection of a linear system with a feed-forward NN. Numerical examples demonstrate the reduced conservatism of the strengthened criteria, particularly for large systems.

3.1 Introduction

As mentioned in Chapters 1 and 2, various systems involving NNs can be modelled as a Lurie system (Figure 3.1) where the nonlinearity, $\Phi(\cdot)$, is a vector of the NN activation functions. Examples include the interconnection of a linear time-invariant (LTI) system with a feed-forward NN (Section 2.2.3) and a Hopfield network [92]. Section 2.3 discussed how a Lurie system can be analysed with a range of criteria from absolute stability: the classical Circle and Popov Criteria [14; 93; 71], other Lyapunov-based criteria [70; 72] and Zames-Falb multipliers [75; 58; 61; 94]. Typically multipliers are limited to single-input-single-output systems; however, there is some work which extends the multiplier framework to repeated nonlinearities [84; 85; 86; 87]. Limitations of these approaches are discussed in Section 2.3. Each of these criteria are posed as semi-definite programming (SDP) problems involving linear matrix inequalities (LMIs); the benefits and implementations of SDPs are discussed in Section 2.1.3. What differentiates




FIGURE 3.1: Lurie system with static nonlinearity.

the criteria is how they balance the trade-off between conservatism and computational complexity.

Recent work has used the absolute stability framework, and the associated SDP tools, for addressing a number of problems in NN analysis: estimation of the region of attraction [20; 95], synthesis of NN controllers [21; 96], and robustness analysis [97; 98; 99]. The main challenge in NN analysis is that the number of activation functions, m , is typically large, meaning that the resulting absolute stability problems suffer from greater computational complexity than traditional absolute stability problems where m is normally small. Furthermore, at this scale, even some of the less conservative absolute stability tools, such as Zames-Falb multiplier analysis, become quite conservative, limiting their practical use.

Contribution: The main contribution of this chapter is to derive less conservative stability criteria which are computationally tractable for high-dimensional problems, such as those involving NNs. To do this, Lurie systems where the nonlinearity is of the ReLU type are focused on. ReLU is a popular choice of activation function in deep learning as it does not suffer from saturated outputs or increased computational inefficiency, unlike the Logistic sigmoid and Tanh functions [100]. Properties of the ReLU function are leveraged to derive tailored QCs for the repeated ReLU. As these characterise the nonlinearity more accurately than sector/slope bounds, the stability analysis holds for fewer nonlinearities. Generality is therefore sacrificed to increase the freedom in the SDP optimisation and, as a consequence, reduce the conservatism of the stability analysis. Thus, the contribution is three-fold: tailored QCs are derived for the repeated ReLU; the low complexity Circle and Popov Criteria are strengthened for this specialised Lurie system; and, convex relaxations are proposed for converting the strengthened Popov Criterion into an SDP problem involving LMIs, which can be solved efficiently.

Chapter structure: Section 3.2 presents the problem setup and highlights the properties of the ReLU function. Section 3.3 derives the tailored QCs for the repeated ReLU. Section 3.4 presents the strengthened Circle and Popov Criteria for this specialised Lurie

system, these are referred to as the Circle-like and Popov-like Criteria. Section 3.5 proposes convex relaxations for the Popov-like Criterion. Finally, Section 3.6 presents some numerical examples.

3.2 Preliminaries

3.2.1 Problem Setup

Consider the interconnection in Figure 3.1, where $P(s) \in \mathcal{RH}_\infty$ is a finite dimensional LTI system, with state space realisation (A, B, C, D) ¹. The system is modelled by (3.1) with $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, $C \in \mathbb{R}^{m \times n}$, $D \in \mathbb{R}^{m \times m}$ and $\Phi(\cdot) : \mathbb{R}^m \rightarrow \mathbb{R}^m$ being the repeated ReLU nonlinearity. Such an interconnection is a type of *Lurie system*. As the repeated ReLU satisfies $\Phi(0) = 0$, the origin is an equilibrium point² of (3.1).

$$\begin{aligned} \dot{x} &= Ax + B\Phi(y) \\ y &= Cx + D\Phi(y) \end{aligned} \tag{3.1}$$

The following assumption is made throughout the chapter.

Assumption 3.1 (Well-posedness). *A unique solution $x(t)$ exists to (3.1) for all $x(0) \in \mathbb{R}^n$ and all $t \geq 0$.*

A discussion on well-posedness is included in Section 2.2.2. For this chapter, it is sufficient to say that well-posedness is equivalent to the existence of a unique solution to the state space equations (3.1). Since $\Phi(\cdot)$ is globally Lipschitz (and differentiable almost everywhere), a sufficient condition is given by Lemma 2.1.

In many absolute stability results (e.g., the standard Circle Criterion and Zames-Falb multipliers) the LMI (2.66), in Lemma 2.1, is an intrinsic part of the stability conditions and thus well-posedness is guaranteed. For the new results presented in this chapter, this is not the case, so well-posedness needs to be verified by other means. Generally, this involves verifying (2.66) directly. Fortunately, this is straightforward for the case of many NN stability problems. For example, in an L -layer feed-forward NN, the D -matrix in (3.1) takes the following form (Section 2.2.3)

¹Section 2.1.1 discusses the fundamentals of signals and systems.

²Section 2.1.2 discusses Lyapunov stability analysis and defines an equilibrium point.

$$D = \begin{bmatrix} 0 & 0 & \dots & 0 & 0 \\ W_1 & 0 & \dots & 0 & 0 \\ 0 & W_2 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & W_{L-1} & 0 \end{bmatrix} \in \mathbb{R}^{m \times m} \quad (3.2)$$

where $W_i \in \mathbb{R}^{m_i \times m_i}$ are weight matrices. In this case, inequality (2.66) becomes

$$\begin{bmatrix} 2\mathbf{U}_1 & -W_1' \mathbf{U}_2 & 0 & \dots & 0 & 0 \\ -\mathbf{U}_2 W_1 & 2\mathbf{U}_2 & -W_2' \mathbf{U}_3 & \dots & 0 & 0 \\ 0 & -\mathbf{U}_3 W_2 & 2\mathbf{U}_3 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & 2\mathbf{U}_{L-1} & -W_{L-1}' \mathbf{U}_L \\ 0 & 0 & 0 & \dots & -\mathbf{U}_L W_{L-1} & 2\mathbf{U}_L \end{bmatrix} \succ 0 \quad (3.3)$$

where $\mathbf{U}_i \in \mathbb{D}_+^{m_i}$ and $\sum_{i=1}^L m_i = m$. Hence, by a Schur complement argument (see e.g. [38]), for an arbitrary choice of $\mathbf{U}_L \in \mathbb{D}_+^{m_L}$, one can always choose $\mathbf{U}_{L-1} \in \mathbb{D}_+^{m_{L-1}}$ such that the lower-right block matrix is positive definite; the remaining \mathbf{U}_i can then be chosen recursively. Thus, well-posedness for this class of Lurie systems is unconditionally guaranteed.

Similarly, for Hopfield network's, the system (3.1) has the form $A = -I, B = W, C = I, D = 0$ [92] and hence is, trivially, well-posed. A detailed discussion of well-posedness is beyond the scope of the chapter; it suffices to say that many systems featuring NNs are naturally well-posed.

The following problem is addressed in the remainder of the chapter.

Problem 3.1. *Find convex Lyapunov-based conditions which ensure the origin of the Lurie system (3.1) is globally asymptotically stable when $\Phi(\cdot)$ is the repeated ReLU.*

3.2.2 Properties of the ReLU Function

The ReLU function (Figure 3.2) is continuous over its domain. The repeated ReLU is a vectorised version of the ReLU function.




FIGURE 3.2: ReLU function.

Definition 3.1 (ReLU function). $\phi(\cdot) : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$

$$\phi(y_i) = \begin{cases} y_i & y_i \geq 0 \\ 0 & y_i < 0 \end{cases} \quad (3.4)$$

Definition 3.2 (Repeated ReLU). If $\phi(\cdot) : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ is the ReLU function, the repeated ReLU is $\Phi(\cdot) : \mathbb{R}^m \rightarrow \mathbb{R}_{\geq 0}^m$

$$\Phi(\cdot) = \left[\phi(\cdot) \quad \dots \quad \phi(\cdot) \right]^\top \quad (3.5)$$

It is well known that the ReLU function satisfies a number of properties [97; 101]; some of these are summarised in Table 3.1. Although the final two properties in Table 3.1 hold for many common static nonlinearities, the first four are less typical. In fact, the complementarity condition holds for few activation functions other than the ReLU function.

3.3 Quadratic Constraints

This section derives two tailored QCs for the repeated ReLU, using properties from Table 3.1. In the next section, the presented QCs are leveraged to prove the strengthened criteria. Section 2.3.1 presents some well-known QCs and Section 2.3.2 outlines the procedure for leveraging the QCs within the Lyapunov stability analysis.

Fact 3.1 (Sector-like QC). Let $\Phi(\cdot) : \mathbb{R}^m \rightarrow \mathbb{R}_{\geq 0}^m$ be the repeated ReLU. If $\mathbf{V} \in \mathcal{Z}^m$ then the following QC holds

$$\Phi(y)' \mathbf{V} [y - \Phi(y)] \geq 0 \quad \forall y \in \mathbb{R}^m \quad (3.6)$$

TABLE 3.1: Properties satisfied by the ReLU function.

$\phi(y_i) \geq 0$	$\forall y_i \in \mathbb{R}$	positivity
$\phi(\beta y_i) = \beta \phi(y_i)$	$\forall y_i \in \mathbb{R}, \beta \in \mathbb{R}_{\geq 0}$	positive homogeneity
$\phi(y_i) - y_i \geq 0$	$\forall y_i \in \mathbb{R}$	positive complement
$\phi(y_i)(y_i - \phi(y_i)) = 0$	$\forall y_i \in \mathbb{R}$	complementarity
$0 \leq \frac{\phi(y_i)}{y_i} \leq 1$	$\forall y_i \in \mathbb{R}$	sector-boundedness (Sector[0, 1])
$0 \leq \frac{\phi(y_i) - \phi(\tilde{y}_i)}{y_i - \tilde{y}_i} \leq 1$	$\forall y_i, \tilde{y}_i \neq y_i \in \mathbb{R}$	slope-restriction (Slope[0, 1])

Proof: By scaling the product of the positivity and positive complement properties, inequality (3.7) follows. When $i = j$ the less restrictive equation (3.8) is implied.

$$\phi(y_i)v_{ij}(y_j - \phi(y_j)) \geq 0 \quad \forall y_i, y_j \in \mathbb{R} \quad \text{and} \quad v_{ij} \leq 0 \quad (3.7)$$

$$\phi(y_i)v_{ii}(y_i - \phi(y_i)) = 0 \quad \forall y_i \in \mathbb{R} \quad \text{and} \quad v_{ii} \in \mathbb{R} \quad (3.8)$$

Summing these inequalities yields

$$\sum_{i=1}^m \sum_{\substack{j=1 \\ j \neq i}}^m \phi(y_i)v_{ij}(y_j - \phi(y_j)) + \sum_{i=1}^m \phi(y_i)v_{ii}(y_i - \phi(y_i)) \geq 0 \quad (3.9)$$

and majorising this expression leads to (3.6) with $\mathbf{V} = [v_{ij}]$. \square

Remark 3.1. The sector-like QC (3.6) takes the same form as the QC associated with the Sector[0, I] ie., where the sector-boundedness property in Table 3.1 holds for all m elements of $\Phi(\cdot)$. The QC associated with the Sector[0, I] holds for a more general class of static nonlinearities, when $\mathbf{V} \in \mathbb{D}_+^m$. However, it should be observed that for nonlinearities which possess the positivity and positive complement properties, such as the ReLU function, considerably more freedom in the choice of $\mathbf{V} \in \mathbb{Z}^m$ is introduced. This manifests itself by permitting unconstrained diagonal elements and non-positive off-diagonal elements. $\square\square$

Fact 3.2 (Positivity QC). Let $\Phi(\cdot) : \mathbb{R}^m \rightarrow \mathbb{R}_{\geq 0}^m$ be the repeated ReLU. If $\mathbf{Q}_{11}, \mathbf{Q}_{12}, \mathbf{Q}_{21}, \mathbf{Q}_{22} \in \mathbb{R}_{\geq 0}^{m \times m}$ then the following QC holds

$$\begin{bmatrix} \Phi(y) \\ \Phi(y) - y \end{bmatrix}' \begin{bmatrix} \mathbf{Q}_{11} & \mathbf{Q}_{12} \\ \mathbf{Q}_{21} & \mathbf{Q}_{22} \end{bmatrix} \begin{bmatrix} \Phi(\tilde{y}) \\ \Phi(\tilde{y}) - \tilde{y} \end{bmatrix} \geq 0 \quad \forall y, \tilde{y} \in \mathbb{R}^m \quad (3.10)$$

Proof: The positivity QC is derived by scaling the product of the positivity and positive complement properties with different arguments. Thus, for any $q_{11,ij}, q_{12,ij}, q_{21,ij}, q_{22,ij} \in \mathbb{R}_{\geq 0}$ it follows that

$$\phi(y_i)q_{11,ij}\phi(\tilde{y}_j) \geq 0 \quad \forall y_i, \tilde{y}_j \in \mathbb{R} \quad (3.11)$$

$$\phi(y_i)q_{12,ij}(\phi(\tilde{y}_j) - \tilde{y}_j) \geq 0 \quad \forall y_i, \tilde{y}_j \in \mathbb{R} \quad (3.12)$$

$$(\phi(y_i) - y_i)q_{21,ij}\phi(\tilde{y}_j) \geq 0 \quad \forall y_i, \tilde{y}_j \in \mathbb{R} \quad (3.13)$$

$$(\phi(y_i) - y_i)q_{22,ij}(\phi(\tilde{y}_j) - \tilde{y}_j) \geq 0 \quad \forall y_i, \tilde{y}_j \in \mathbb{R} \quad (3.14)$$

Summing over these inequalities and majorising the resulting expressions leads to (3.10) with $\mathbf{Q}_k = [q_{k,ij}]$ for $k = \{11, 12, 21, 22\}$. \square

Remark 3.2. *It is clear that the positivity QC (3.10) contains four individual QCs. In particular, (3.15) is the majorised expression of (3.11) included as the top left element of (3.10).*

$$\Phi(y)' \mathbf{Q}_{11} \Phi(\tilde{y}) \geq 0 \quad \forall y, \tilde{y} \in \mathbb{R}^m \quad (3.15)$$

Also, note that the positivity QC (3.10) holds for $\tilde{y} = y$. Two extra QCs (3.16), (3.17) can be extracted from this case; the other two possible QCs are redundant since they would take the same form as the sector-like QC, but with matrices from a more constrained set.

$$\Phi(y)' \mathbf{Q}_{11} \Phi(y) \geq 0 \quad \forall y \in \mathbb{R}^m \quad (3.16)$$

$$[\Phi(y) - y]' \mathbf{Q}_{22} [\Phi(y) - y] \geq 0 \quad \forall y \in \mathbb{R}^m \quad (3.17)$$

Later in the chapter, it will become clear that inequality (3.10) in its entirety is rather unwieldy and the special cases (3.15), (3.16) are easier to apply. $\square\square$

3.4 Global Stability Analysis

This section derives the strengthened stability criteria. Two Lyapunov candidates are proposed and the tailored QCs are used to derive matrix inequalities which, by the Barbashin-Krasovskii Theorem (Theorem 2.2), are sufficient to verify the origin of the specialised Lurie system (3.1) is globally asymptotically stable (GAS). The first Lyapunov candidate has a quadratic form, as in the Circle Criterion, and the second is of the Lurie-type, as in the Popov Criterion (Refer to Section 2.3 for more details on both). Due to the integral term in the Lurie-type Lyapunov candidate, the strengthened Popov Criterion may only be applied to systems with $D = 0$, as is the case with the standard Popov Criterion.

Theorem 3.1 (Circle-like Criterion). *Consider the Lurie system (3.1) with $\Phi(\cdot)$ the repeated ReLU. Let Assumption 3.1 be satisfied. If there exists $\mathbf{P} \in \mathcal{S}_+^n$, $\mathbf{V} \in \mathcal{Z}^m$, $\mathbf{Q}_{11} \in \mathbb{R}_{\geq 0}^{m \times m}$ such that*

$$\begin{bmatrix} He(A'\mathbf{P}) & \mathbf{PB} + C'\mathbf{V}' \\ \star & He(\mathbf{Q}_{11} - \mathbf{V}(I - D)) \end{bmatrix} \prec 0 \quad (3.18)$$

then the origin of (3.1) is GAS.

Proof: The quadratic Lyapunov candidate (3.19) is radially unbounded and satisfies $V_c(x) > 0 \quad \forall x \neq 0$ if $\mathbf{P} \in \mathcal{S}_+^n$.

$$V_c(x) = x'\mathbf{P}x \quad (3.19)$$

Since $\Phi(\cdot)$ is the repeated ReLU, QCs (3.6) and (3.16) are satisfied. Appending these to the time derivative of $V_c(\cdot)$ leads to (3.20).

$$\dot{V}_c(x) \leq \dot{x}'\mathbf{P}x + x'\mathbf{P}\dot{x} + 2\Phi(y)'\mathbf{V}[y - \Phi(y)] + 2\Phi(y)'\mathbf{Q}_{11}\Phi(y) \quad (3.20)$$

Substituting (3.1) into (3.20) and putting into quadratic form results in (3.21). Therefore, $\dot{V}_c(x) < 0 \quad \forall x \neq 0$ if (3.18) holds.

$$\dot{V}_c(x) \leq \begin{bmatrix} x \\ \Phi \end{bmatrix}' \begin{bmatrix} He(A'\mathbf{P}) & \mathbf{PB} + C'\mathbf{V}' \\ \star & He(\mathbf{Q}_{11} - \mathbf{V}(I - D)) \end{bmatrix} \begin{bmatrix} x \\ \Phi \end{bmatrix} \quad (3.21)$$

□

Remark 3.3. *The Circle-like Criterion is a specialisation of the Circle Criterion when $\Phi(\cdot)$ is the repeated ReLU. Since the solution space of the Circle Criterion is a subset of (3.18) when $\mathbf{V} \in \mathcal{Z}^m$, $\mathbf{Q}_{11} \in \mathbb{R}_{\geq 0}^{m \times m}$ are reduced to $\mathbf{V} \in \mathbb{D}_+^m$, $\mathbf{Q}_{11} = 0$, one expects Theorem 3.1 to be less conservative than the Circle Criterion.* □□

Theorem 3.2 (Popov-like Criterion). *Consider the Lurie system (3.1) with $\Phi(\cdot)$ the repeated ReLU and let $D = 0$. If there exists $\mathbf{P} \in \mathcal{S}_+^n$; $\mathbf{H} \in \mathbb{R}^{m \times m}$; $\mathbf{\Lambda}, \mathbf{W} \in \mathbb{D}_+^m$; $\mathbf{V} \in \mathcal{Z}^m$; $\mathbf{Q}_{11}, \tilde{\mathbf{Q}}_{11} \in \mathbb{R}_{\geq 0}^{m \times m}$ such that*

$$\begin{bmatrix} He(A'\mathbf{P}) & \mathbf{PB} + C'\mathbf{V}' + A'C'\mathbf{H}'\mathbf{\Lambda} & A'C'\mathbf{H}'\mathbf{\Lambda} + C'(\mathbf{H}' - I)\mathbf{W} \\ \star & He(\tilde{\mathbf{Q}}_{11} + \mathbf{Q}_{11} - \mathbf{V} + \mathbf{\Lambda}\mathbf{H}\mathbf{C}\mathbf{B}) & B'C'\mathbf{H}'\mathbf{\Lambda} + \tilde{\mathbf{Q}}_{11} \\ \star & \star & -2\mathbf{W} \end{bmatrix} \prec 0 \quad (3.22)$$

then the origin of (3.1) is GAS.

Proof: The following Lyapunov candidate is considered

$$V_p(x) = x'Px + 2 \int_0^{\mathbf{H}y} \Lambda \Phi(\sigma) \cdot d\sigma \quad (3.23)$$

$$= x'Px + \Phi(\mathbf{H}y)' \Lambda \Phi(\mathbf{H}y) \quad (3.24)$$

This is a more general version of the standard Lurie-type Lyapunov candidate since the line integral features an additional, *unstructured* matrix $\mathbf{H} \in \mathbb{R}^{m \times m}$, to be determined.

A necessary and sufficient condition for the line integral to be path independent is the existence of a scalar function of which the integrand is the gradient [102; 103]. Since $\Phi(\cdot)$ is the repeated ReLU, the scalar function $\theta(\cdot)$ can be chosen as

$$\theta(\sigma) = \Phi(\sigma)' \Lambda \Phi(\sigma) \quad (3.25)$$

The gradient of θ can be calculated as

$$\nabla_\sigma(\theta(\sigma)) = 2 \frac{\partial \Phi}{\partial \sigma} \Lambda \Phi(\sigma) \quad (3.26)$$

where

$$\frac{\partial \Phi}{\partial \sigma} = \text{diag}\left(\frac{\partial \phi(\sigma_1)}{\partial \sigma_1}, \dots, \frac{\partial \phi(\sigma_m)}{\partial \sigma_m}\right) \quad (3.27)$$

Since $\frac{\partial \Phi}{\partial \sigma}$ and Λ are both diagonal it follows that

$$\nabla_\sigma(\theta(\sigma)) = 2\Lambda \frac{\partial \Phi}{\partial \sigma} \Phi(\sigma) = 2\Lambda \Phi(\sigma) \quad \text{a.e.} \quad (3.28)$$

The final equality holds because $\frac{\partial \phi_i(\sigma_i)}{\partial \sigma_i} \phi_i(\sigma_i) = \phi_i(\sigma_i)$ almost everywhere (to see this set $\sigma_i < 0$ and then $\sigma_i \geq 0$). Therefore, by the Gradient Theorem, the Lyapunov candidate (3.23) can equally be expressed as (3.24). This is clearly radially unbounded and satisfies $V_p(x) > 0 \quad \forall x \neq 0$, if $\mathbf{P} \in \mathcal{S}_+^n$ and $\Lambda \in \mathcal{D}_+^m$. Furthermore, it is *independent of the choice of \mathbf{H}* .

The time derivative of $V_p(\cdot)$ is given by (3.29)-(3.31) where, for convenience, $\tilde{y} := \mathbf{H}y$.

$$\dot{V}_p(x) = \dot{x}'\mathbf{P}x + x'\mathbf{P}\dot{x} + \nabla_{\tilde{y}}(\theta(\tilde{y})) \cdot \frac{\partial \tilde{y}}{\partial x} \dot{x} \quad (3.29)$$

$$= \dot{x}'\mathbf{P}x + x'\mathbf{P}\dot{x} + (2\mathbf{\Lambda}\Phi(\tilde{y}))'\mathbf{H}C\dot{x} \quad (3.30)$$

$$= 2\left(x'\mathbf{P} + \Phi(\tilde{y})'\mathbf{\Lambda}\mathbf{H}C\right)(Ax + B\Phi(y)) \quad (3.31)$$

The final equality features the repeated ReLU with two different arguments. However, $\Phi(\tilde{y})$ can be expressed as

$$\Phi(\tilde{y}) = \underbrace{\Phi(\tilde{y}) - \Phi(y)}_{=:\Psi(\tilde{y},y)} + \Phi(y) \quad (3.32)$$

Both the sector-like QC (3.6) and the standard slope-restricted QC (2.82) can be appended to equation (3.31). It should be noted that for the repeated ReLU, the slope-restricted QC is valid globally with $K_1 = 0$ and $K_2 = I$. This results in

$$\begin{aligned} \dot{V}_p(x) \leq & 2\left(x'\mathbf{P} + \Phi(\tilde{y})'\mathbf{\Lambda}\mathbf{H}C\right)(Ax + B\Phi(y)) + 2\Phi(y)'\mathbf{V}[y - \Phi(y)] \\ & + 2\Psi(\tilde{y},y)'\mathbf{W}[\tilde{y} - y - \Psi(\tilde{y},y)] \end{aligned} \quad (3.33)$$

Using (3.32) and $\tilde{\mathbf{Q}}_{11} \in \mathbb{R}_{\geq 0}^{m \times m}$ the positivity QC (3.15) can be re-written as

$$\Phi(y)'\tilde{\mathbf{Q}}_{11}\left(\Psi(\tilde{y},y) + \Phi(y)\right) \geq 0 \quad \forall y, \tilde{y} \in \mathbb{R}^m \quad (3.34)$$

Appending this and QC (3.16) to (3.33) gives

$$\begin{aligned} \dot{V}_p(x) \leq & 2\left(x'\mathbf{P} + \Phi(\tilde{y})'\mathbf{\Lambda}\mathbf{H}C\right)(Ax + B\Phi(y)) \\ & + 2\Phi(y)'\mathbf{V}[y - \Phi(y)] + 2\Psi(\tilde{y},y)'\mathbf{W}[\tilde{y} - y - \Psi(\tilde{y},y)] \\ & + 2\Phi(y)'\tilde{\mathbf{Q}}_{11}[\Psi(\tilde{y},y) + \Phi(y)] + 2\Phi(y)'\mathbf{Q}_{11}\Phi(y) \end{aligned} \quad (3.35)$$

This can be majorised to get

$$\dot{V}_p(x) \leq \begin{bmatrix} x \\ \Phi \\ \Psi \end{bmatrix}' \begin{bmatrix} \text{He}(A'\mathbf{P}) & \mathbf{P}\mathbf{B} + \mathbf{C}'\mathbf{V}' + A'\mathbf{C}'\mathbf{H}'\mathbf{\Lambda} & A'\mathbf{C}'\mathbf{H}'\mathbf{\Lambda} + \mathbf{C}'(\mathbf{H}' - \mathbf{I})\mathbf{W} \\ * & \text{He}(\tilde{\mathbf{Q}}_{11} + \mathbf{Q}_{11} - \mathbf{V} + \mathbf{\Lambda}\mathbf{H}\mathbf{C}\mathbf{B}) & \mathbf{B}'\mathbf{C}'\mathbf{H}'\mathbf{\Lambda} + \tilde{\mathbf{Q}}_{11} \\ * & * & -2\mathbf{W} \end{bmatrix} \begin{bmatrix} x \\ \Phi \\ \Psi \end{bmatrix} \quad (3.36)$$

□

Remark 3.4. *The Popov-like Criterion is a specialisation of the Popov Criterion when $\Phi(\cdot)$ is the repeated ReLU. Since the solution space of the Popov Criterion is a subset of (3.22) when $\mathbf{V} \in \mathcal{Z}^m$, $\mathbf{H} \in \mathbb{R}^{m \times m}$, $\mathbf{W} \in \mathbb{D}_+^m$, $\mathbf{Q}_{11}, \tilde{\mathbf{Q}}_{11} \in \mathbb{R}_{\geq 0}^{m \times m}$ are reduced to $\mathbf{V} \in \mathbb{D}_+^m$, $\mathbf{H} = \mathbf{I}$, $\mathbf{W} = \mathbf{Q}_{11} = \tilde{\mathbf{Q}}_{11} = 0$, one expects Theorem 3.2 to be less conservative than the standard Popov Criterion.* □□

Remark 3.5. *Although the Popov-like Criterion is less conservative than the Popov Criterion, the arising matrix inequality (3.22) is not linear in the matrix variables.* □□

Remark 3.6. *To further tailor the stability analysis to the repeated ReLU, the positivity QC (3.10) could have replaced the special case positivity QC (3.15) in (3.35). However, this would have introduced an additional nonlinear term to (3.22) and increased the complexity of the matrix inequality.* □□

Remark 3.7. *The Circle-like Criterion can be recovered from the Popov-like Criterion. As the solution space of (3.18) is a subset of (3.22) when $\mathbf{H} \in \mathbb{R}^{m \times m}$, $\mathbf{\Lambda} \in \mathbb{D}_+^m$, $\tilde{\mathbf{Q}}_{11} \in \mathbb{R}_{\geq 0}^{m \times m}$ are reduced to $\mathbf{H} = \mathbf{I}$, $\mathbf{\Lambda} = \tilde{\mathbf{Q}}_{11} = 0$, one expects Theorem 3.2 to be less conservative than Theorem 3.1. The trade-off for this reduced conservatism is increased complexity.* □□

3.5 Convex Relaxations for Theorem 3.2

Inequality (3.22) is a bilinear matrix inequality (BMI) which is difficult to convexify and impractical to solve for high-dimensional systems. Consequently, two convex relaxations are described below, which enable the matrix inequality to be expressed as an LMI. The approaches suggested below make specific choices for certain variables; however, less conservative relaxations may exist.

3.5.1 Specific Choice of $\mathbf{\Lambda}$ and \mathbf{W}

Corollary 3.1 (Relaxed Popov-like Criterion 1). *Consider the Lurie system (3.1) with $\Phi(\cdot)$ being the repeated ReLU and $D = 0$. If there exists $\mathbf{P} \in \mathcal{S}_+^n$; $\mathbf{H} \in \mathbb{R}^{m \times m}$; $\mathbf{V} \in \mathcal{Z}^m$; $\mathbf{Q}_{11}, \tilde{\mathbf{Q}}_{11} \in \mathbb{R}_{\geq 0}^{m \times m}$ such that*

$$\begin{bmatrix} \text{He}(A'\mathbf{P}) & \mathbf{P}\mathbf{B} + \mathbf{C}'\mathbf{V}' + A'\mathbf{C}'\mathbf{H}' & A'\mathbf{C}'\mathbf{H}' + \eta\mathbf{C}'(\mathbf{H}' - I) \\ \star & \text{He}(\tilde{\mathbf{Q}}_{11} + \mathbf{Q}_{11} - \mathbf{V} + \mathbf{H}\mathbf{C}\mathbf{B}) & \mathbf{B}'\mathbf{C}'\mathbf{H}' + \tilde{\mathbf{Q}}_{11} \\ \star & \star & -2\eta I \end{bmatrix} \prec 0 \quad (3.37)$$

then the origin of (3.1) is GAS.

Proof: Matrix inequality (3.37) is a special case of (3.22) where $\mathbf{\Lambda} = I$ and $\mathbf{W} = \eta I$. \square

Since $\mathbf{\Lambda}$ always appears in a product with \mathbf{H} , the choice $\mathbf{\Lambda} = I$ may be made without loss of generality. However, some conservatism is introduced since \mathbf{H} appears in a product with \mathbf{W} in the upper right element, without $\mathbf{\Lambda}$. As the only other appearance of \mathbf{W} is in the lower right element, the choice $\mathbf{W} = \eta I$ is made, with the choice of η guided by the Schur complement conditions needed to satisfy (3.22) (see e.g. [38]).

The advantage of this corollary, over the one presented below, is that it makes use of \mathbf{H} being a *full matrix*. The disadvantage is that restrictive choices for $\mathbf{\Lambda}$ and \mathbf{W} have been made which reduce the solution space.

3.5.2 Specific Choice of \mathbf{H}

Corollary 3.2 (Relaxed Popov-like Criterion 2). *Consider the Lurie system (3.1) with $\Phi(\cdot)$ being the repeated ReLU and $D = 0$. If there exists $\mathbf{P} \in \mathcal{S}_+^n$, $\mathbf{\Lambda} \in \mathbb{D}_+^m$, $\mathbf{V} \in \mathcal{Z}^m$, $\mathbf{Q}_{11} \in \mathbb{R}_{\geq 0}^{m \times m}$ such that*

$$\begin{bmatrix} \text{He}(A'\mathbf{P}) & \mathbf{P}\mathbf{B} + \mathbf{C}'\mathbf{V}' + A'\mathbf{C}'\mathbf{\Lambda} \\ \star & \text{He}(\mathbf{Q}_{11} + \mathbf{\Lambda}\mathbf{C}\mathbf{B} - \mathbf{V}) \end{bmatrix} \prec 0 \quad (3.38)$$

then the origin of (3.1) is GAS.

Proof: By setting $\mathbf{H} = I$ the matrix in (3.22) collapses to a 2 by 2 block matrix. This is a result of $\Phi(\tilde{y}) = \Phi(y) \Rightarrow \Psi(\tilde{y}, y) = 0$, which reduces (3.35) to

$$\begin{aligned} \dot{V}_p(x) &\leq 2(x'\mathbf{P} + \Phi(y)'\mathbf{\Lambda}\mathbf{C})(Ax + B\Phi(y)) \\ &\quad + 2\Phi(y)'\mathbf{V}[y - \Phi(y)] + 2\Phi(y)'(\tilde{\mathbf{Q}}_{11} + \mathbf{Q}_{11})\Phi(y) \end{aligned} \quad (3.39)$$

Since $\tilde{\mathbf{Q}}_{11} + \mathbf{Q}_{11} \in \mathbb{R}_{\geq 0}^{m \times m}$, one may set $\tilde{\mathbf{Q}}_{11} = 0$ without loss of generality. Putting (3.39) into quadratic form shows (3.36) has been relaxed to

$$\dot{V}_p(x) \leq \begin{bmatrix} x \\ \Phi \end{bmatrix}' \begin{bmatrix} \text{He}(A'\mathbf{P}) & \mathbf{P}\mathbf{B} + \mathbf{C}'\mathbf{V}' + A'\mathbf{C}'\mathbf{\Lambda} \\ \star & \text{He}(\mathbf{Q}_{11} + \mathbf{\Lambda}\mathbf{C}\mathbf{B} - \mathbf{V}) \end{bmatrix} \begin{bmatrix} x \\ \Phi \end{bmatrix} \quad (3.40)$$

TABLE 3.2: Example state space models (A, B, C, D) .

Ex	n	m	Source
1	9	3	[72] Ex. 3
2	3	3	[104] Ex.3
3	3	4	[83] Ex. 4.9
4	8	4	[105] Ex. 22
5	6	4	[105] Ex. 17
6	6	4	[105] Ex. 19
7	8	4	[105] Ex. 23
8	5	5	[104] Ex. 2
9	40	40	[32]
10	60	60	[32]
11	80	80	[32]
12	100	100	[32]

□

If the restriction $\mathbf{H} \in \mathbb{D}_+^m$ is made, the product of \mathbf{H} and $\mathbf{\Lambda}$ will always be a member of \mathbb{D}_+^m . Therefore, the product of \mathbf{H} and $\mathbf{\Lambda}$ may equivalently be represented by the matrix $\mathbf{\Lambda} \in \mathbb{D}_+^m$; $\mathbf{H} = I$ may be chosen without loss of generality.

The advantage of this corollary, over the one above, is that it has lower complexity and does not require any parameters to be chosen. Clearly, the disadvantage is that the flexibility introduced by $\mathbf{H} \in \mathbb{R}^{m \times m}$ has been completely removed.

3.6 Numerical Examples

The maximum series gain (also known as the maximum sector/slope size) was used to compare the conservatism of the criteria developed in this chapter against the classical Circle and Popov Criteria [14], and the more recent developments using Zames-Falb multiplier's [106] and quadratic forms [72]. The Projective method [107] was implemented to find the maximum series gain, whilst the LMIs were posed using YALMIP [43] and solved using MOSEK [42].

3.6.1 Numerical Setup

The setup involved inserting a series gain (α) into the feedback loop by replacing $\Phi(y)$ with $\alpha\Phi(y)$ in (3.1) where $\alpha \in \mathbb{R}_{\geq 0}$. It is clear from (3.1) that this is equivalent to replacing B with αB and D with αD in the LMIs of each criterion. The maximum series gain is the largest α for which each criterion can certify the origin of (3.1) is GAS. The Nyquist gain provides an upper bound on this quantity.

TABLE 3.3: Comparison of the maximum series gain.

Ex	Circle	Theorem 3.1	Popov	Corollary 3.1	Corollary 3.2	Park	Zames-Falb	Nyquist Gain
1	20.8800	39.5200	434.2800	10,953.77	100,000+	448.6800	434.2800	100,000+
2	89.9000	89.9000	89.9000	89.9000	89.9000	89.9000	89.9000	89.9000
3	0.5236	0.6818	0.5236	0.6818	0.6818	0.5236	0.5522	0.6983
4	0.0010	0.0012	0.0010	0.0012	0.0015	0.0010	0.0010	0.0020
5	0.0813	0.0814	0.0824	0.0814	0.0830	0.0845	0.0845	0.0869
6	0.1946	0.3901	0.1947	0.4155	0.5048	0.2266	0.2773	0.8202
7	0.0966	0.1232	0.0968	0.1232	0.1462	0.1035	0.1035	0.2002
8	2.0221	2.0221	2.0221	2.0221	2.0221	2.0221	2.0221	2.0221
9	1.4695	2.0516	1.4695	2.0516	2.0516	1.4695	1.4695	2.0600
10	1.3820	2.0224	1.3820	2.0343	2.2240	1.3820	1.3820	2.7000
11	1.4531	2.1189	1.4605	2.0449	2.2092	1.4605	1.4605	2.2200
12	1.4613	2.2144	1.4613	2.1582	2.3380	1.4613	1.4613	2.7500

Table 3.2 lists 12 state space models (A, B, C, D) chosen to benchmark the conservatism of our criteria. The first 8 are from the literature, whereas the last 4 were GAS Hopfield network's, randomly generated according to the SVD Combo parametrisation in [32]. All chosen models have $D = 0$ to allow comparison of all criteria. The table also lists the associated values of (n, m) where $x \in \mathbb{R}^n$ and $\Phi(\cdot) : \mathbb{R}^m \rightarrow \mathbb{R}^m$. The examples used can be found by referring to the references or looking at the related code³.

For each example, Corollary 3.1 was applied with $\eta = 10,000$ and, for Example's 10 - 12, the matrices $\mathbf{H}, \mathbf{V}, \mathbf{Q}_{11}, \tilde{\mathbf{Q}}_{11}$ were restricted to being symmetric in order to reduce the computation time. The diagonal matrices of the Zames-Falb method varied between examples and were often set to the identity due to lack of guidance in selection. Naturally, different choices of η and the Zames-Falb parameters may lead to different results; furthermore, the additional symmetry restrictions to some parameters of Corollary 3.1 may lead to increased conservatism.

3.6.2 Discussion

Table 3.3 presents the maximum series gain and Table 3.4 presents the number of decision variables associated with each criterion. It should be noted that total decision variable count of Corollary 3.1 was presented, even though for some examples, some variables were restricted to being symmetric during implementation. The following observations were noted:

- Corollary 3.2 is of equal or less conservatism than all existing criteria in 7 out of the 8 low-dimensional examples and all the high-dimensional examples. Examples 9-12 emphasise that Corollary 3.2 strikes a very appealing balance of reduced

³<https://github.com/CR-Richardson/Max-Series-Gain>

TABLE 3.4: Comparison of the decision variable count.

Ex	Circle	Theorem 3.1	Popov	Corollary 3.1	Corollary 3.2	Park	Zames-Falb
1	48	63	51	81	66	87	252
2	9	24	12	42	27	30	45
3	10	38	14	70	42	40	48
4	40	68	44	100	72	90	208
5	25	53	29	85	57	67	129
6	25	53	29	85	57	67	129
7	40	68	44	100	72	90	208
8	20	65	25	115	70	70	100
9	860	4,020	900	7,220	4,060	3,360	4,300
10	1,890	9,030	1,950	16,230	9,090	7,440	9,450
11	3,320	16,040	3,400	28,840	16,120	13,120	16,600
12	5,150	25,050	5,250	45,050	25,150	20,400	25,750

conservatism and tractable complexity compared to existing criteria. Corollary 3.1 is never less conservative than Corollary 3.2 and always has higher complexity. Although, it should be noted on the high-dimensional examples, the conservatism of Corollary 3.1 is competitive with Corollary 3.2, even with the symmetric matrix restriction.

- Theorem 3.1 is of equal or less conservatism than all existing criteria in 75% of low-dimensional examples and all high-dimensional examples, although still inferior to Corollary 3.2. Theorem 3.1 also strikes an appealing balance between reduced conservatism and tractable complexity. An advantage over Corollary 3.2 is that it can be applied when $D \neq 0$.
- In the high-dimensional examples, the new criteria are significantly less conservative than all existing criteria. In fact, for each high-dimensional example, all existing criteria recover identical values for the maximum series gain.
- Corollary 3.1 is more conservative than Popov in Example 5. This highlights the effect of the convex relaxation to Theorem 3.2. This effect is not observed for Corollary 3.2.
- Example 8: a positive system for which the multivariable Aizerman Conjecture holds [104]. In this case, all methods achieved the Nyquist gain upper bound on α . The same applies for Example 2.
- In the first 8 examples, Theorem 3.1 and Corollary 3.2 were of similar complexity, slightly higher than Popov. Corollary 3.1 had higher complexity than Corollary 3.2 and Park. Zames-Falb had notably higher complexity than Corollary 3.1, excluding Examples 3 and 8. As larger systems were considered, it becomes clear that the complexity of the Circle and Popov Criteria scales the best. Theorem 3.1, Corollary 3.2, [72], and Zames-Falb [106] all follow similar scaling trends. Finally, the complexity of Corollary 3.1 scales the worst.




FIGURE 3.3: Time taken to solve LMIs for Hopfield network's in Examples 9-12.

Of course, the number of decision variables is only one of several factors which contribute to the time taken to solve an LMI (Section 2.1.3). To provide a more comprehensive comparison of this, the time taken for each criterion to be solved was presented in Figure 3.3. This is equivalent to setting $\alpha = 1$, which for Example's 9-12, GAS was verified by each criterion. The computation was performed on an Apple M1 Pro and the * denotes that Theorem 3.1 and Corollary 3.2 were applied with \mathbf{V} , \mathbf{Q}_{11} reduced to being symmetric matrices; furthermore, the $^+$ denotes Corollary 3.1 was applied with all matrices being reduced to symmetric. The figure highlights that Theorem 3.1, Corollary 3.1 (with symmetric restrictions), and Corollary 3.2 are the slowest criteria to be solved and have similar scaling trends to Zames-Falb [106]. The time taken to solve Theorem 3.1 and Corollary 3.2 became comparable to [72] once the symmetry restrictions were enforced. Finally, the Circle and Popov Criterion were by far the quickest to be solved.

Taking into account the conservatism and computational factors associated with each criterion, the following guidance is suggested for practitioners working with high-dimensional systems of the form (3.1). It is recommended to first try the fast, but conservative, Circle and Popov Criteria. If these are unable to verify stability, then try the restricted versions of Theorem 3.1 and Corollary 3.2, as these strike a good balance between speed and conservatism. If these also fail, then try the corresponding full criterion. For these systems, it does not make sense to apply [72] or Zames-Falb [106] as these are more conservative, but take a similar amount of time as the respective versions of Theorem 3.1 and Corollary 3.2.

3.7 Conclusion

This chapter proposed the strengthened Circle and Popov Criteria for the analysis of Lurie systems with repeated ReLU nonlinearities. The criteria were built upon new,

tailored quadratic constraints derived for the ReLU function. The new criteria have, potentially, much lower levels of conservatism than the standard Circle and Popov Criteria and, in some cases, much lower than Park and Zames-Falb. The appeal of the new criteria has been demonstrated with numerical examples, where Corollary 3.2 in particular, demonstrated low levels of conservatism in all examples. When focusing on high-dimensional systems, the new criteria were significantly less conservative and Theorem 3.1 was consistently competitive with Corollary 3.2. The new results have two deficiencies: (i) the results are limited to the ReLU nonlinearity; (ii) the number of variables in the LMI grow quadratically. Despite this, it is hoped that these results may help bring neural network based control into the domain of safety critical systems.

Chapter 4

Stability Analysis of Discrete Lurie Systems with ReLU Nonlinearities

This chapter addresses the stability analysis of a discrete-time (DT) Lurie system featuring a static repeated ReLU nonlinearity. Such systems often arise in the analysis of Hopfield networks and other neural feedback loops. Custom quadratic constraints, satisfied by the repeated ReLU, are employed to strengthen the standard DT Circle and DT Popov Criteria for this specific Lurie system. The criteria can be expressed as a set of linear matrix inequalities (LMIs) with less restrictive conditions on the matrix variables. It is further shown that if the Lurie system under consideration has a unique equilibrium point at the origin, then this equilibrium point is in fact globally stable or unstable, meaning that local stability analysis will provide no additional benefit. Numerical examples demonstrate that the strengthened criteria achieve a desirable balance between reduced conservatism and complexity when compared to existing criteria.

4.1 Introduction

Chapters 1 and 2 highlighted that various systems involving neural networks (NNs) can be modelled as a Lurie system (Section 2.2.1). Notable instances are a Hopfield network [92] and the interconnection of a linear time-invariant (LTI) system with a feed-forward NN (Section 2.2.3). Whilst Chapter 3 directly analysed the stability of these system in *continuous-time*, the more traditional focus in the literature, most NN control systems would be implemented digitally. To fill this gap, this chapter focused on *discrete-time* (DT) Lurie systems (Figure 4.1).




FIGURE 4.1: Discrete-time Lurie system with static nonlinearity.

To ensure the stability of a DT Lurie system, one can employ various absolute stability criteria, including the classical DT Circle and DT Popov¹ Criteria [69], Tsytkin Criterion [108], alternative Lyapunov-based criteria [109; 73; 74], and Zames-Falb multipliers [110; 81; 82]. Each of these criteria are posed as semi-definite programming (SDP) problems involving linear matrix inequalities (LMIs); the benefits and implementations of SDPs are discussed in Section 2.1.3. These criteria vary in their approach to balancing computational complexity and conservatism. In parallel to the continuous-time case, the DT Circle and DT Popov Criteria offer lower complexity, while Park and Zames-Falb multipliers have higher complexity, but tend to be less conservative (Section 2.3).

Recent research has utilised the absolute stability framework and associated SDP tools to address various challenges in NN analysis. This includes: estimating the basin of attraction [20; 95; 111], synthesising NN controllers [21; 96], and robustness analysis [97; 98; 99; 112]. The major difficulty in NN analysis stems from the large number of activation functions, m . This leads to increased computational complexity in absolute stability problems, compared to those traditionally studied. Some less conservative tools, like the Park Criterion [73], exhibit extremely poor scalability with m . Conversely, tools like the DT Circle Criterion become overly conservative, also limiting their utility in NN analysis.

Contribution: This chapter focuses on the *specialised DT Lurie system*, where the nonlinearity is the repeated ReLU, commonly used in deep learning. The first contribution confronts the challenge of balancing conservatism and computational complexity in absolute stability problems encountered in NN analysis. It achieves this by enhancing the low complexity DT Circle and DT Popov Criteria tailored for this specialised Lurie system (Theorem 4.1 and Theorem 4.2). The second contribution is the remarkable discovery that, under certain conditions, local stability at the origin of the specialised DT Lurie system is equivalent to global stability (Theorem 4.4). This implies that if global stability is not provable, then attempts to prove local stability will be similarly futile.

¹Several versions of the DT Popov Criterion exist. When the DT Popov Criterion is referenced in this chapter, we always mean the version presented in [69, Theorem 4.3].

4.2 Preliminaries

This section states some fundamental results and properties which are later used for global and local stability analysis (Section 2.1.2). The first two results are propositions related to the mean value theorem (MVT).

4.2.1 Mean Value Theorem

Proposition 4.1 (Mean Value Theorem). *If $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is continuous on $[a, b]$, there is a point $\epsilon \in (a, b)$ such that*

$$\int_a^b \phi(\sigma) d\sigma = (b - a)\phi(\epsilon) \quad (4.1)$$

Proof: [103, Page 100]. □

Proposition 4.2 (MVT for slope-restricted functions). *If $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is continuous on $[a, b]$ and slope-restricted on $[0, \mu]$, then the following inequality must hold*

$$\int_a^b \phi(\sigma) d\sigma \leq \mu(b - a)^2 + \phi(a)(b - a) \quad (4.2)$$

Proof: If $\epsilon \in (a, b)$ and $\phi(\cdot)$ is slope-restricted on the same domain, then (4.3) must hold since $0 < \epsilon - a < b - a$. Subbing (4.3) into (4.1) results in Proposition 4.2.

$$0 \leq \frac{\phi(\epsilon) - \phi(a)}{\epsilon - a} \leq \mu \implies \phi(a) \leq \phi(\epsilon) \leq \mu(b - a) + \phi(a) \quad (4.3)$$

□

4.2.2 Properties of the ReLU Function

The global stability analysis results are constructed on the observation that the ReLU function satisfies a number of properties, summarised in Table 3.1. Although the slope-restricted property holds for many static nonlinearities, the first four properties are much less typical. In fact, the complementarity property holds for few activation functions other than ReLU, which is defined below for convenience.

Definition 4.1 (Repeated ReLU). If $\phi(\cdot) : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ is the ReLU function, the repeated ReLU is $\Phi(\cdot) : \mathbb{R}^m \rightarrow \mathbb{R}_{\geq 0}^m$

$$\phi(y_i) := \begin{cases} y_i & y_i \geq 0 \\ 0 & y_i < 0 \end{cases} \quad \Phi(\cdot) := \begin{bmatrix} \phi(\cdot) \\ \vdots \\ \phi(\cdot) \end{bmatrix} \quad (4.4)$$

Fact 4.1. For $U \in \mathcal{U}$ and by definition of the ReLU function, the repeated ReLU can be expressed as

$$\Phi(y) = Uy \quad (4.5)$$

where $\mathcal{U} = \{ \text{diag}(u_1, \dots, u_m) \mid u_i \in 0, 1 \text{ and } i \in 1, \dots, m \}$.

Proof: $\phi(y_i) = u(y_i)y_i$, where $u(\cdot)$ is the unit step function. Expressing this in vector form is equivalent to Fact 4.1. \square

4.2.3 Positively Homogenous Functions

A key property of the ReLU function, and the repeated ReLU by extension, which enables one to prove the remarkable results in Section 4.4 is *positive homogeneity*. Several facts about positively homogenous functions (Table 3.1) are introduced below.

Fact 4.2. If $\theta(\cdot) : \mathbb{R}^m \rightarrow \mathbb{R}^m$ is bijective and positively homogenous, then $\theta^{-1}(\cdot) : \mathbb{R}^m \rightarrow \mathbb{R}^m$ is positively homogenous too.

Proof: Assume that $\theta^{-1}(\cdot)$ exists and is not positively homogenous, that is $\alpha\theta^{-1}(v) \neq \theta^{-1}(\alpha v)$. However, because $\theta(\cdot)$ is positively homogenous and bijective, it follows that for all $\alpha \in \mathbb{R}_{\geq 0}$

$$v = \theta^{-1}\left(\frac{1}{\alpha}\theta(\alpha v)\right) \quad (4.6)$$

Now, by the assumption that $\theta^{-1}(\cdot)$ is not positively homogenous, this means that

$$v \neq \frac{1}{\alpha}\theta^{-1} \circ \theta(\alpha v) = v \quad (4.7)$$

Clearly this is a contradiction and hence $\theta^{-1}(\cdot)$ must be positively homogenous. \square

Fact 4.3. Let $\theta(\cdot) : \mathbb{R}^m \rightarrow \mathbb{R}^m$ be defined by $\theta(v) := v - D\Phi(v)$ for some matrix $D \in \mathbb{R}^{m \times m}$. If $\Phi(\cdot) : \mathbb{R}^m \rightarrow \mathbb{R}^m$ is positively homogenous, then so is $\theta(\cdot)$.

Proof: $\alpha\theta(v) = \alpha v - D\alpha\Phi(v) = \alpha v - D\Phi(\alpha v) = \theta(\alpha v)$ for all $\alpha \in \mathbb{R}_{\geq 0}$. \square

4.2.4 Quadratic Constraints Satisfied by the Repeated ReLU

Using the properties from Table 3.1, two novel and less restrictive quadratic constraints (QCs) were constructed in Section 3.3 (Fact 3.1 and Fact 3.2) which are satisfied by the repeated ReLU. The same QCs are leveraged in Section 4.3 of this work, along with the standard slope-restricted QC (Fact 2.11), which holds globally for the repeated ReLU with $K_1 = 0$ and $K_2 = I$.

4.2.5 Problem Setup

Consider the DT Lurie system in Figure 4.1, where $P(z) \in \mathcal{RH}_\infty$ is a finite dimensional DT LTI system with state space realisation $(A, B, C, D)^2$ and the static nonlinearity $\Phi(\cdot)$ is the repeated ReLU. The Lurie system is modelled by (4.8) with $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, $C \in \mathbb{R}^{m \times n}$ and $D \in \mathbb{R}^{m \times m}$. As the repeated ReLU satisfies $\Phi(0) = 0$, the origin is an equilibrium point³ of (4.8).

$$\begin{aligned} x_{k+1} &= Ax_k + B\Phi(y_k) \\ y_k &= Cx_k + D\Phi(y_k) \end{aligned} \tag{4.8}$$

Assumption 4.1 (Well-posedness). *A unique solution x_k exists to (4.8) for all $x_0 \in \mathbb{R}^n$ and all $k \in \mathbb{N}$.*

Well-posedness is equivalent to the existence of a unique solution to the state space equations (4.8). Since $\Phi(\cdot)$ is globally Lipschitz (and differentiable almost everywhere), this is ensured if there exists a unique solution $y_k = \theta^{-1}(Cx_k)$ to $\theta(y_k) := y_k - D\Phi(y_k) = Cx_k$. A sufficient condition for this is given by Lemma 2.1.

In many absolute stability results (e.g., DT Circle Criterion) the LMI (2.66), from Lemma 2.1, is an intrinsic part of the stability conditions, so well-posedness is guaranteed. For instance, for the zero order hold (ZOH) discretisation of the Hopfield network [92], system (4.8) has the form $A = e^{-\tau I}$, $B = \int_0^\tau e^{-\tau I} W d\tau$, $C = I$, $D = 0$ and hence is, trivially, well-posed. A comprehensive examination of well-posedness is outside the

²Section 2.1.1 discusses the fundamentals of signals and systems.

³Section 2.1.2 discusses Lyapunov stability analysis and defines an equilibrium point.

scope of this chapter; it is enough to state that numerous DT systems incorporating NNs inherently exhibit well-posedness. Refer to Section 2.2.2 for further discussion.

Throughout the remainder of the chapter, Assumption 4.1 is assumed to hold. This implies a solution $y_k = \theta^{-1}(Cx_k)$ exists to (4.8); furthermore Fact 4.2 and Fact 4.3 show the solution is positively homogenous.

4.3 Global Stability Analysis

This section applies the DT counterpart of the Barbashin-Krasovskii Theorem (Section 2.1.2) to derive two LMIs which verify the origin of the specialised Lurie system (4.8) is globally asymptotically stable (GAS). The main results are constructed using quadratic and Lurie-type Lyapunov candidates, as is respectively the case in the DT Circle and DT Popov Criteria, and the novel QCs discussed in Section 4.2.4.

Theorem 4.1 (DT Circle-like Criterion). *Consider the DT Lurie system (4.8) with $\Phi(\cdot)$ the repeated ReLU. Let Assumption 4.1 be satisfied. If there exists $\mathbf{P} \in \mathbb{S}_+^n$, $\mathbf{V} \in \mathbb{Z}^m$, and $\mathbf{Q}_{11} \in \mathbb{R}_{\geq 0}^{m \times m}$ such that*

$$\begin{bmatrix} A'\mathbf{P}A - \mathbf{P} & A'\mathbf{P}B + C'\mathbf{V}' \\ \star & B'\mathbf{P}B + He(\mathbf{Q}_{11} - \mathbf{V}(I - D)) \end{bmatrix} \prec 0 \quad (4.9)$$

then the origin of (4.8) is GAS.

Proof: Choosing a quadratic Lyapunov candidate $V_{cl}(x) = x'\mathbf{P}x$ with $\mathbf{P} \in \mathbb{S}_+^n$ and looking at the difference along the trajectories of system (4.8) gives

$$\begin{aligned} \Delta V_{cl} &= (Ax_k + B\Phi_k)'\mathbf{P}(Ax_k + B\Phi_k) - x_k'\mathbf{P}x_k \\ &= x_k'(A'\mathbf{P}A - \mathbf{P})x_k + x_k'A'\mathbf{P}B\Phi_k + \Phi_k'B'\mathbf{P}Ax_k + \Phi_k'B'\mathbf{P}B\Phi_k \end{aligned} \quad (4.10)$$

where $\Delta V_{cl} := V_{cl}(x_{k+1}) - V_{cl}(x_k)$ and $\Phi_k := \Phi(y_k)$. Appending the sector-like QC (3.6) and the positivity QC (3.15) for the special case $\tilde{y}_k = y_k$, leads to

$$\begin{aligned} \Delta V_{cl} &\leq x_k'(A'\mathbf{P}A - \mathbf{P})x_k + x_k'A'\mathbf{P}B\Phi_k + \Phi_k'B'\mathbf{P}Ax_k + \Phi_k'B'\mathbf{P}B\Phi_k \\ &\quad + 2\Phi_k'\mathbf{V}(y_k - \Phi_k) + 2\Phi_k'\mathbf{Q}_{11}\Phi_k \end{aligned} \quad (4.11)$$

The right hand side of (4.11) will be negative definite if Theorem 4.1 is satisfied. This can be seen by substituting (4.8) in for y_k , then expressing (4.11) in quadratic form. \square

Remark 4.1. Theorem 4.1 strengthens the DT Circle Criterion when $\Phi(\cdot)$ is the repeated ReLU. One expects Theorem 4.1 to verify GAS for a larger space of (A, B, C, D) matrices since the DT Circle Criterion has more restrictions on the LMI variables: $\mathbf{V} \in \mathbb{D}_+^m$, $\mathbf{Q}_{11} = 0$. $\square\square$

Theorem 4.2 (DT Popov-like Criterion). Consider the Lurie system (4.8) with $\Phi(\cdot)$ the repeated ReLU and let $D = 0$. If there exists $\mathbf{P} \in \mathbb{S}_+^n$; $\mathbf{H} \in \mathbb{R}^{m \times m}$; $\mathbf{\Lambda}, \mathbf{W} \in \mathbb{D}_+^m$; $\mathbf{V} \in \mathbb{Z}^m$; $\tilde{\mathbf{Q}}_{11}, \mathbf{Q}_{11} \in \mathbb{R}_{\geq 0}^{m \times m}$ such that

$$\begin{bmatrix} X_{11} & X_{12} & (A - I)'C'H'\mathbf{\Lambda} + C'(\mathbf{H} - I)'\mathbf{W} \\ \star & X_{22} & B'C'H'\mathbf{\Lambda} + \tilde{\mathbf{Q}}_{11} \\ \star & \star & -2\mathbf{W} \end{bmatrix} \prec 0 \quad (4.12)$$

$$\begin{aligned} X_{11} &= A'\mathbf{P}A - \mathbf{P} + He((A - I)'C'H'\mathbf{\Lambda}HC(A - I)) \\ X_{12} &= A'\mathbf{P}B + C'\mathbf{V}' + (A - I)'C'H'\mathbf{\Lambda} + 2(A - I)'C'H'\mathbf{\Lambda}HC B \\ X_{22} &= B'\mathbf{P}B + He(\mathbf{Q}_{11} - \mathbf{V} + \mathbf{\Lambda}HC B + B'C'H'\mathbf{\Lambda}HC B + \tilde{\mathbf{Q}}_{11}) \end{aligned}$$

then the origin of (4.8) is GAS.

Proof: A generalised Lurie-type Lyapunov candidate

$$V_{pl}(x) = x'\mathbf{P}x + 2 \int_0^{\mathbf{H}y_k} \mathbf{\Lambda}\Phi(\sigma) \cdot d\sigma \quad (4.13)$$

with $\mathbf{P} \in \mathbb{S}_+^n$, $\mathbf{H} \in \mathbb{R}^{m \times m}$ and $\mathbf{\Lambda} \in \mathbb{D}_+^m$ is selected. Due to the integral term in the generalised Lurie-type Lyapunov candidate, Theorem 4.2 may only be applied to systems with $D = 0$, as is the case with the DT Popov Criterion. Now, looking at the difference along the trajectories of system (4.8) results in

$$\begin{aligned} \Delta V_{pl} &= 2 \int_0^{\mathbf{H}y_{k+1}} \mathbf{\Lambda}\Phi(\sigma) \cdot d\sigma - 2 \int_0^{\mathbf{H}y_k} \mathbf{\Lambda}\Phi(\sigma) \cdot d\sigma \\ &\quad + \underbrace{(Ax_k + B\Phi_k)'\mathbf{P}(Ax_k + B\Phi_k) - x_k'\mathbf{P}x_k}_{=\Delta V_{cl}} \end{aligned} \quad (4.14)$$

where $\Delta V_{pl} := V_{pl}(x_{k+1}) - V_{pl}(x_k)$ and $\Phi_k := \Phi(y_k)$. Since the line integrals are path independent (see Section 3.4 for more discussion on this), assume the system follows a trajectory which begins at the origin, passes through $\mathbf{H}y_k$ and ends at $\mathbf{H}y_{k+1}$, without loss of generality. As a result

$$\Delta V_{pl} = \Delta V_{cl} + 2 \int_{\tilde{y}_k}^{\tilde{y}_{k+1}} \mathbf{\Lambda}\Phi(\sigma) \cdot d\sigma \quad (4.15)$$

where $\tilde{y}_k := \mathbf{H}y_k$. As $\mathbf{\Lambda} = \text{diag}(\lambda_1 \dots \lambda_m)$, the line integral can be expressed as

$$\int_{\tilde{y}_k}^{\tilde{y}_{k+1}} \mathbf{\Lambda} \Phi(\sigma) \cdot d\sigma = \sum_{i=1}^m \lambda_i \int_{\tilde{y}_{k,i}}^{\tilde{y}_{k+1,i}} \phi(\sigma_i) d\sigma_i \quad (4.16)$$

The ReLU function, $\phi(\cdot)$, is slope-restricted on $[0, 1]$; hence, Proposition 4.2 can be applied with $a = \tilde{y}_{k,i}$, $b = \tilde{y}_{k+1,i}$ and $\mu = 1$. Expressing the resulting summation in quadratic form and substituting into (4.15) leads to

$$\Delta V_{pl} \leq \Delta V_{cl} + 2\tilde{\Phi}'_k \mathbf{\Lambda} (\tilde{y}_{k+1} - \tilde{y}_k) + 2(\tilde{y}_{k+1} - \tilde{y}_k)' \mathbf{\Lambda} (\tilde{y}_{k+1} - \tilde{y}_k) \quad (4.17)$$

where $\tilde{\Phi}_k := \Phi_k(\tilde{y}_k)$. Appending the sector-like QC (3.6), both cases of the positivity QC (3.15) and the slope-restricted QC (2.82) gives

$$\begin{aligned} \Delta V_{pl} \leq & 2\tilde{\Phi}'_k \mathbf{\Lambda} (\tilde{y}_{k+1} - \tilde{y}_k) + 2(\tilde{y}_{k+1} - \tilde{y}_k)' \mathbf{\Lambda} (\tilde{y}_{k+1} - \tilde{y}_k) \\ & + \underbrace{\Delta V_{cl} + 2\tilde{\Phi}'_k \mathbf{V} (y_k - \Phi_k) + 2\tilde{\Phi}'_k \mathbf{Q}_{11} \Phi_k}_{(4.9) \text{ when } D=0} \\ & + 2\tilde{\Phi}'_k \tilde{\mathbf{Q}}_{11} \tilde{\Phi}_k + 2\Psi'_k \mathbf{W} (\tilde{y}_k - y_k - \Psi_k) \end{aligned} \quad (4.18)$$

where $\Psi_k := \Psi(\tilde{y}_k, y_k) = \tilde{\Phi}_k - \Phi_k$. Now, substituting $\tilde{\Phi}_k = \Psi_k + \Phi_k$ and replacing $\tilde{y}_{k+1}, \tilde{y}_k, y_k$ with the respective functions of x_k gives

$$\begin{aligned} \Delta V_{pl} \leq & 2(\Psi_k + \Phi_k)' \mathbf{\Lambda} \mathbf{H} \mathbf{C} \left((A - I)x_k + B\Phi_k \right) + 2\Psi'_k \mathbf{W} \left((\mathbf{H} - I)Cx_k - \Psi_k \right) \\ & + 2\left((A - I)x_k + B\Phi_k \right)' \mathbf{C}' \mathbf{H}' \mathbf{\Lambda} \mathbf{H} \mathbf{C} \left((A - I)x_k + B\Phi_k \right) \\ & + 2\tilde{\Phi}'_k \tilde{\mathbf{Q}}_{11} (\Psi_k + \Phi_k) + \underbrace{\Delta V_{cl} + 2\tilde{\Phi}'_k \mathbf{V} (y_k - \Phi_k) + 2\tilde{\Phi}'_k \mathbf{Q}_{11} \Phi_k}_{(4.9) \text{ when } D=0} \end{aligned} \quad (4.19)$$

The right hand side of (4.19) is guaranteed to be negative definite if Theorem 4.2 is satisfied. This can be seen by substituting the associated quadratic form of LMI (4.9) into (4.19) when $D = 0$, then manipulating the full expression into quadratic form. \square

Remark 4.2. Theorem 4.2 contains matrix variable products which prevents the matrix inequality from being linear. Setting $\mathbf{H} = I$ is one convex relaxation which reduces (4.12) to an

LMI. Furthermore, this relaxation forces $\Psi(\tilde{y}_k, y_k) \equiv 0$, hence the block matrix from (4.12) is still required to be negative definite, but has collapsed to the (2×2) block matrix

$$\begin{bmatrix} A'PA - P + He((A - I)'C'\Lambda C(A - I)) & A'PB + C'V' + (A - I)'C'\Lambda + 2(A - I)'C'\Lambda CB \\ * & B'PB + He(Q_{11} - V + \Lambda CB + B'C'\Lambda CB) \end{bmatrix} \prec 0 \quad (4.20)$$

□□

Remark 4.3. Remark 4.2 strengthens the DT Popov Criterion when $\Phi(\cdot)$ is the repeated ReLU. One expects Remark 4.2 to verify GAS for a larger space of (A, B, C) matrices since the DT Popov Criterion has more restrictions on the LMI variables: $V \in \mathbb{D}_+^m$, $Q_{11} = \tilde{Q}_{11} = 0$. □□

4.4 Equivalence of Local and Global Stability

As many properties of the ReLU function hold globally, limiting the scope of the stability analysis to a local region did not seem promising, as done in [20]. For this reason, conditions which ensured an asymptotically stable equilibrium point was, in fact, GAS were investigated.

Theorem 4.3. System (4.8) has a unique equilibrium point at the origin if all matrices in the set \mathcal{H} are full rank where

$$\mathcal{H} = \left\{ A - I + BU(I - DU)^{-1}C \mid U \in \mathcal{U} \right\} \quad (4.21)$$

and the 2^m possible permutations of U are captured by the set

$$\mathcal{U} = \left\{ \text{diag}(u_1, \dots, u_m) \mid u_i \in \{0, 1\} \text{ and } i \in \{1, 2, \dots, m\} \right\} \quad (4.22)$$

Proof: Any equilibrium point, x_{eq} , of (4.8) must satisfy $x_{k+1} - x_k = 0$. This is equivalent to

$$0 = (A - I)x_{k,eq} + B\Phi(y_{k,eq}) \quad y_{k,eq} = Cx_{k,eq} + D\Phi(y_{k,eq}) \quad (4.23)$$

Using Fact 4.1, the output equation of (4.23) may equivalently be expressed as

$$(I - DU)y_{k,eq} = Cx_{k,eq} \Leftrightarrow y_{k,eq} = \underbrace{(I - DU)^{-1}C}_{:=K} x_{k,eq} \quad (4.24)$$

Equation (4.24) does not provide the general solution to the output equation, but instead a relationship between an equilibrium state and the corresponding system output. Using Fact 4.1 and substituting (4.24) into the state equation of (4.23) results in a single equation which any equilibrium point must satisfy

$$\begin{aligned} 0 &= (A - I)x_{k,eq} + BUy_{k,eq} \\ &= (A - I + BUK)x_{k,eq} =: \check{A}x_{k,eq} \end{aligned} \quad (4.25)$$

Since $U \in \mathcal{U}$, then $\check{A} \in \mathcal{H}$. Therefore, if all matrices in \mathcal{H} have full rank, the only solution to (4.25) is $x_{k,eq} = 0$. \square

This condition is sufficient, but may be conservative as the state equations may only allow U to enter a subset of the 2^m possible permutations.

Remark 4.4. *As a square matrix is only full rank if and only if it is also invertible, Theorem 4.3 can be verified by computing the determinant for each matrix in the set \mathcal{H} . If all are non-zero, the set \mathcal{H} is full rank.* $\square\square$

Theorem 4.4. *If the origin of (4.8) is a unique equilibrium point and a ball of any radius r_x can be established as a region of attraction $\mathcal{B}_{r_x} = \{x : \|x\| \leq r_x\}$ then, the origin is actually a GAS equilibrium point.*

Proof: Under Assumption 4.1, system (4.8) can be expressed as

$$x_{k+1} = Ax_k + B\Phi \circ \theta^{-1}(Cx_k) \quad (4.26)$$

where $\theta(y_k) := y_k - D\Phi(y_k) = Cx_k$. Now assume Theorem 4.3 is satisfied, which guarantees the origin is a unique equilibrium point of (4.26). Furthermore, assume it can be established the origin of (4.26) has a region of attraction given by $\mathcal{B}_{r_x} := \{x : \|x\| \leq r_x\}$. Now, define a positively scaled initial state $z_0 := \alpha x_0$ where $\alpha > 0$. The initial state z_0 will follow the trajectory

$$z_k(z_0) = \alpha x_k(x_0) \quad (4.27)$$

Using Fact 4.2 and Fact 4.3, it is clear that this is an instance of the general solution to the scaled system

$$\begin{aligned} z_{k+1} &= \alpha Ax_k + \alpha B\Phi \circ \theta^{-1}(Cx_k) \\ &= A(\alpha x_k) + B\Phi \circ \theta^{-1}(C(\alpha x_k)) \\ &= Az_k + B\Phi \circ \theta^{-1}(Cz_k) \end{aligned} \quad (4.28)$$

As the scaled system (4.28) is of the same form as (4.26), a region of attraction $\mathcal{B}_{r_z} = \{z : \|z\| \leq r_z\}$ must exist around the origin of (4.28). Using (4.27), a relationship between \mathcal{B}_{r_x} and \mathcal{B}_{r_z} can be established. First, \mathcal{B}_{r_x} may equivalently be expressed in terms of z

$$\mathcal{B}_{r_x} = \left\{x : \|x\| \leq r_x\right\} = \left\{\frac{1}{\alpha}z : \frac{1}{\alpha}\|z\| \leq r_x\right\} \quad (4.29)$$

A new set may then be defined which contains the unscaled vectors, z , satisfying (4.29)

$$\mathcal{B}_{r_z} = \left\{z : \|z\| \leq \alpha r_x\right\} \quad (4.30)$$

This is the region of attraction, where $r_z = \alpha r_x$. As α can represent any positive scalar, one may set $\mathcal{B}_{r_z} = \mathbb{R}^n$. Thus, if one can show the origin of (4.26) has a region of attraction, by the equivalence of (4.26) and (4.28), the origin of (4.26) is GAS. \square

Remark 4.5. *The implications of Theorem 4.4 are profound. Provided the origin is a unique equilibrium of system (4.26), then if global stability cannot be established, it is futile to attempt to establish local stability. This is somewhat counterintuitive. Typically in stability analysis, if one cannot establish global stability, one attempts a local stability analysis. Theorem 4.4 implies that, for ReLU problems, this will not be fruitful.* $\square\square$

4.5 Numerical Examples

The maximum series gain was used to compare the conservatism of the criteria developed in this chapter against criteria with low (DT Circle and DT Popov [69]) and high ([73]) complexity. The Projective method [107] was utilised to find the maximum series gain in conjunction with YALMIP [43] and MOSEK [42] for solving the LMIs.

4.5.1 Numerical Setup

The setup involved inserting a series gain ($\alpha \in \mathbb{R}_{\geq 0}$) into the feedback loop by replacing $\Phi(y_k)$ with $\alpha\Phi(y_k)$ in (4.8). It is clear from (4.8) that this is equivalent to replacing B with αB and D with αD in the LMIs of each criterion. The maximum series gain represents the highest value of α for which each criterion can verify the origin of (4.8) is GAS. The Nyquist gain provides an upper bound on this quantity.

The benchmark examples were discretised counterparts of the Lurie systems (A, B, C, D) used in Chapter 3, where the ZOH method was used with a sampling period of $T_s \in \{10^{-2}, 10^{-4}\}$ seconds. Table 3.2 lists the examples (prior to discretisation), taken from the literature, and the respective values of (n, m) where $x \in \mathbb{R}^n$ and $\Phi(\cdot) : \mathbb{R}^m \rightarrow \mathbb{R}^m$. Each state space model had a zero D matrix to facilitate comparison with the DT Popov

TABLE 4.1: Comparison of the maximum series gain.

Ex	DT Circle	Theorem 4.1	DT Popov	Remark 4.2	Park	Nyquist Gain
1	20.8659	39.4280	411.582	3310.38	450.117	6666.67
2	89.9000	89.9000	89.9000	89.9000	89.9000	89.9000
3	0.5236	0.6818	0.5236	0.6818	0.5236	0.6983
4	0.0010	0.0012	0.0010	0.0014	0.0010	0.0020
5	0.0813	0.0814	0.0824	0.0829	0.0845	0.0869
6	0.1951	0.3916	0.1951	0.5007	0.2272	0.8212
7	0.0967	0.1230	0.0969	0.1448	0.1037	0.2008
8	2.0221	2.0221	2.0221	2.0221	2.0221	2.0221
9	1.4681	2.0519	1.4789	2.0536	1.4766	2.0600
10	1.3781	1.5559	1.3892	1.6978	N/A	2.7000
11	1.4500	2.1161	1.4500	2.2113	N/A	2.2200
12	1.4584	2.1977	1.4693	2.3107	N/A	2.7500

and DT Popov-like Criteria. Consult the related code⁴ to see the specific examples. Finally, as the DT Popov-like Criterion was a bilinear matrix inequality (BMI), the convex relaxation stated in Remark 4.2 was deployed, to avoid losing the benefits of SDP problems.

4.5.2 Discussion

Table 4.1 presents the maximum series gain computed by each criterion and Table 4.2 states the number of decision variables associated with each criterion. Each system with α set less than the Nyquist gain has a unique equilibrium point; hence, by Theorem 4.4, it will be futile to attempt a local stability analysis on such systems.

In 7 out of the 8 low-dimensional examples, Remark 4.2 exhibits conservatism that is either equal to or less than all other criteria. In 6 out of the same 8 examples, Theorem 4.1 demonstrates conservatism that is either equal to or less than all existing criteria, albeit remaining inferior to Remark 4.2 in 5 out of 8 instances. However, one obvious advantage of Theorem 4.1 is its applicability when $D \neq 0$. Example 8 demonstrates a positive system wherein the multivariable Aizerman Conjecture is valid for the continuous time counterpart [104]. In this scenario, all approaches attained the linear upper limit on α . This was also achieved in Example 2.

In the four Hopfield network's, Examples 9-12, the new criteria were significantly less conservative than all other approaches. In fact for Examples 10 - 12, the Park Criterion [73] was impractical as it had such a large number of variables. For this reason, maximum series gain results were not obtained for these examples. This reduced conservatism came at the expense of an increased number of decision variables, compared

⁴<https://github.com/CR-Richardson/DT-Max-Series-Gain>

TABLE 4.2: Comparison of the decision variable count.

Ex	DT Circle	Theorem 4.1	DT Popov	Remark 4.2	Park
1	48	63	51	66	579
2	9	24	12	27	159
3	10	38	14	42	207
4	40	68	44	72	572
5	25	53	29	57	402
6	25	53	29	57	402
7	40	68	44	72	572
8	20	65	25	70	395
9	860	4,020	900	4,060	21,360
10	1,890	9,030	1,950	9,090	47,640
11	3,320	16,040	3,400	16,120	84,320
12	5,150	25,050	5,250	25,150	131,400

to the DT Circle and Popov Criteria, but far less than the Park Criterion. Hence, the new criteria balance conservatism and complexity well.

To further compare the complexity of the criteria, the time for each criterion to verify GAS of Example's 9-12 was plotted in Figure 4.2. It was known a priori that each tractable criterion could successfully solve these problems since they each found a maximum series gain greater than 1. The plot supports the earlier claim that the Circle and Popov Criteria scale the best, whilst Theorem 4.1 and Remark 4.2 were slower, but followed a similar scaling trend. Restricting the \mathbf{V} and \mathbf{Q}_{11} matrices to being symmetric, sacrifices some conservatism, but makes some significant computational savings. Finally, the Park Criterion was only tractable for Example 9 and was far slower than all other criteria.

4.6 Conclusion

This work extends the results from Chapter 3 to discrete-time Lurie systems, accounting for the digital implementation of a neural network controller. The DT Circle and DT Popov Criteria are strengthened for analysing DT Lurie systems with the repeated ReLU nonlinearity. These refined criteria offer the potential for significantly lower conservatism levels compared to the standard DT Circle and DT Popov Criteria, whilst retaining a practical level of computational efficiency for larger systems. Local stability was also investigated, but this resulted in conditions which show that if the Lurie system under consideration has a unique equilibrium point at the origin, then this equilibrium point is in fact globally stable or unstable, meaning that local stability analysis will provide no additional benefit. The main deficiency of the new results is their limitation to the ReLU nonlinearity. Despite this, it is hoped that these results may help bring neural network based control into the domain of safety critical systems.




FIGURE 4.2: Time taken to solve LMIs for Hopfield network's in Examples 9-12.

Chapter 5

Stability Analysis of Lurie Systems with Magnitude Nonlinearities

This chapter considers the interconnection of a continuous time linear time-invariant system and a multivariable magnitude nonlinearity. A number of different quadratic constraints are established for the magnitude nonlinearity and then used to derive stability criteria based on quadratic and Lurie-type Lyapunov functions. The new stability criteria are cast as matrix inequalities and in some cases solved using semi-definite programming. Connections are made between the magnitude nonlinearity and neural network activation functions, such as the (leaky) ReLU, effectively allowing the stability criteria derived here to be used to analyse interconnections of dynamical systems and neural networks. Using the positive homogeneity property, shared by the (leaky) ReLU and magnitude functions, mild conditions are also established to show that the existence of a unique equilibrium point is sufficient for local and global stability to be equivalent. Finally, numerical examples illustrate how the new global stability criteria compare favourably with competing results from the literature. For high-dimensional Hopfield networks, the conservatism is competitive with the strengthened Popov Criterion, whilst being applicable to a wider range of nonlinearities.

5.1 Introduction

Lurie systems are feedback interconnections involving a linear time-invariant (LTI) system and a, normally static, nonlinear element (Figure 3.1). They have been studied extensively in the control systems literature due to their relevance to a number of practical problems and because there are several elegant and computationally efficient ways of analysing their stability, as discussed in Section 2.3. Approaches include the Circle and Popov Criteria [93; 14; 71], as well as the techniques of Zames and Falb [58], Park [72] and others [61; 83]. Each criterion is formulated as a semi-definite programming (SDP)

problem involving linear matrix inequalities (LMIs). Advantages of SDPs and implementation details are covered in Section 2.1.3. The key distinction between the criteria lies in which class of nonlinearities they can be applied to, alongside how they manage the trade-off between conservatism and complexity.

Over recent years, a resurgence of interest in Lurie systems has occurred, due to the recognition that dynamical systems involving neural networks (NNs) can often be modelled in the Lurie framework (Section 2.2.1). Indeed recent analyses of such systems have either adopted established approaches [97; 55; 20; 21; 113] or have derived new but related theoretical results [112; 114; 115; 116]. An approach for handling activation functions which invalidate the slope-restricted assumption was detailed in [117]. Furthermore, in [118; 119], piecewise affine systems have been written in Lurie form, significantly expanding the class of systems covered.

In this chapter, the normal assumptions above are replaced with the assumption that the nonlinearity is the repeated (applied element-wise) magnitude function. This makes the results relevant to systems where the magnitude naturally arises such as electronic circuits containing full-bridge rectifiers, and a class of systems able to produce chaotic behaviour [120]. Furthermore, it is noted that some activation functions, such as the (leaky) ReLU, can also be represented using the magnitude nonlinearity. Therefore, the results introduced here can also be used to analyse the stability of systems containing such activation functions and can be considered a more general alternative to the results of Chapter 3.

Contribution: Novel loop transformations are presented in Section 5.2.2 which enable systems with magnitude nonlinearities to be represented as systems with (leaky) ReLU nonlinearities and vice versa. New quadratic constraints (QCs) for the repeated magnitude are presented in Section 5.3 and leveraged to establish new stability criteria (Section 5.4). Similar, but different, relaxations to those used in Chapter 3 are presented in Section 5.5 to convexify one of the main results. Section 5.6 presents some simple conditions for when global and local stability are equivalent; the proof is different but analogous to the discrete-time result from Chapter 4. Finally, Section 5.7 compares numerically the conservatism and complexity of the new criteria against existing results.

5.2 Preliminaries

5.2.1 Problem Setup

To describe the problem, the repeated magnitude function must first be defined.

Definition 5.1 (Repeated Magnitude). *If $|\cdot| : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ is the magnitude function, the repeated magnitude is $|\cdot| : \mathbb{R}^m \rightarrow \mathbb{R}_{\geq 0}^m$*

$$|u_i| := \begin{cases} +u_i & u_i \geq 0 \\ -u_i & u_i < 0 \end{cases} \quad |u| := \begin{bmatrix} |u_1| \\ \vdots \\ |u_m| \end{bmatrix} \quad (5.1)$$

Consider the Lurie system in Figure 3.1, where $P(s) \in \mathcal{RH}_\infty$, is a finite dimensional LTI system, with state-space realisation (A, B, C, D) ¹. The system is modelled by (5.2)

$$\begin{aligned} \dot{x} &= Ax + B\Phi(y) \\ y &= Cx + D\Phi(y) \end{aligned} \quad (5.2)$$

where $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, $C \in \mathbb{R}^{m \times n}$, $D \in \mathbb{R}^{m \times m}$ and $\Phi(\cdot) : \mathbb{R}^m \mapsto \mathbb{R}^m$. In this chapter, a special case of Lurie system is studied, where $\Phi(y) \equiv |y|$. That is

$$\begin{aligned} \dot{x} &= Ax + B|y| \\ y &= Cx + D|y| \end{aligned} \quad (5.3)$$

As the repeated magnitude satisfies $|0| = 0$, the origin is an equilibrium point² of (5.3).

The following assumption is made throughout the chapter.

Assumption 5.1 (Well-posedness). *A unique solution $x(t)$ exists to (5.3) for all $x(0) \in \mathbb{R}^n$ and all $t \geq 0$.*

As the magnitude function is globally Lipschitz, Assumption 5.1 will hold unconditionally for $D = 0$; when $D \neq 0$, some sufficient conditions are discussed in Section 5.2.3. The main problem addressed in the remainder of the paper is stated next.

Problem 5.1. *Find tractable Lyapunov conditions which ensure the origin of (5.3) is globally asymptotically stable (GAS).*

5.2.2 Loop Transformations Between Magnitude and ReLU

Prior work, detailed in Section 2.2.3, has shown that many NN stability problems can be expressed as a Lurie system (5.2), where $\Phi(\cdot)$ contains the NN activation functions. When the popular leaky ReLU is chosen as the NN activation function, such problems

¹Section 2.1.1 discusses the fundamentals of signals and systems.

²Section 2.1.2 discusses Lyapunov stability analysis and defines an equilibrium point.

can be framed in the form (5.3), as shown below. Consider the repeated *leaky ReLU* function, $\Phi_{\text{LR}}(\cdot)$, defined by

$$\Phi_{\text{LR}}(u) := \begin{bmatrix} \phi_{\text{LR}}(u_1) \\ \vdots \\ \phi_{\text{LR}}(u_m) \end{bmatrix} \quad \phi_{\text{LR}}(u_i) := \begin{cases} u_i & u_i \geq 0 \\ \epsilon u_i & u_i < 0 \end{cases} \quad (5.4)$$

where $0 \leq \epsilon \ll 1$ is typically assumed. For the special case $\epsilon = 0$, Equation (5.4) defines the repeated ReLU function, $\Phi_{\text{ReLU}}(\cdot)$. Since all functions can be written in terms of their odd and even components, it is easy to see that

$$\Phi_{\text{LR}}(u) = \frac{1+\epsilon}{2}u + \frac{1-\epsilon}{2}|u| \quad (5.5)$$

where the leaky ReLU is a sum of a linear term and magnitude term. Hence, it is clear that NN stability problems involving the repeated leaky ReLU can be written as problems involving the repeated magnitude by leveraging (5.5). To simplify the exposition, $\epsilon = 0$ is assumed and transformations between (5.3) and (5.6) are derived. These transformations can easily be extended to the leaky ReLU where $\epsilon \neq 0$.

$$\begin{aligned} \dot{x} &= \tilde{A}x + \tilde{B}\Phi_{\text{ReLU}}(y) \\ y &= \tilde{C}x + \tilde{D}\Phi_{\text{ReLU}}(y) \end{aligned} \quad (5.6)$$

The following two results capture the relationships between the systems (5.3) and (5.6).

Proposition 5.1. *Assume $\tilde{E} = I - \frac{1}{2}\tilde{D} \in \mathbb{R}^{m \times m}$ is nonsingular. Then the feedback system (5.6) can be expressed in the form (5.3) via the following relationship*

$$\left[\begin{array}{c|c} A & B \\ \hline C & D \end{array} \right] := \left[\begin{array}{c|c} \tilde{A} + \frac{1}{2}\tilde{B}\tilde{E}^{-1}\tilde{C} & \frac{1}{2}\tilde{B}(I + \frac{1}{2}\tilde{E}^{-1}\tilde{D}) \\ \hline \tilde{E}^{-1}\tilde{C} & \frac{1}{2}\tilde{E}^{-1}\tilde{D} \end{array} \right] \quad (5.7)$$

Proof: Using (5.5), with $\epsilon = 0$, Equation (5.6) can be rearranged as

$$(I - \frac{1}{2}\tilde{D})y = \tilde{C}x + \frac{1}{2}\tilde{D}|y| \quad (5.8)$$

From the definition of \tilde{E} and its assumed nonsingularity

$$y = \tilde{E}^{-1}\tilde{C}x + \frac{1}{2}\tilde{E}^{-1}\tilde{D}|y| \quad (5.9)$$

The state equation of (5.6) can be re-written using (5.5) and (5.9)

$$\dot{x} = \left(\tilde{A} + \frac{1}{2}\tilde{B}\tilde{E}^{-1}\tilde{C}\right)x + \frac{1}{2}\tilde{B}\left(I + \frac{1}{2}\tilde{E}^{-1}\tilde{D}\right)|y| \quad (5.10)$$

Defining (5.7) based on (5.9) and (5.10) leads to the equivalent representation of (5.6) in the form (5.3). \square

Conversely, magnitude feedback systems (5.3) can be expressed in terms of (leaky) ReLU feedback systems (5.6).

Proposition 5.2. *Assume $E = I + D \in \mathbb{R}^{m \times m}$ is nonsingular. Then the feedback system (5.3) can be expressed in the form (5.6) via the following relationship*

$$\left[\begin{array}{c|c} \tilde{A} & \tilde{B} \\ \hline \tilde{C} & \tilde{D} \end{array} \right] := \left[\begin{array}{c|c} A - BE^{-1}C & 2B(I - E^{-1}D) \\ \hline E^{-1}C & 2E^{-1}D \end{array} \right] \quad (5.11)$$

Proof: Reversing (5.5), the output equation of (5.3) becomes

$$(I + D)y = Cx + 2D\Phi_{\text{ReLU}}(y) \quad (5.12)$$

Noting the definition and nonsingularity of E then gives

$$y = E^{-1}Cx + 2E^{-1}D\Phi_{\text{ReLU}}(y) \quad (5.13)$$

The state equation of (5.3) can also be re-written using the reversal of (5.5) and (5.13)

$$\dot{x} = (A - BE^{-1}C)x + 2B(I - E^{-1}D)\Phi_{\text{ReLU}}(y) \quad (5.14)$$

Thus, defining (5.11) based on (5.13) and (5.14) leads to the equivalent representation of (5.3) in the form (5.6). \square

Propositions 5.1 and 5.2 show that systems with (leaky) ReLU nonlinearities (5.6) can be analysed with the results derived here; and systems with magnitude nonlinearities (5.3) can be analysed with the results of Chapter 3. Both Propositions assume the existence of an inverse matrix. For an L -layer feed-forward NN, the \tilde{D} -matrix in (5.6) has a sparse strictly triangular structure (Section 2.2.3), in which case \tilde{E} exists. There are many cases when E exists, for example, if D in (5.3) is positive definite or strictly triangular. The loop transformations used in Propositions 5.1 and 5.2 are illustrated by Figure 5.1.

Remark 5.1. *The loop transformations in Figure 5.1 resemble those in [121, Figure 5] for illustrating the equivalence between the Small Gain Theorem and Passivity Theorem. Whilst*

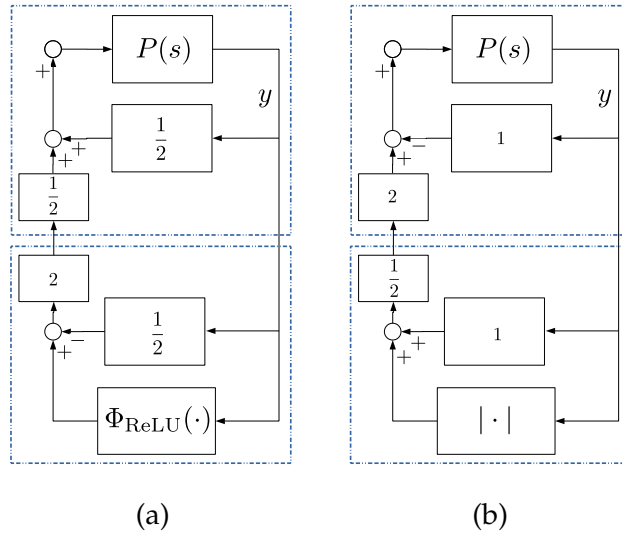


FIGURE 5.1: Loop transformations from (a) ReLU to magnitude (Proposition 5.1) and (b) magnitude to ReLU (Proposition 5.2).

each block in Figure 5.1 only requires one feedback or feed-forward connection; in [121], each block requires both. Loop transformations are used extensively in absolute stability problems, such as modifying the representation of a NN controller [21] and computing the upper bound of a NN Lipschitz constant [122]. $\square\square$

5.2.3 Well-posedness

Traditionally, absolute stability problems assume $D = 0$. This is true for Hopfield networks e.g., [92]; however, systems involving feed-forward NNs can sometimes be expressed as (5.2) where the NN weights form a sparse, non-zero D matrix, as shown in Section 2.2.3. To handle such systems, well-posedness of (5.2) must be addressed for $D \neq 0$. When $\Phi(\cdot)$ is a sector bounded nonlinearity in the Sector $[0, I]$ (Figure 5.2 shows an example where $\Phi(\cdot)$ is a 1-dimensional map), Section 2.2.2 shows that well-posedness is guaranteed if there exists a matrix $\mathbf{Y} \in \mathbb{D}_+^m$ such that

$$2\mathbf{Y} - \mathbf{Y}\tilde{D} - \tilde{D}'\mathbf{Y} \succ 0 \quad (5.15)$$

Since $\Phi_{\text{LR}}(\cdot)$ is in the Sector $[0, I]$, it follows that satisfaction of (5.15) guarantees (5.6) is well-posed. Thus, using Proposition 5.2, it is clear that (5.3) will be well-posed if (5.15) is satisfied with $\tilde{D} = 2E^{-1}D = 2(I + D)^{-1}D$. That is

$$\mathbf{Y} - \mathbf{Y}E^{-1}D - (E^{-1}D)'\mathbf{Y} \succ 0 \quad (5.16)$$

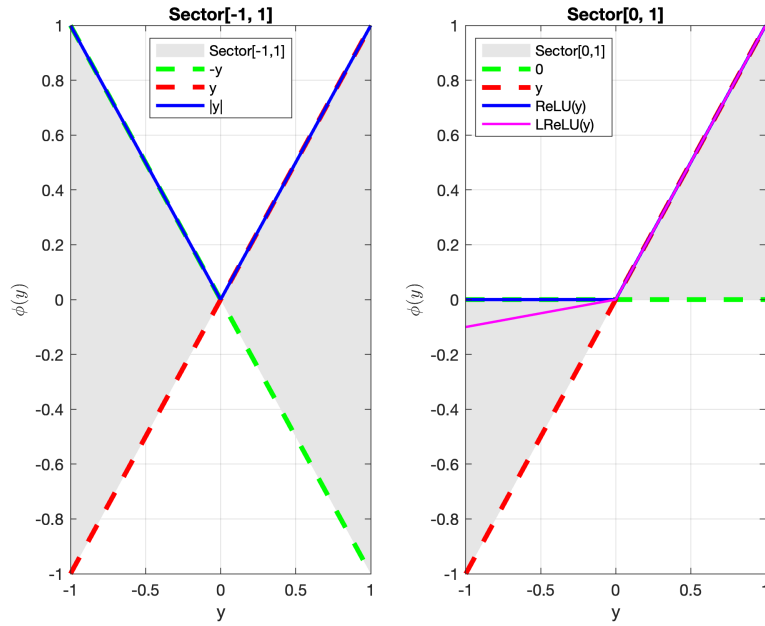


FIGURE 5.2: Sector[-1, 1] & Sector[0, 1] alongside the magnitude & (leaky) ReLU.

Applying the identity $E^{-1}D = I - E^{-1}$, followed by a congruence transform with transformation matrix E , it can be seen that this inequality is equivalent to

$$\mathbf{Y} - D'YD \succ 0 \quad (5.17)$$

Fact 5.1. *Let $E = I + D \in \mathbb{R}^{m \times m}$ be nonsingular. If there exists a $\mathbf{Y} \in \mathbb{D}_+^m$ satisfying (5.17), then Assumption 5.1 holds.*

It was noted in Chapter 3 that many feedback loops involving NNs naturally ensure that (5.15) is satisfied and thus are well-posed.

5.3 Quadratic Constraints

Quadratic constraints (QCs), which characterise the nonlinearity of the Lurie system, are the key ingredient in the derivation of many absolute stability results. They are adjoined to the derivative of the Lyapunov candidate to obtain matrix inequalities. This is exploited in the standard Circle/Popov Criteria (Section 2.3), but also other QCs have been established for different activation functions in [97], and specific QCs for the ReLU function were derived in Chapter 3. In this section, new QCs are derived for the magnitude nonlinearity.

There are three obvious properties the scalar magnitude function satisfies, for all $u_i \in \mathbb{R}$

$$|u_i| \geq 0 \quad (5.18)$$

$$|u_i| + u_i \geq 0 \quad (5.19)$$

$$|u_i| - u_i \geq 0 \quad (5.20)$$

Since these hold component-wise, it is easy to assemble these as the following QCs, characterising the repeated magnitude.

Fact 5.2 (Positivity QC). *If $\mathbf{Q} \in \mathbb{R}_{\geq 0}^{3m \times 3m}$ then the following inequality holds*

$$\begin{bmatrix} |u| \\ |u| + u \\ |u| - u \end{bmatrix}' \underbrace{\begin{bmatrix} \mathbf{Q}_{11} & \mathbf{Q}_{12} & \mathbf{Q}_{13} \\ \mathbf{Q}_{21} & \mathbf{Q}_{22} & \mathbf{Q}_{23} \\ \mathbf{Q}_{31} & \mathbf{Q}_{32} & \mathbf{Q}_{33} \end{bmatrix}}_{\mathbf{Q}} \begin{bmatrix} |u| \\ |u| + u \\ |u| - u \end{bmatrix} \geq 0 \quad \forall u \in \mathbb{R}^m \quad (5.21)$$

For any $\alpha \in \mathbb{R}$ a further property satisfied by the scalar magnitude function is

$$\alpha(|u_i| + u_i)(|u_i| - u_i) = \alpha(|u_i|^2 - u_i^2) = 0 \quad \forall u_i \in \mathbb{R} \quad (5.22)$$

This, combined with (5.19) and (5.20) yields the following fact.

Fact 5.3 (Metzler QC). *If $\mathbf{V} \in \mathbb{M}^m$, then the following inequality holds*

$$(|u| + u)' \mathbf{V} (|u| - u) \geq 0 \quad \forall u \in \mathbb{R}^m \quad (5.23)$$

Note that Fact 5.3 affords substantially more freedom in the choice of \mathbf{V} than a “small gain” bound on $|u|$. In the small gain approach, see for instance [38, Section 4.2.4], one observes that $|u_i|^2 \leq u_i^2$ and then obtains a QC of the form (5.23) with \mathbf{V} diagonal. In Fact 5.3, \mathbf{V} needs only to be Metzler, providing more freedom when solving the resulting matrix inequalities.

Remark 5.2. *Other, seemingly less useful, QCs also exist for the repeated magnitude; for example, if $\mathbf{W} \in \mathbb{D}_+^m$, it is straightforward to observe that the following inequalities*

$$u' \mathbf{W} (u + |u|) \geq 0 \quad u' \mathbf{W} (u - |u|) \geq 0 \quad \forall u \in \mathbb{R}^m$$

lack negative definite terms in $|\cdot|$, hampering the derivation of useful LMIs. □□

5.3.1 Simplifying the Positivity QC

Fact 5.2 provides nine different QCs and Fact 5.3 presents another. While it is possible to incorporate all ten constraints into the stability criteria, each constraint is associated with an extra m^2 decision variables, so it is desirable to consolidate these constraints.

Lemma 5.1 (Simplified Positivity QC). *If \mathbf{Q} satisfies Fact 5.2 and $\mathbf{Q}_{22} \in \mathbb{S}_{\geq 0}^m$, then the following inequality holds*

$$\begin{bmatrix} u \\ |u| \end{bmatrix}' \begin{bmatrix} 2\mathbf{Q}_{22} & 2\mathbf{Q}_{22} - \mathbf{Q}_{13} \\ \star & He(\mathbf{Q}_{13} + \mathbf{Q}_{22}) \end{bmatrix} \begin{bmatrix} u \\ |u| \end{bmatrix} \geq 0 \quad \forall u \in \mathbb{R}^m \quad (5.24)$$

Proof: The first step is to note the following observations

- The QCs involving \mathbf{Q}_{21} , \mathbf{Q}_{31} and \mathbf{Q}_{32} can be omitted since they duplicate those involving \mathbf{Q}_{12} , \mathbf{Q}_{13} and \mathbf{Q}_{23} .
- Secondly, the QC involving \mathbf{Q}_{23}

$$(|u| + u)' \mathbf{Q}_{23} (|u| - u) \geq 0 \quad (5.25)$$

can also be removed since it is a more restrictive version of the QC presented in Fact 5.3. That is, $\mathbf{Q}_{23} \in \mathbb{R}_{\geq 0}^{m \times m}$ whereas $\mathbf{V} \in \mathbb{M}^m$.

- Note that the matrices \mathbf{Q}_{11} , \mathbf{Q}_{22} and \mathbf{Q}_{33} can, without loss of generality, be taken as symmetric since

$$w_i' \mathbf{Q}_{ii} w_i = \frac{1}{2} w_i' (\mathbf{Q}_{ii} + \mathbf{Q}_{ii}') w_i \quad (5.26)$$

where $i \in \{1, 2, 3\}$ and w_i represents the associated term $|u|$ or $|u| \pm u$.

- Finally, it is also possible to remove the QC involving \mathbf{Q}_{11} ($|u|' \mathbf{Q}_{11} |u| \geq 0$). When adjoined to the derivative of the Lyapunov candidate, this QC does nothing to assist in making the derivative negative definite.

Applying each of these observations, condenses (5.21) to

$$\begin{bmatrix} u \\ |u| \end{bmatrix}' \begin{bmatrix} \mathbf{Q}_{22} + \mathbf{Q}_{33} & \frac{1}{2}(\mathbf{Q}_{12} - \mathbf{Q}_{13}) + \mathbf{Q}_{22} - \mathbf{Q}_{33} \\ \star & \frac{1}{2} He(\mathbf{Q}_{12} + \mathbf{Q}_{13}) + \mathbf{Q}_{22} + \mathbf{Q}_{33} \end{bmatrix} \begin{bmatrix} u \\ |u| \end{bmatrix} \geq 0 \quad \forall u \in \mathbb{R}^m \quad (5.27)$$

This can be further simplified, with the loss of some generality, by noting that \mathbf{Q}_{33} plays the same role as \mathbf{Q}_{22} in the diagonal entries and \mathbf{Q}_{12} plays the same role as \mathbf{Q}_{13} in the (2,2) entry. In the off-diagonal blocks, \mathbf{Q}_{33} plays the same role as \mathbf{Q}_{13} and \mathbf{Q}_{12} plays the same role as \mathbf{Q}_{22} . Therefore, the above inequality can be replaced by (5.24). \square

5.3.2 Sector-like Lemma for Magnitude Nonlinearities

The following Lemma will be a useful addition to the previous QCs.

Lemma 5.2. *Suppose that $\mathbf{H} \in \mathbb{R}_{\geq 0}^{m \times m}$ is invertible. Define $\Psi(\mathbf{H}, u) := \mathbf{H}^{-1}|\mathbf{H}u|$, then the following inequality holds*

$$\Psi(\mathbf{H}, u)' \mathbf{W}(|u| - \Psi(\mathbf{H}, u)) \geq 0 \quad \forall y \in \mathbb{R}^m \quad (5.28)$$

for all $\mathbf{W} = \mathbf{H}'\mathbf{M}\mathbf{H}$ where $\mathbf{M} \in \mathbb{R}_{\geq 0}^{m \times m}$.

Proof: Denoting each row of $\mathbf{H} \in \mathbb{R}_{\geq 0}^{m \times m}$ by \mathbf{H}'_i

$$\begin{aligned} |\mathbf{H}'_i u| &= \left| \sum_{j=1}^m \mathbf{H}_{ij} u_j \right| \leq \sum_{j=1}^m |\mathbf{H}_{ij} u_j| \\ &= \sum_{j=1}^m \mathbf{H}_{ij} |u_j| = \mathbf{H}'_i |u| \end{aligned} \quad (5.29)$$

As this applies to each row of \mathbf{H} , one can conclude $|\mathbf{H}u| \leq \mathbf{H}|u|$.

Next, \mathbf{H} is introduced to define a generalised form of the QC involving \mathbf{Q}_{13} , from Fact 5.2, in terms of Ψ .

$$\begin{aligned} S(\mathbf{W}, \mathbf{H}, u) &:= \Psi(\mathbf{H}, u)' \mathbf{W}(|u| - \Psi(\mathbf{H}, u)) \\ &= (\mathbf{H}^{-1}|\mathbf{H}u|)' \mathbf{W}(|u| - \mathbf{H}^{-1}|\mathbf{H}u|) \end{aligned} \quad (5.30)$$

Expanding the terms leads to

$$S(\mathbf{W}, \mathbf{H}, u) = |\mathbf{H}u|' (\mathbf{H}')^{-1} \mathbf{W} \mathbf{H}^{-1} \mathbf{H} |u| - |\mathbf{H}u|' (\mathbf{H}')^{-1} \mathbf{W} \mathbf{H}^{-1} |\mathbf{H}u| \quad (5.31)$$

Substituting for \mathbf{M} and leveraging $|\mathbf{H}u| \leq \mathbf{H}|u|$

$$\begin{aligned} S(\mathbf{W}, \mathbf{H}, u) &= |\mathbf{H}u|' \mathbf{M} \mathbf{H} |u| - |\mathbf{H}u|' \mathbf{M} |\mathbf{H}u| \\ &\geq |\mathbf{H}u|' \mathbf{M} |\mathbf{H}u| - |\mathbf{H}u|' \mathbf{M} |\mathbf{H}u| = 0 \end{aligned} \quad (5.32)$$

□

It should be noted that this QC has the same form as the sector QC, so it is referred to as the “sector-like” inequality.

5.4 Global Stability Analysis

This section reports global stability conditions based on quadratic and Lurie-type Lyapunov candidates (Section 2.3).

5.4.1 Quadratic Lyapunov Function

Theorem 5.1 (Quadratic Criterion). *Consider the feedback system (5.3) and let Assumption 5.1 be satisfied. If there exists $\mathbf{P} \in \mathbf{S}_+^n$, $\mathbf{V} \in \mathbb{M}^m$, $\mathbf{Q}_{13} \in \mathbb{R}_{\geq 0}^{m \times m}$ and $\mathbf{Q}_{22} \in \mathbf{S}_{\geq 0}^m$ such that*

$$F_Q(\mathbf{P}, \mathbf{V}, \mathbf{Q}_{13}, \mathbf{Q}_{22}) \prec 0 \quad (5.33)$$

where

$$F_Q = \begin{bmatrix} \text{He}(\mathbf{P}\mathbf{A} - \mathbf{C}'(\mathbf{V} - \mathbf{Q}_{22})\mathbf{C}) & \mathbf{P}\mathbf{B} - \mathbf{C}'\mathbf{V}'(\mathbf{I} + \mathbf{D}) + \mathbf{C}'\mathbf{V}(\mathbf{I} - \mathbf{D}) + \mathbf{C}'\mathbf{Q}_{22}\mathbf{D} + \mathbf{C}'(2\mathbf{Q}_{22} - \mathbf{Q}_{13}) \\ \star & \text{He}((\mathbf{I} + \mathbf{D})'\mathbf{V}(\mathbf{I} - \mathbf{D}) + \mathbf{Q}_{22} + \mathbf{Q}_{13} + \mathbf{D}'\mathbf{Q}_{22}\mathbf{D} + \mathbf{D}'(2\mathbf{Q}_{22} - \mathbf{Q}_{13})) \end{bmatrix} \quad (5.34)$$

then the origin of (5.3) is GAS.

Proof: Consider the quadratic Lyapunov function $V_Q(x) = x'\mathbf{P}x$ and consider its derivative along the trajectories of (5.3). Appending the QCs from Fact 5.3, denoted by $M(y, \mathbf{V})$, and Lemma 5.1, denoted by $SP(y, \mathbf{Q}_{13}, \mathbf{Q}_{22})$, is equivalent to the final quadratic inequality below.

$$\begin{aligned} \dot{V}_Q(x) &= 2x'\mathbf{P}(Ax + B|y|) \\ &\leq 2x'\mathbf{P}(Ax + B|y|) + 2M(y, \mathbf{V}) + SP(y, \mathbf{Q}_{13}, \mathbf{Q}_{22}) \\ &= \begin{bmatrix} x \\ |y| \end{bmatrix}' F_Q(\mathbf{P}, \mathbf{V}, \mathbf{Q}_{13}, \mathbf{Q}_{22}) \begin{bmatrix} x \\ |y| \end{bmatrix} \end{aligned} \quad (5.35)$$

□

Remark 5.3. *Theorem 5.1 is in the form of an LMI and bears some resemblance to the LMI associated with the Small Gain Theorem [14]. In the small gain approach, which for the magnitude nonlinearity is equivalent to considering diagonal norm-bounded constraints as described in [38, Section 5.1], one has the matrix \mathbf{V} simply being diagonal and constraints associated with the matrices \mathbf{Q}_{ij} do not exist; that is $\mathbf{Q}_{ij} \equiv 0$. In this case, Theorem 5.1 simply reduces to*

$$\begin{bmatrix} \mathbf{P}\mathbf{A} + \mathbf{A}'\mathbf{P} - \mathbf{C}'\mathbf{V}\mathbf{C} & \mathbf{P}\mathbf{B} - \mathbf{C}'\mathbf{V}'\mathbf{D} \\ \star & \mathbf{V} - \mathbf{D}'\mathbf{V}\mathbf{D} \end{bmatrix} \prec 0 \quad (5.36)$$

This is in the same form as the matrix inequality in [38, Eq. 5.15] and is clearly a special case of (5.33). Therefore, Theorem 5.1 is less conservative than the Small Gain Theorem for systems of the form (5.3). Satisfying (5.36) also automatically implies well-posedness of (5.3) since the well-posedness condition (5.17) naturally arises in the (2, 2) block. $\square\square$

5.4.2 Lurie-type Lyapunov Function

The Popov Criterion (see, for instance [70; 71]) uses the following Lurie-type Lyapunov function to derive a stability criterion

$$V_a(x) = x'Px + 2 \int_0^y \Lambda \Phi(\sigma) \cdot d\sigma \quad (5.37)$$

with $\Lambda \in \mathbb{D}_+^m$. However, when $\Phi(\sigma) \equiv |\sigma|$, it is clear that such a choice is not appropriate. As shown by the scalar case, the integral term

$$2 \int_0^y |\sigma| d\sigma = |y| y \quad (5.38)$$

is not guaranteed to be non-negative. Instead, the following alternative choice of Lyapunov function could be made

$$V_b(x) = x'Px + 2 \int_0^{\mathbf{H}y} \Lambda(|\sigma| + \sigma) \cdot d\sigma \quad (5.39)$$

again with $\Lambda \in \mathbb{D}_+^m$. The second term can be directly integrated to get

$$V_b(x) = x'Px + 2 \sum_{i=1}^m \int_0^{\mathbf{H}_i y} \Lambda_{ii}(|\sigma_i| + \sigma_i) d\sigma_i = x'Px + \sum_{i=1}^m \Lambda_{ii} \left[|\sigma_i|(\sigma_i + |\sigma_i|) \right]_0^{\mathbf{H}_i y} \quad (5.40)$$

where each row of \mathbf{H} is denoted by \mathbf{H}_i' . The final term is positive by virtue of inequality (5.19). In fact, positivity of this Lyapunov function is not surprising since (5.5) implies

$$V_b(x) = x'Px + 4 \int_0^{\mathbf{H}y} \Lambda \Phi_{\text{ReLU}}(\sigma) \cdot d\sigma \quad (5.41)$$

This has the same form as the Lurie-type Lyapunov function used in Chapter 3.

Unfortunately, the Lyapunov function (5.39) still requires the matrix Λ to be positive definite, despite the extra freedom afforded by the matrix \mathbf{H} . For this reason, the following Lyapunov function is proposed

$$V_L(x) = x' \mathbf{P} x + 2 \int_0^{\mathbf{H}_1 y} \boldsymbol{\Lambda}_1 (|\sigma| + \sigma) \cdot d\sigma + 2 \int_0^{\mathbf{H}_2 y} \boldsymbol{\Lambda}_2 (\sigma - |\sigma|) \cdot d\sigma \quad (5.42)$$

with $\boldsymbol{\Lambda}_1, \boldsymbol{\Lambda}_2 \in \mathbb{D}_+^m$. These two integral terms are useful because they lead to terms $\boldsymbol{\Lambda}_1 - \boldsymbol{\Lambda}_2$ in the resulting matrix inequality, potentially inducing a reduction in conservatism. The same approach was used in [70; 71] to, effectively, enable $\boldsymbol{\Lambda}$ to be indefinite. The integrals in $V_L(x)$ can be evaluated to obtain

$$V_L(x) = x' \mathbf{P} x + \sum_{i=1}^m \boldsymbol{\Lambda}_{1,ii} \left[|\sigma_i| (|\sigma_i| + \sigma_i) \right]_0^{\mathbf{H}_{1,i} y} + \sum_{i=1}^m \boldsymbol{\Lambda}_{2,ii} \left[|\sigma_i| (|\sigma_i| - \sigma_i) \right]_0^{\mathbf{H}_{2,i} y} \quad (5.43)$$

The final term is positive by virtue of inequality (5.20). Furthermore, this final expression verifies that the integrals in (5.42) are path independent [102; 103]. Using the Lyapunov function $V_L(x)$, the following result can be derived.

Theorem 5.2 (Lurie-based Criterion). *Consider the feedback system (5.3) and let $D = 0$. If there exists $\mathbf{P} \in \mathbb{S}_+^n$; $\boldsymbol{\Lambda}_1, \boldsymbol{\Lambda}_2 \in \mathbb{D}_+^m$; $\mathbf{V} \in \mathbb{M}^m$; $\mathbf{Q}_{22} \in \mathbb{S}_{\geq 0}^m$; $\mathbf{Q}_{13}, \mathbf{M}, \mathbf{H} \in \mathbb{R}_{\geq 0}^{m \times m}$ with \mathbf{H} nonsingular such that*

$$F_L(\mathbf{P}, \mathbf{H}, \boldsymbol{\Lambda}_1, \boldsymbol{\Lambda}_2, \mathbf{M}, \mathbf{V}, \mathbf{Q}_{22}, \mathbf{Q}_{13}) \prec 0 \quad (5.44)$$

where

$$F_L(\cdot) = \begin{bmatrix} \text{He}(\mathbf{P}\mathbf{A} + \mathbf{C}'\mathbf{H}'(\boldsymbol{\Lambda}_1 + \boldsymbol{\Lambda}_2)\mathbf{H}\mathbf{C}\mathbf{A} - \mathbf{C}'(\mathbf{V} - \mathbf{Q}_{22})\mathbf{C}) & * & * \\ \mathbf{B}'\mathbf{P} + \mathbf{B}'\mathbf{C}'\mathbf{H}'(\boldsymbol{\Lambda}_1 + \boldsymbol{\Lambda}_2)\mathbf{H}\mathbf{C} - \mathbf{V}\mathbf{C} + \mathbf{V}'\mathbf{C} + (2\mathbf{Q}_{22} - \mathbf{Q}_{13})'\mathbf{C} & \text{He}(\mathbf{V} + \mathbf{Q}_{22} + \mathbf{Q}_{13}) & * \\ \mathbf{H}'(\boldsymbol{\Lambda}_1 - \boldsymbol{\Lambda}_2)\mathbf{H}\mathbf{C}\mathbf{A} & \mathbf{H}'(\boldsymbol{\Lambda}_1 - \boldsymbol{\Lambda}_2)\mathbf{H}\mathbf{C}\mathbf{B} + \mathbf{H}'\mathbf{M}\mathbf{H} & -\text{He}(\mathbf{H}'\mathbf{M}\mathbf{H}) \end{bmatrix} \quad (5.45)$$

then the origin of (5.3) is GAS.

Proof: At the expense of some conservatism, it is assumed that $\mathbf{H} := \mathbf{H}_1 = \mathbf{H}_2$, leading to the following derivative of the Lyapunov function (5.42)

$$\dot{V}_L(x) = 2x' \mathbf{P} \dot{x} + 2 \left(y' \mathbf{H}' (\boldsymbol{\Lambda}_1 + \boldsymbol{\Lambda}_2) + |\mathbf{H}y|' (\boldsymbol{\Lambda}_1 - \boldsymbol{\Lambda}_2) \right) \mathbf{H}\mathbf{C} \dot{x} \quad (5.46)$$

Appending the QCs from Fact 5.3, Lemma 5.1 and Lemma 5.2, respectively denoted by $M(y, \mathbf{V})$, $SP(y, \mathbf{Q}_{13}, \mathbf{Q}_{22})$ and $S(y, \mathbf{W}, \mathbf{H})$ results in

$$\dot{V}_L(x) \leq \dot{V}_L(x) + 2M(y, \mathbf{V}) + SP(y, \mathbf{Q}_{13}, \mathbf{Q}_{22}) + 2S(y, \mathbf{W}, \mathbf{H}) \quad (5.47)$$

After some algebra, the above inequality can be expressed in the quadratic form

$$\dot{V}_L(x) \leq \begin{bmatrix} x \\ |y| \\ \Psi \end{bmatrix}' F_L(\cdot) \begin{bmatrix} x \\ |y| \\ \Psi \end{bmatrix} \quad (5.48)$$

where $\Psi = \Psi(\mathbf{H}, y)$ and $F_L(\cdot)$ is defined in (5.45). \square

Remark 5.4. It is possible to enhance Theorem 5.2 further by noting that, since \mathbf{H} is assumed positive and invertible, the quadratic inequalities (5.18)-(5.20) can be used to obtain

$$\mathbf{H}\Psi(\mathbf{H}, u) \geq 0 \quad (5.49)$$

$$\mathbf{H}\Psi(\mathbf{H}, u) + \mathbf{H}u \geq 0 \quad (5.50)$$

$$\mathbf{H}\Psi(\mathbf{H}, u) - \mathbf{H}u \geq 0 \quad (5.51)$$

where $\Psi(\mathbf{H}, u)$ is defined in Lemma 5.2. By application of Fact 5.2, this implies that for all $\mathbf{R} \in \mathbb{R}_{\geq 0}^{3m \times 3m}$ and for all $u \in \mathbb{R}^m$

$$\begin{bmatrix} \mathbf{H}\Psi \\ \mathbf{H}\Psi + \mathbf{H}u \\ \mathbf{H}\Psi - \mathbf{H}u \end{bmatrix}' \underbrace{\begin{bmatrix} \mathbf{R}_{11} & \mathbf{R}_{12} & \mathbf{R}_{13} \\ \mathbf{R}_{21} & \mathbf{R}_{22} & \mathbf{R}_{23} \\ \mathbf{R}_{31} & \mathbf{R}_{32} & \mathbf{R}_{33} \end{bmatrix}}_{\mathbf{R}} \begin{bmatrix} \mathbf{H}\Psi \\ \mathbf{H}\Psi + \mathbf{H}u \\ \mathbf{H}\Psi - \mathbf{H}u \end{bmatrix} \geq 0 \quad (5.52)$$

where the shorthand $\Psi = \Psi(\mathbf{H}, u)$ has been used. Simplifying in the same way as Lemma 5.1, the reduced set of QCs becomes

$$\begin{bmatrix} u \\ \Psi \end{bmatrix}' \begin{bmatrix} 2\mathbf{H}'\mathbf{R}_{22}\mathbf{H} & \mathbf{H}'(2\mathbf{R}_{22} - \mathbf{R}_{13})\mathbf{H} \\ \star & \mathbf{H}'\text{He}(\mathbf{R}_{13} + \mathbf{R}_{22})\mathbf{H} \end{bmatrix} \begin{bmatrix} u \\ \Psi \end{bmatrix} \geq 0 \quad (5.53)$$

Using the S-procedure, this constraint could additionally be adjoined to (5.48), although at the expense of adding several more decision variables to the arising matrix inequality. $\square\square$

Remark 5.5. For $D = 0$, Theorem 5.1 is a special case of Theorem 5.2 with $\Lambda_1 = \Lambda_2 = \mathbf{W} = 0$. Hence, Theorem 5.2 will always be less conservative under this condition. $\square\square$

Remark 5.6. The key advantage of Theorem 5.1 over Theorem 5.2 is that it accounts for $D \neq 0$ and can be applied to systems involving feed-forward NNs. When $D = 0$, such as RNNs and bridge rectifiers, Theorem 5.2 will be less conservative. $\square\square$

5.5 Convex Relaxations for Theorem 5.2

Theorem 5.2 is potentially much less conservative than Theorem 5.1, but suffers from a crucial weakness: (5.44) is a nonlinear matrix inequality since it contains products of two or three matrix variables. This section presents two relaxations which recover LMIs from (5.44) and make Theorem 5.2 tractable.

5.5.1 Specific Choice of Λ and M

Corollary 5.1. *Consider the feedback system (5.3) and let $D = 0$. If there exist $\mathbf{P} \in \mathbb{S}_+^n$; $\mathbf{X}, \mathbf{Q}_{22} \in \mathbb{S}_{\geq 0}^m$; $\mathbf{V} \in \mathbb{M}^m$; $\mathbf{Q}_{13} \in \mathbb{R}_{\geq 0}^{m \times m}$ and a scalar $\mu > 0$ such that*

$$G(\mathbf{P}, \mathbf{X}, \mathbf{V}, \mathbf{Q}_{22}, \mathbf{Q}_{13}, \mu) \prec 0 \quad (5.54)$$

where

$$G(\cdot) = \begin{bmatrix} \text{He}(\mathbf{P}\mathbf{A} + 2\mathbf{C}'\mathbf{X}\mathbf{C}\mathbf{A} - \mathbf{C}'(\mathbf{V} - \mathbf{Q}_{22})\mathbf{C}) & \star & \star \\ \mathbf{B}'\mathbf{P} + 2\mathbf{B}'\mathbf{C}'\mathbf{X}\mathbf{C} - \mathbf{V}\mathbf{C} + \mathbf{V}'\mathbf{C} + (2\mathbf{Q}_{22} - \mathbf{Q}_{13})'\mathbf{C} & \text{He}(\mathbf{V} + \mathbf{Q}_{22} + \mathbf{Q}_{13}) & \star \\ 0 & \mu\mathbf{X} & -\text{He}(\mu\mathbf{X}) \end{bmatrix} \quad (5.55)$$

then the origin of (5.3) is GAS.

Proof: The matrix inequality (5.54) is a special case of (5.44) where $\Lambda_1 = \Lambda_2 = I$, $\mathbf{M} = \mu I$, and $\mathbf{X} := \mathbf{H}'\mathbf{H} \in \mathbb{S}_{\geq 0}^m$. \square

With this choice, the only remaining bilinear term is $\mu\mathbf{X}$. Hence, (5.54) can be solved as an LMI plus a line search.

5.5.2 Specific Choice of \mathbf{H}

Corollary 5.2. *Consider the feedback system (5.3) and let $D = 0$. If there exists $\mathbf{P} \in \mathbb{S}_+^n$; $\Lambda_1, \Lambda_2 \in \mathbb{D}_+^m$; $\mathbf{V} \in \mathbb{M}^m$; $\mathbf{Q}_{22} \in \mathbb{S}_{\geq 0}^m$ and $\mathbf{Q}_{13} \in \mathbb{R}_{\geq 0}^{m \times m}$ such that*

$$J(\mathbf{P}, \Lambda_1, \Lambda_2, \mathbf{V}, \mathbf{Q}_{22}, \mathbf{Q}_{13}) \prec 0 \quad (5.56)$$

where

$$J(\cdot) = \begin{bmatrix} \text{He}(\mathbf{P}\mathbf{A} + \mathbf{C}'(\Lambda_1 + \Lambda_2)\mathbf{C}\mathbf{A} - \mathbf{C}'(\mathbf{V} - \mathbf{Q}_{22})\mathbf{C}) & \star \\ \mathbf{B}'\mathbf{P} + \mathbf{B}'\mathbf{C}'(\Lambda_1 + \Lambda_2)\mathbf{C} - \mathbf{V}\mathbf{C} + \mathbf{V}'\mathbf{C} + (2\mathbf{Q}_{22} - \mathbf{Q}_{13})'\mathbf{C} + (\Lambda_1 - \Lambda_2)\mathbf{C}\mathbf{A} & \text{He}(\mathbf{V} + \mathbf{Q}_{22} + \mathbf{Q}_{13} + (\Lambda_1 - \Lambda_2)\mathbf{C}\mathbf{B}) \end{bmatrix} \quad (5.57)$$

then the origin of (5.3) is GAS.

Proof: The LMI (5.56) is a special case of (5.44) when $\mathbf{H} = I$. Note that this relaxation implies $\Psi(I, y) = |y|$ which makes Lemma 5.2 redundant. \square

To motivate this choice, first restrict $\mathbf{H} \in \mathbb{D}_+^m$. This implies that the products $\mathbf{H}'\Lambda_1\mathbf{H}$, $\mathbf{H}'\Lambda_2\mathbf{H} \in \mathbb{D}_+^m$ and the product $\mathbf{H}'\mathbf{M}\mathbf{H} \in \mathbb{R}_{\geq 0}^{m \times m}$. Since these matrices do not appear elsewhere in inequality (5.45), the choice $\mathbf{H} = I$ can be made without loss of generality. The appeal of the matrix inequality (5.56) is that it is entirely *linear*, making Corollary 5.2 preferable to Corollary 5.1.

Remark 5.7. *Theorem 5.1 is also a special case of Corollary 5.2, when $D = 0$, under the same relaxations as Remark 5.5. Hence, Corollary 5.2 will always be less conservative.* $\square\square$

5.6 Equivalence of Local and Global Stability

It is obvious that the magnitude nonlinearity is *positive homogenous*, as is the (leaky) ReLU nonlinearity; this implies $\alpha|u| = |\alpha u|$ for all scalars $\alpha > 0$. This fact can be harnessed to prove that, under mild conditions, if $x = 0$ is a local equilibrium point of the system (5.3), then it will in fact be a global equilibrium point. This is unusual for typical Lurie systems where, if a global analysis fails, one might still try to establish local stability for a given region of attraction. The first step in establishing that global and local stability are equivalent is the following lemma.

Lemma 5.3. *Assume $(I - DU)$ is invertible and define*

$$\mathcal{H} = \left\{ A + BU(I - DU)^{-1}C : U \in \mathcal{U} \right\} \quad (5.58)$$

where

$$\mathcal{U} = \left\{ \text{diag}(u_1, \dots, u_m) \mid u_i \in -1, 1 \text{ and } i \in 1, \dots, m \right\} \quad (5.59)$$

If all matrices in \mathcal{H} are full rank, then $x = 0$ is a unique equilibrium point of (5.3).

Proof: For $U \in \mathcal{U}$ and by definition of the magnitude function, the repeated magnitude may be expressed as

$$|y| = Uy \quad (5.60)$$

Now, at equilibrium, the state equation of (5.3) becomes

$$0 = Ax + B|y| \quad (5.61)$$

Assume $x \neq 0$ is an equilibrium point, then (5.60) is leveraged to equivalently express (5.61) as

$$0 = (A + BU(I - DU)^{-1}C)x \quad \forall U \in \mathcal{U} \quad (5.62)$$

However, as all matrices in \mathcal{H} are full rank; there is a contradiction and hence $x \neq 0$ cannot be an equilibrium point. Thus, $x = 0$ is a unique equilibrium point of (5.3). \square

Remark 5.8. *With a minor modification, Lemma 5.3 also holds for Lurie systems with ReLU nonlinearities. That is, if $(I - D\tilde{U})$ is invertible, where*

$$\tilde{U} = \left\{ \text{diag}(u_1, \dots, u_m) \mid u_i \in 0, 1 \text{ and } i \in 1, \dots, m \right\} \quad (5.63)$$

then $x = 0$ is a unique equilibrium point of (3.1) if all matrices in \mathcal{H} are full rank. $\square\square$

Remark 5.9. *A sufficient condition for $(I - DU)^{-1}$ to exist is for Fact 5.1 to be satisfied i.e., (5.3) is well-posed. This can simply be seen by leveraging (5.60) to show that under Fact 5.1, a solution to $y = Cx + DUy$ must exist for all $x \in \mathbb{R}^n$ and $U \in \mathcal{U}$; hence, $(I - DU)$ must be nonsingular $\forall U \in \mathcal{U}$.* $\square\square$

As an example, Hopfield networks [92] with ReLU activations can be expressed in the form (5.3) where $A = -I + \frac{1}{2}\tilde{B}$, $B = \frac{1}{2}\tilde{B}$, $C = I$, $D = 0$, and $\tilde{B} \in \mathbb{R}^{n \times n}$. Therefore, Lemma 5.3 requires $0.5\tilde{B}(U + I) - I$ to be full rank for all $U \in \mathcal{U}$; this will typically be the case unless \tilde{B} has a very special structure. Lemma 5.3 can then be used to establish the following interesting result.

Theorem 5.3. *Consider the Lurie system (5.3) and let Assumption 5.1 be satisfied. If $x = 0$ is a unique equilibrium point and it is locally stable, then it is also globally stable.*

Proof: Assume $x = 0$ is a locally stable equilibrium point of (5.3), this implies the existence of a ball

$$\mathcal{B}(x, c) := \left\{ x \in \mathbb{R}^n : \|x\| < c \right\} \quad (5.64)$$

such that $\forall x(0) \in \mathcal{B}(x, c)$ it follows that $\lim_{t \rightarrow \infty} x(t) = 0$; that is, $\mathcal{B}(x, c)$ is a region of attraction of $x = 0$. Assumption 5.1 implies that a unique solution, y , exists to the implicit equation

$$\theta(y) := y - D|y| = Cx \quad (5.65)$$

for all x ; equivalently, $\theta(\cdot) : \mathbb{R}^m \mapsto \mathbb{R}^m$ is invertible. It is easy to see that $\theta(\cdot)$ is positive homogenous and as $\theta^{-1}(\cdot)$ is assumed to exist, it must also be positive homogenous (see Fact 4.2). The Lurie system (5.3) can thus be written as

$$\dot{x} = Ax + B|\theta^{-1}(Cx)| \quad (5.66)$$

Defining the change of coordinates $z := \alpha x$ for $\alpha > 0$, the dynamics in the scaled coordinates are

$$\dot{z} = \alpha Ax + \alpha B|\theta^{-1}(Cx)| = Az + B|\theta^{-1}(Cz)| \quad (5.67)$$

As this is of the same form as (5.66), it must have an equilibrium point $z = 0$ and associated region of attraction $\mathcal{B}(z, c)$. Scrutinising the region of attraction

$$\begin{aligned} \mathcal{B}(z, c) &= \left\{ z \in \mathbb{R}^n : \|z\| < c \right\} \\ &= \left\{ x \in \mathbb{R}^n : \|x\| < \frac{c}{\alpha} \right\} \\ &= \mathcal{B}\left(x, \frac{c}{\alpha}\right) \end{aligned} \quad (5.68)$$

Since α can be arbitrary small, it is clear that $\mathcal{B}(x, c/\alpha) \rightarrow \mathbb{R}^n$ as $\alpha \rightarrow 0$. □

Remark 5.10. *Since the ReLU function is also positively homogenous, a minor adaptation of Theorem 5.3 also holds for Lurie systems with ReLU nonlinearities. To state this explicitly, consider the Lurie system (3.1) and let Assumption 3.1 be satisfied. If $x = 0$ is a unique equilibrium point and it is locally stable, then it is also globally stable.* □□

If it is possible to establish $x = 0$ as a unique locally stable equilibrium point of (5.3), then this section has shown that it is in fact globally stable. Equivalently, if the conditions of Lemma 5.3 are satisfied, but it is not possible to prove global stability of (5.3), then it is futile to attempt a local stability analysis. Remark 5.8 and Remark 5.10 prove similar results for Lurie systems containing the ReLU nonlinearity.

5.7 Numerical Examples

The conservatism and complexity of Theorem 5.1 and Corollary 5.2, which take the form of LMIs, were tested using the same numerical examples that appeared in Chapter 3. In Chapter 3, the systems under consideration were Lurie systems with a ReLU nonlinearity, as in (5.6). To provide a useful comparison, these examples were transformed

TABLE 5.1: Comparison of the maximum series gain.

Ex	SGT	Theorem 3.1	Corollary 3.1	Corollary 3.2	Park	Zames-Falb	Theorem 5.1	Corollary 5.2	Nyquist Gain
1	20.8766	39.5200	10,953.77	100,000+	448.6800	434.2800	40.1000	90,283.38	100,000+
2	89.9000	89.9000	89.9000	89.9000	89.9000	89.9000	89.9000	89.9000	89.9000
3	0.5236	0.6818	0.6818	0.6818	0.5236	0.5522	0.6818	0.6818	0.6983
4	0.0010	0.0012	0.0012	0.0015	0.0010	0.0010	0.0012	0.0015	0.0020
5	0.0813	0.0814	0.0814	0.0830	0.0845	0.0845	0.0814	0.0845	0.0869
6	0.1946	0.3901	0.4155	0.5048	0.2266	0.2773	0.4082	0.6129	0.8202
7	0.0966	0.1232	0.1232	0.1462	0.1035	0.1035	0.1232	0.1502	0.2002
8	2.0221	2.0221	2.0221	2.0221	2.0221	2.0221	2.0221	2.0221	2.0221
9	1.4695	2.0516	2.0516	2.0516	1.4695	1.4695	2.0516	2.0516	2.0600
10	1.3820	2.0224	2.0343	2.2240	1.3820	1.3820	2.0699	2.1088	2.7000
11	1.4531	2.1189	2.0449	2.2092	1.4605	1.4605	2.1189	2.2067	2.2200
12	1.4613	2.2144	2.1582	2.3380	1.4613	1.4613	2.2256	2.3099	2.7500

into the form (5.3) using the loop transformation from Section 5.2.2. The comparison was performed by inserting a series gain ($\alpha \in \mathbb{R}_{\geq 0}$) into the feedback loop, replacing B with αB and D with αD , then computing the maximum series gain for which stability was guaranteed. The Nyquist gain provides an upper bound on this quantity.

For each criterion, finding the maximum series gain was implemented as a bisection method in conjunction with the YALMIP parser [43] and MOSEK solver [42]. Various criteria were compared to Theorem 5.1 and Corollary 5.2: the classical Small Gain Theorem [38, Section 5.1], the ReLU results of Chapter 3 (where Corollary 3.1 was applied as described in Section 3.6.1), which improved on the standard Circle and Popov Criteria, and the more advanced approaches of Park [72] and Zames-Falb, using the multiplier search of [106].

The maximum series gain results are presented in Table 5.1 and the number of decision variables are detailed in Table 5.2. The example systems are detailed in Table 3.2 and can be found by looking at the related code³. The first eight examples are typical low-dimensional systems from the control literature, whereas the last four examples are high-dimensional Hopfield network's.

5.7.1 Discussion

It is noteworthy that, in 7 out of the first 8 examples, Corollary 5.2 is better than or equal to all other criteria at predicting the largest value of α for which stability is guaranteed. However, Corollary 3.2 marginally outperforms Corollary 5.2 in the high-dimensional examples. Hence, the main advantage of Corollary 5.2 over Corollary 3.2, for high-dimensional systems, is that it can be applied to a wider class of activation functions (ReLU and Leaky ReLU) via the loop transformations in Section 5.2.2.

Another interesting observation is that Theorem 5.1 is better than or equal to Theorem 3.1 in all examples. This is an important comparison since these are the least conservative criteria which can be applied to systems with $D \neq 0$. Finally, when considering the high-dimensional systems, all new criteria significantly outperform the previous state of the art in Park [72] and Zames-Falb [106].

³<https://github.com/CR-Richardson/Max-Series-Gain-Magnitude>

TABLE 5.2: Comparison of the decision variable count.

Ex	SGT	Theorem 3.1	Corollary 3.1	Corollary 3.2	Park	Zames-Falb	Theorem 5.1	Corollary 5.2
1	48	63	81	66	87	252	69	75
2	9	24	42	27	30	45	30	36
3	10	38	70	42	40	48	48	56
4	40	68	100	72	90	208	78	86
5	25	53	85	57	67	129	63	71
6	25	53	85	57	67	129	63	71
7	40	68	100	72	90	208	78	86
8	20	65	115	70	70	100	80	90
9	860	4,020	7,220	4,060	3,360	4,300	4,840	4,920
10	1,890	9,030	16,230	9,090	7,440	9,450	10,860	10,980
11	3,320	16,040	28,840	16,120	13,120	16,600	19,280	19,440
12	5,150	25,050	45,050	25,150	20,400	25,750	30,100	30,300

The LMI in Corollary 5.2 contains more decision variables than that in Corollary 3.2. This is also true when comparing the LMI in Theorem 5.1 with Theorem 3.1. Whilst this is an unwelcome deficiency, the number of decision variables is only one of several factors which contribute to the time taken to solve an LMI (Section 2.1.3). To offer a more thorough comparison, the time required to satisfy each criterion, for Example's 9-12, is shown in Figure 5.3. This corresponds to setting $\alpha = 1$, where each criterion was able to verify GAS. The computation was performed on an Apple M1 Pro. The * denotes that Theorem 3.1, Corollary 3.2, Theorem 5.1, and Corollary 5.2 were applied with $\mathbf{V}, \mathbf{Q}_{ij}$ reduced to being symmetric; furthermore, the $^+$ denotes Corollary 3.1 was applied with all matrices being reduced to symmetric.

Figure 5.3 reinforces that Theorem 5.1 and Corollary 5.2 are the slowest to verify GAS of each Hopfield network. These criteria follow similar scaling trends to Corollary 3.1 $^+$, Corollary 3.2, and Theorem 3.1, which are slightly slower than Zames-Falb when $n = m = 80$ and $n = m = 100$. The restricted versions of Theorem 5.1 and Corollary 5.2 fall equidistant between time taken to solve the Zames-Falb and Park criteria.

Considering both the conservatism and computational demands of each criterion, the following recommendations are proposed for practitioners dealing with high dimensional systems that can be represented in the form (5.3). It is recommended to first try the, fast but conservative, Small Gain Theorem. If this is unable to verify GAS, then to try the restricted versions of Theorem 5.1 and Corollary 5.2 and finally the full versions if these also fail. For these systems, it does not make sense to apply Park [72] and Zames-Falb [106] since they are far more conservative. In the event that the system can equivalently be represented in the form (5.6) and has $D = 0$, then both versions of Corollary 3.2 are preferable to the corresponding versions of Corollary 5.2. Both versions of Theorem 5.1 are preferable when $D \neq 0$.




FIGURE 5.3: Time taken to solve LMIs for Hopfield network's in Examples 9-12.

5.8 Conclusion

This chapter analysed the stability of Lurie systems with repeated magnitude nonlinearities. The relationship with Lurie systems having (leaky) ReLU nonlinearities was highlighted through loop transformations between the two systems; as a result, Lurie systems with (leaky) ReLU nonlinearities can be analysed using the results in this chapter. By positive homogeneity of the magnitude function, it was shown that if the Lurie system had a unique equilibrium point, then it must either be globally stable or unstable, rendering local stability analysis futile. Finally, global stability criteria, in the form of linear matrix inequalities, were established based on novel quadratic constraints which characterised the repeated magnitude. The new criteria were compared against other state of the art criteria and Theorem 5.1 was the least conservative criterion which can be applied when $D \neq 0$. Furthermore, Corollary 5.2 was competitive with Corollary 3.2, whilst being applicable to a wider range of nonlinearities.

Chapter 6

Lurie Networks with Robust Convergent Dynamics

This chapter introduces the Lurie network as a novel and unifying time invariant neural ODE. Many existing models, including recurrent neural networks and neural oscillators, are special cases of the Lurie network in this context. In order for a Lurie network to be sufficiently expressive, it may be too restrictive to require it to be stable in the traditional sense understood in control theory. Instead, this chapter seeks to impose mild constraints on the weights and biases of the Lurie network to ensure k -contraction is guaranteed. This generalised stability measure permits global convergence to a point, line or plane in the neural state-space. This includes global convergence to one of multiple equilibrium points or limit cycles, as observed in many dynamical systems including associative and working memory. Weights and biases of the Lurie network, which satisfy the k -contraction constraints, are encoded through unconstrained parametrisations. The novel stability results and parametrisations provide a toolset for training over the space of k -contracting Lurie network's using standard nonlinear optimisation algorithms. These results are also applied to construct and train a graph Lurie network satisfying the same convergence properties. Empirical results show the improvement in prediction accuracy, generalisation and robustness on a range of simulated dynamical systems, when the graph structure and k -contraction conditions are introduced. Further experiments highlight the models ability to form interpretable representations when applied to classification tasks.

6.1 Introduction

The Lurie network, proposed in this chapter, is a time-invariant model and a special case of the *Forced Lurie System* (Section 2.2.1). When modelling time-invariant dynamical systems, many machine learning (ML) models including linear state space models

(LSSM), recurrent neural networks (RNN), and some graph neural networks (GNN) are special cases of a Lurie network (Section 6.3.1). Such models have proven to be highly expressive as demonstrated by their successful application on a wide range of tasks such as sequential processing [32; 31; 30; 22], computer vision [31; 30; 22], language modelling [22] and computational chemistry [27].

Consider dynamical systems in neuroscience. The brain organises its representations of the world and carries out complex functions through collective interactions of simpler modules [123]. More abstractly, this can be viewed as a graph structured dynamical system. Convergent dynamics in widespread regions of the central nervous system are thought to play a crucial role in: forming some of these representations [124], processing information over extended periods [125], learning [29; 126], memory storage [92; 127; 128; 129] and enhancing the robustness of each of these functions [124]. As summarised in [124], convergent dynamics in the brain take several forms including neural circuits with multiple equilibrium points, such as circuits implementing associative memory [130; 131; 132], and neural circuits exhibiting limit cycles, leveraged for working memory [128]. The graph structure and convergent dynamics are also shared with many other dynamical systems such as chemical processes [133], opinion dynamics and power systems [134]. As a result, encoding such structural and dynamical properties as an inductive bias is motivated by neuroscience, for a general ML framework, and for learning robust models of dynamical systems.

The convergence and stability analysis of dynamical systems has been well-studied in the control theory literature. A pertinent example is the absolute stability problem (Section 2.3). Approaches to this problem can be classified as Lyapunov analysis [70; 72; 14], Zames-Falb multipliers [58; 61; 91; 64] or k -contraction analysis [135; 133; 134]. Lyapunov and Zames-Falb multiplier methods are primarily designed to analyse the convergence to an equilibrium point, whereas k -contraction methods (Section 6.2) analyse a variety of global convergence behaviours including convergence to points, lines and planes.

Contribution: Although designing networks with convergent dynamics is well motivated, ensuring such a property requires constraints on the network parameters which can be detrimental if too restrictive. With this in mind, Section 6.4 focuses on using k -contraction analysis to derive mild constraints on the weights of the Lurie network which ensure global convergence to a point, line or plane in the neural state space, for all Lurie networks with slope-restricted activation functions. Section 6.5 then establishes unconstrained parametrisations of these conditions which allows the Lurie network to be trained using gradient-based nonlinear optimisation algorithms, whilst limiting the search space to weights which satisfy the k -contraction conditions. Section 6.4.2 and Section 6.5.2, respectively, extend both of these results for constructing and parametrisating a graph Lurie network (GLN) from individual k -contracting Lurie




FIGURE 6.1: Trajectories from k -contracting dynamical systems. For $k \in \{1, 2\}$, time-invariant k -contracting systems exponentially converge to equilibrium points. 3-contracting systems exponentially converge to limit cycles.

networks. In Section 6.6, the difference in prediction accuracy, generalisation and robustness of the unconstrained Lurie network, k -contracting Lurie network, and GLN is empirically benchmarked on a range of simulated dynamical systems datasets. Furthermore, to illustrate the generality of the model, its ability to form representations was also investigated using the fashion MNIST dataset.

6.2 Preliminaries: k -contraction Analysis

In this chapter, k -contraction analysis [136; 48], the geometrical generalisation of contraction analysis [49], is applied to control convergence in the neural state space. Intuitively, k -contraction implies the volume of k -dimensional bodies exponentially converges to zero when governed by the system dynamics. Alternatively, this could be thought of as exponential convergence to a $(k - 1)$ -dimensional subspace. When $k = 1$, this reduces to standard contraction [49], which implies that all trajectories exponentially converge to a single trajectory. For a general time-varying dynamical system, satisfying the k -contraction property does not guarantee stability. However, for time-invariant dynamical systems, it has been shown that for every bounded solution: 1-contraction implies global convergence to a unique equilibrium point [49], 2-contraction implies global convergence to an equilibrium point, which is not necessarily unique but must be connected along a line [136], and 3-contraction, under certain assumptions, implies convergence to a non-unique attractor in a 2d subspace [137]. Three examples of k -contracting dynamics are presented in Figure 6.1.

Time-invariant dynamical systems which satisfy the k -contraction property for $k \in \{1, 2, 3\}$ have several desirable properties for ML models. They can exhibit a wide range of complex convergent behaviours such as multi-stable and orbitally stable systems [138]. This suggests that a model can be *expressive* whilst satisfying the k -contraction conditions, particularly for higher values of k where the constraints are less restrictive,

as highlighted in Section 6.4. The k -contraction property also implies an inherent *robustness* as the trajectories can only converge to a finite number of long term behaviours. Next, the fundamental k -contraction result from [48] is restated. Section 2.1.8 provides further background and a formal definition of k -contraction.

Theorem 6.1 ([48]). *Fix $k \in [1, n]_{\mathbb{N}}$ and consider the nonlinear system $\dot{x} = f(t, x)$ with $f : \mathbb{R}_{\geq 0} \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ continuously differentiable. If there exists $\eta > 0$ and an invertible matrix $\Theta \in \mathbb{R}^{n \times n}$ such that*

$$\mu_{2, \Theta^{(k)}}(J_f^{[k]}(t, x)) \leq -\eta \quad \forall x \in \mathbb{R}^n \text{ and } t \in \mathbb{R}_{\geq 0} \quad (6.1)$$

then the nonlinear system is k -contracting in the 2-norm w.r.t the metric $P := \Theta^{\top} \Theta$.

This result has two features: (i) it requires the existence of an invertible matrix P . In the simplest case, one can expect a solution $P = pI_n$ to exist. For other systems, such simple solutions will not exist and more general matrices such as $P \in \mathbb{S}_+^n$ will be required, making the proofs more difficult; (ii) it requires the use of compound matrices (Section 2.1.6). For a matrix $W \in \mathbb{R}^{n \times m}$, the matrix $W^{[k]}$ with $k \in [1, \min(n, m)]_{\mathbb{N}}$ will have the size $\binom{n}{k} \times \binom{m}{k}$ which is typically much larger and more computationally difficult to work with. A technical introduction to compound matrices, k -contraction analysis and how they relate is presented in Section 2.1.6 - Section 2.1.8. Results which verify (6.1), for the special case of nonlinear systems (6.2), are derived in Section 6.4.

6.3 Lurie Networks

A *Lurie network* is defined by (6.2) with weights $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, $C \in \mathbb{R}^{m \times n}$ and biases $b_x \in \mathbb{R}^n$, $b_y \in \mathbb{R}^m$. Its relationship to the forced Lurie system is discussed in Section 2.2.1.

$$\dot{x} = Ax + B\Phi(y) + b_x \quad y = Cx + b_y \quad x(0) = x_0 \quad (6.2)$$

The model has a biased linear component interconnected with a nonlinearity of the form $\Phi(y) := [\phi_1(y_1), \dots, \phi_m(y_m)]^{\top}$ where $\phi_i(y_i)$ is assumed to be slope-restricted with an upper bound $g > 0$, such that $0 \preceq J_{\Phi}(y) \preceq gI_m$. Some activation functions which satisfy this property are presented in Table 2.1. For simplicity, it is assumed the same scalar nonlinearity is applied element-wise and the subscript is dropped. The proposed Lurie network is a very general model as highlighted by the special cases in Section 6.3.1 and it's relationship to deep feed-forward neural networks (Section 2.2.4).

Finally, it is important to observe that the model is time-invariant, this implies that if Theorem 6.1 is satisfied, the model will inherit the appealing convergence and robustness properties stated in Section 6.2. This may be viewed as a limitation for modelling neural dynamics since the brain is subject to various external inputs; however, it is common to assume that, at least on the time scale of interest, the dynamics evolve in a time-invariant manner [124].

6.3.1 Example Lurie Networks

When applied to time-invariant dynamical systems, many models from the ML literature become special cases of (6.2). In this setting, the time-varying external inputs are replaced with trainable biases. As the results in Section 6.4 apply to any model of the form (6.2), they can also be applied to these special cases.

Lipschitz RNN: A stability constrained RNN [31] where the parameters A, C are expressed as a weighted sum of symmetric and skew-symmetric terms in order to control the eigenvalues of the Jacobian. The remaining components are $B = I_n, b_x = 0$ and $\phi(\cdot) \equiv \tanh(\cdot)$.

$$A, C \in \left\{ (1 - \beta)(W + W^\top) + \beta(W - W^\top) - \gamma I_n \mid W \in \mathbb{R}^{n \times n}, 0.5 \leq \beta \leq 1, \gamma > 0 \right\} \quad (6.3)$$

Antisymmetric RNN: A constrained RNN [30] with the same motivations as the Lipschitz RNN. Related to (6.2) by $A = 0, B = I_n, b_x = 0, \phi(\cdot) \equiv \tanh(\cdot)$ and $C \in \{W - \gamma I_n \mid W \in \text{Skew}(n), \gamma > 0\}$.

SVD Combo: A 1-contracting graph coupled RNN [32] where q denotes the number of individual RNNs and n denotes the state dimension of each RNN. A special case of (6.2) with $A = L - aI_{qn}, C = I_{qn}$ and $b_y = 0$. The graph coupling matrix is denoted by L and B is block diagonal where each block contains the synaptic weights of the individual RNNs. Both of these matrices are expressed by special parametrised forms to ensure the individual RNNs and graph coupled RNN are 1-contracting. As in (6.2), the nonlinearity is required to be slope-restricted.

Neural Oscillators: This example is from the graph ML literature [24]. The state of the general neural oscillator is governed by a second order ODE; however, its equivalent first order representation takes the form (6.2) with one possible realisation given by $C_{21} \in \mathbb{R}^{n \times n}, b_x = 0$ and

$$A = \begin{bmatrix} 0 & I \\ 0 & 0 \end{bmatrix} \quad B = \begin{bmatrix} 0 & 0 \\ 0 & I \end{bmatrix} \quad C = \begin{bmatrix} 0 & 0 \\ C_{21} & 0 \end{bmatrix} \quad b_y = \begin{bmatrix} 0 \\ b_{y2} \end{bmatrix} \quad (6.4)$$

The solution is then passed through an affine readout layer.

Graph Coupled Oscillators: Another example of a second order ODE from the graph ML literature [27]. The state is defined as a matrix, but this can simply be recast in a vectorised form which relates to (6.2) if a linear coupling function is chosen. This requires the weights to have a block matrix form, where one possible realisation is defined by $C_{21} \in \mathbb{R}^{n \times n}$, $b_x = b_y = 0$ and

$$A = \begin{bmatrix} 0 & I \\ -\gamma I & -\alpha I \end{bmatrix} \quad B = \begin{bmatrix} 0 & 0 \\ 0 & I \end{bmatrix} \quad C = \begin{bmatrix} 0 & 0 \\ C_{21} & 0 \end{bmatrix} \quad (6.5)$$

LSSM: When the external input is replaced by nonlinear output feedback (i.e., $u(t) \equiv \Phi(y)$) and $D = 0$, the linear state-space layer used in S4 [22] and Hippo [139] is a special case of (6.2) with A being a lower triangular Hippo matrix, $B \in \mathbb{R}^{n \times m}$, $C \in \mathbb{R}^{m \times n}$ and $b_x = b_y = 0$.

6.3.2 Graph Lurie Networks

Many larger scale dynamical systems such as molecular, social, biological, and financial networks [140] naturally have a graph structure. To make the Lurie network more applicable to these problems, a graph coupling term is introduced to model a set of q interacting Lurie networks. To illustrate this, q independent Lurie networks can be modelled by

$$\dot{x} = A_G x + B_G \Phi(y) + b_x \quad y = C_G x + b_y \quad x(0) = x_0 \quad (6.6)$$

where $A_G := \text{blockdiag}(A_1, \dots, A_q) \in \mathbb{R}^{qn \times qn}$, $B_G := \text{blockdiag}(B_1, \dots, B_q) \in \mathbb{R}^{qn \times qm}$, $C_G := \text{blockdiag}(C_1, \dots, C_q) \in \mathbb{R}^{qm \times qn}$ and the biases are $b_x \in \mathbb{R}^{qn}$, $b_y \in \mathbb{R}^{qm}$. The graph Lurie network (GLN) is then defined by

$$\dot{x} = (A_G + L)x + B_G \Phi(y) + b_x \quad y = C_G x + b_y \quad x(0) = x_0 \quad (6.7)$$

where $L := [L_{jl}]$ is a block matrix with block $L_{jl} \in \mathbb{R}^{n \times n}$ connecting Lurie network l to Lurie network j . The state and nonlinearity of the GLN are defined by $x \in \mathbb{R}^{qn}$ and $\Phi : \mathbb{R}^{qm} \rightarrow \mathbb{R}^{qm}$ where the states of the independent Lurie networks have been

stacked into a single state. Interestingly, the GLN (6.7) is actually a special case of a Lurie network; however, as the networks get larger, so does the search space. Thus, imposing any assumptions which respect the structure of the problem can reduce the search space and lead to more robust and generalisable models. Any prior knowledge about the graph can be encoded through constraints on the graph coupling matrix.

6.4 k -contracting Lurie Networks

6.4.1 k -contracting Lurie Networks

Two sufficient results which satisfy Theorem 6.1 and guarantee (6.2) is k -contracting are presented next. Conditions were derived in [134, Theorem 2] which verify Theorem 6.1 for a Lurie network with $A \in \mathbb{D}^n$ and $b_y = 0$. Theorem 6.2 extends them to account for $A \in \mathbb{R}^{n \times n}$ and $b_y \neq 0$.

Theorem 6.2. *Consider the Lurie network (6.2) with $\Phi(y) := [\phi_1(y_1), \dots, \phi_m(y_m)]^\top$ being slope-restricted such that $0 \preceq J_\Phi(y) \preceq gI_m$. Now define $\alpha_k := (2k)^{-1} \sum_{i=1}^k \lambda_i(A + A^\top)$ where $k \in [1, n]_{\mathbb{N}}$. If $\alpha_k < 0$ and*

$$g^2 \sum_{i=1}^k \sigma_i^2(B) \sigma_i^2(C) < \alpha_k^2 k \quad (6.8)$$

then (6.2) is k -contracting in the 2-norm with respect to the metric $P = -\alpha_k^{-1} I_n$.

Proof:

Step 1

By (6.8) there exists $\gamma < 0$ such that

$$0 < \gamma^2 < \alpha_k^2 \quad \text{and} \quad g^2 \sum_{i=1}^k \sigma_i^2(B) \sigma_i^2(C) < \gamma^2 k \quad (6.9)$$

Now, define the equivalent representation of the Lurie network (6.2)

$$\dot{x} = \bar{A}x - \bar{B}\Psi(x) \quad y = \bar{C}x \quad (6.10)$$

with

$$\bar{A} := A \quad \bar{B} := \gamma I_n \quad \bar{C} := I_n \quad \Psi(x) := -\gamma^{-1} B \Phi(Cx + b_y) - \gamma^{-1} b_x \quad (6.11)$$

Also, define $P := pI_n$ and $\Theta := p^{\frac{1}{2}}I_n$, with $p > 0$.

Inserting P, Θ , and (6.10) into the left hand side of (A.1) yields

$$\begin{aligned} & P^{(k)} \bar{A}^{[k]} + (\bar{A}^{[k]})^\top P^{(k)} + \Theta^{(k)} \left((\Theta \bar{B} \bar{B}^\top \Theta)^{[k]} + (\Theta^{-1} \bar{C}^\top \bar{C} \Theta^{-1})^{[k]} \right) \Theta^{(k)} \\ &= P^{(k)} A^{[k]} + (A^{[k]})^\top P^{(k)} + \Theta^{(k)} \left((\gamma^2 P)^{[k]} + (P^{-1})^{[k]} \right) \Theta^{(k)} \\ &= (pI_n)^{(k)} A^{[k]} + (A^{[k]})^\top (pI_n)^{(k)} + (p^{\frac{1}{2}}I_n)^{(k)} \left((\gamma^2 p I_n)^{[k]} + (p^{-1} I_n)^{[k]} \right) (p^{\frac{1}{2}}I_n)^{(k)} \end{aligned} \quad (6.12)$$

Applying the relevant special cases from (2.41) and (2.48) leads to

$$= p^k (A^{[k]} + (A^{[k]})^\top) + k(\gamma^2 p + p^{-1}) p^k I$$

Next, applying Fact 2.6 and Fact 2.8 results in

$$= p^k ((A + A^\top)^{[k]} + k(\gamma^2 p + p^{-1}) I)$$

Finally, re-applying (2.48) and Fact 2.8 yields

$$= p^k (A + A^\top + (\gamma^2 p + p^{-1}) I_n)^{[k]}$$

If the matrix above is negative definite, then (A.1) is satisfied for some suitably chosen $\eta_1 > 0$. This is true when

$$(A + A^\top + (\gamma^2 p + p^{-1}) I_n)^{[k]} \prec 0 \quad (6.13)$$

By Fact 2.7 and [141, Eq. 285], the inequality above can equivalently be expressed as

$$k(\gamma^2 p + p^{-1}) + \sum_{i=1}^k \lambda_i(A + A^\top) < 0 \quad (6.14)$$

Inserting α_k , this simplifies to

$$\gamma^2 p^2 + 2\alpha_k p + 1 < 0 \quad (6.15)$$

For γ satisfying (6.9), a solution $p = -\alpha_k^{-1}$ always exists.

Step 2

By (6.10)

$$J_{\Psi}(x) = -\gamma^{-1}BJ_{\Phi}(Cx + b_y)C$$

Inserting Θ , J_{Ψ} , and (6.10) into the left hand side of (A.2), and applying Fact 2.7, yields

$$\begin{aligned} \sum_{i=1}^k \lambda_i \left(\Theta^{-1} \bar{C}^{\top} (J_{\Psi}^{\top}(y) J_{\Psi}(y) - I) \bar{C} \Theta^{-1} \right) &= \sum_{i=1}^k \lambda_i \left(p^{-1} J_{\Psi}^{\top}(y) J_{\Psi}(y) - p^{-1} I \right) \\ &= (p^{-1} J_{\Psi}^{\top}(y) J_{\Psi}(y) - p^{-1} I)^{[k]} \end{aligned}$$

If $(p^{-1} J_{\Psi}^{\top}(y) J_{\Psi}(y) - p^{-1} I)^{[k]} \prec 0$, then (A.2) is satisfied for some suitably chosen $\eta_2 > 0$. Applying Fact 2.7, this negative definite requirement reduces to

$$\sum_{i=1}^k \lambda_i \left(p^{-1} J_{\Psi}^{\top}(y) J_{\Psi}(y) \right) - kp^{-1} = p^{-1} \sum_{i=1}^k \sigma_i^2(J_{\Psi}(y)) - kp^{-1} < 0 \quad (6.16)$$

Finally, inserting J_{Ψ} and, at the expense of some conservatism, applying the well-known property of singular values [142, Theorem 3.3.14], the inequality above becomes

$$\gamma^{-2} \sum_{i=1}^k \sigma_i^2(B) \sigma_i^2(J_{\Phi}(y)) \sigma_i^2(C) < k \quad (6.17)$$

By the assumption made on the slope of Φ , this inequality will always be satisfied if (6.8) holds. \square

Remark 6.1. *The additional freedom permitted by k -contraction over standard contraction is highlighted by the summation of the eigenvalues and singular values in Theorem 6.2. In 1-contraction, Theorem 6.2 requires the largest eigenvalue of the symmetric component of A to be negative whereas for $k \in [2, n]_{\mathbb{N}}$, this condition on A becomes incrementally more relaxed as k is increased. Equation (6.8) illustrates a similar relaxation of the constraints on B and C . Theorem 6.2 has several appealing features: (i) it does not require the computation of the troublesome compound matrices; (ii) it provides a way of embedding the k -contraction property into the structure of a Lurie network based on fairly simple unconstrained parametrisations of the weights, as shown in Section 6.5; (iii) the biases are not present in the condition, so are naturally unconstrained; (iv) the result does not rely on symmetries of the parameters. In many Hopfield-based models of associative memory, symmetry in the parameters is needed to make*

the existence of a global energy function mathematically tractable; however, this simplification is biologically unrealistic and limits the model's expressive power. The limitation of the result is that only Lurie networks which are k -contracting in a scalar metric can be verified. $\square\square$

A second result which addresses the scalar metric drawback is presented next; however, it comes at the cost of strong constraints on the weights B and C .

Theorem 6.3. Consider the Lurie network (6.2) with $\Phi(y) := [\phi_1(y_1), \dots, \phi_n(y_n)]^\top$ being slope-restricted such that $0 \preceq J_\Phi(y) \preceq gI_n$. Fix $k \in [1, n]_{\mathbb{N}}$. If $B \in \mathbb{D}^n$, $C = B^{-1}$ and

$$P^{(k)} A^{[k]} + (A^{[k]})^\top P^{(k)} + 2kgP^{(k)} \prec 0 \quad (6.18)$$

then (6.2) is k -contracting in the 2-norm w.r.t the metric $P \in \mathbb{D}_+^n$.

Proof: This proof aims to directly verify Theorem 6.1 for $\Theta \in \mathbb{D}^n$, where f represents the Lurie network (6.2). The proof begins by substituting the Jacobian of the Lurie network into the left hand side of (6.1), followed by the application of Fact 2.8 to obtain the second equality. The subadditivity property of the matrix measure μ_2 was then leveraged to split the terms [35, Section 2.2]. As the second term is difficult to manipulate, the simplifying assumption $C = B^{-1}$ was made in order to apply Fact 2.9. As $\Theta, B, J_\Phi \in \mathbb{D}^n$, so are both of their k -compound counterparts (2.41), (2.48), which means the terms Θ, B cancel out. The property $\mu_2(\cdot) \leq \|\cdot\|_2$ [35, Theorem 16] was then applied followed by Fact 2.7, which allowed the slope restricted assumption on Φ to be utilised. By the relevant special case of the k -additive compound (2.48), the 2-norm was calculated. Finally, kg was incorporated in μ_2 as shown in the final inequality.

$$\begin{aligned} \mu_{2, \Theta^{(k)}}(J_f^{[k]}(t, x)) &= \mu_{2, \Theta^{(k)}}((A + BJ_\Phi C)^{[k]}) \\ &= \mu_{2, \Theta^{(k)}}(A^{[k]} + (BJ_\Phi(y)C)^{[k]}) \\ &\leq \mu_{2, \Theta^{(k)}}(A^{[k]}) + \mu_{2, \Theta^{(k)}}((BJ_\Phi C)^{[k]}) \\ &= \mu_{2, \Theta^{(k)}}(A^{[k]}) + \mu_2(\Theta^{(k)} B^{(k)} J_\Phi^{[k]} B^{-(k)} \Theta^{-(k)}) \\ &= \mu_{2, \Theta^{(k)}}(A^{[k]}) + \mu_2(J_\Phi^{[k]}) \\ &\leq \mu_{2, \Theta^{(k)}}(A^{[k]}) + \|(gI_n)^{[k]}\|_2 \\ &= \mu_{2, \Theta^{(k)}}(A^{[k]}) + kg \\ &= \mu_2(\Theta^{(k)} A^{[k]} \Theta^{-(k)} + kgI_n) \end{aligned} \quad (6.19)$$

If the final inequality is negative, then Theorem 6.1 is satisfied for some suitably chosen η . This is equivalent to the matrix inequality below.

$$\frac{1}{2}(\Theta^{(k)} A^{[k]} \Theta^{-(k)} + \Theta^{-(k)} (A^{[k]})^\top \Theta^{(k)} + 2kgI_n) \prec 0 \quad (6.20)$$

Multiplying on the left by $2\Theta^{(k)}$ and on the right by $\Theta^{(k)}$ results in (6.18). \square

Remark 6.2. *Theorem 6.3 improves upon Theorem 6.2 in the sense that Lurie networks which are k -contracting in a diagonal metric can now be verified; however, only when $C = B^{-1}$. Due to this constraint, Theorem 6.3 has very limited practical use; however, it may prove to be theoretically insightful for addressing the scalar metric drawback in the future. Finally, it is important to highlight that Theorem 6.2 and Theorem 6.3 apply to the class of slope-restricted nonlinearities, so these results address the absolute stability problem with respect to the k -contraction property. $\square\square$*

6.4.2 k -contracting Graph Lurie Networks

In this section, q independent Lurie networks (6.2) are assumed to be k -contracting in the 2-norm with respect to the metrics P_1, \dots, P_q . This is equivalent to (6.6) k -contracting in the 2-norm with respect to the metric $P = \text{blockdiag}(P_1, \dots, P_q)$. Theorem 6.2 and Theorem 6.3 provide two results which can, respectively, verify this with respect to a scalar metric $P_j = p_j I_n$ and a diagonal metric $P_j \in \mathbb{D}_+^n$ for $j \in [1, q]_{\mathbb{N}}$. Other results may be used providing they apply to systems of the form (6.2). Under this assumption, Theorem 6.4 provides a constraint on the graph coupling term which ensures the GLN is k -contracting when constructed from q independently k -contracting Lurie networks.

Theorem 6.4. *Fix $k \in [1, n]_{\mathbb{N}}$ and consider the GLN (6.7) where the q independent Lurie networks are k -contracting in the 2-norm with respect to the metrics P_j with $j \in [1, q]_{\mathbb{N}}$. If the graph coupling matrix $L \in \mathbb{R}^{qn \times qn}$ satisfies*

$$P^{(k)} L^{[k]} + (L^{[k]})^\top P^{(k)} \preceq 0 \quad (6.21)$$

where $P := \text{blockdiag}(P_1, \dots, P_q)$, then (6.7) is k -contracting in the 2-norm with respect to the metric P .

Proof: This proof verifies Theorem 6.1 when f represents the GLN (6.7). The proof starts by expressing the Jacobian as a sum of the Jacobian of the q independent Lurie networks, J_{indep} , and Jacobian of the coupling term, J_{couple}

$$J_f(x) = J_{\text{indep}}(x) + J_{\text{couple}}(x) \quad (6.22)$$

where $J_{\text{indep}}(x) = A_G + B_G J_\Phi C_G$ and $J_{\text{couple}}(x) = L$. Substituting J_f into the left hand side of (6.1) and applying the sub-additivity property of μ_2 results in

$$\mu_{2,\Theta^{(k)}}(J_f^{[k]}) \leq \mu_{2,\Theta^{(k)}}(J_{\text{indep}}^{[k]}) + \mu_{2,\Theta^{(k)}}(J_{\text{couple}}^{[k]}) \quad (6.23)$$

As the q independent Lurie networks are assumed to be k -contracting in the 2-norm with respect to the metric P , then $\mu_{2,\Theta^{(k)}}(J_{\text{indep}}^{[k]}) < 0$. Under this assumption, Theorem 6.1 is satisfied if

$$\mu_{2,\Theta^{(k)}}(J_{\text{couple}}^{[k]}) \leq 0 \quad (6.24)$$

This is equivalent to the matrix inequality below, where J_{couple} has been substituted in.

$$\frac{1}{2}(\Theta^{(k)}L^{[k]}\Theta^{-(k)} + \Theta^{-(k)}(L^{[k]})^\top\Theta^{(k)}) \preceq 0 \quad (6.25)$$

Multiplying on the left by $2\Theta^{(k)}$ and on the right by $\Theta^{(k)}$ results in (6.21). \square

Remark 6.3. *It should be noted that Theorem 6.4 could be applied more generally for constructing k -contracting graph coupled systems. Providing the q subsystems are k -contracting in the 2-norm with respect to some metric $P = \text{blockdiag}(P_1, \dots, P_q)$, then Theorem 6.4 is applicable when coupling them through the graph coupling term. It is not necessary for the q subsystems to be Lurie networks. $\square\square$*

6.5 Parametrisation of k -contracting Lurie Networks

To train a k -contracting Lurie network using gradient based optimisers, parametrisations which express the constrained weights in terms of unconstrained variables must be found. To formalise this idea, the sets $\Omega_2(g, k)$ and $\Omega_4(k, P)$ are defined. As the biases do not appear in these sets, they are naturally unconstrained.

The set $\Omega_2(g, k)$ contains the Lurie network's which satisfy Theorem 6.2.

$$\Omega_2(g, k) := \left\{ (\bar{A}, \bar{B}, \bar{C}) \mid \alpha_k := \frac{1}{2k} \sum_{i=1}^k \lambda_i(\bar{A} + \bar{A}^\top) < 0, z_k := g^2 \sum_{i=1}^k \sigma_i^2(\bar{B})\sigma_i^2(\bar{C}) < \alpha_k^2 k \right\} \quad (6.26)$$

The set $\Omega_4(k, P)$ contains the graph coupling matrices which satisfy Theorem 6.4.

$$\Omega_4(k, P) := \left\{ \bar{L} \in \mathbb{R}^{qn \times qn} \mid P^{(k)}\bar{L}^{[k]} + (\bar{L}^{[k]})^\top P^{(k)} \preceq 0 \right\} \quad (6.27)$$

A parametrisation associated with Theorem 6.3 is not established due to the limitations mentioned earlier. The next section present two different parametrisations of the set

$\Omega_2(g, k)$, both proofs apply the eigenvalue and singular value decompositions, detailed in Section 2.1.5.

6.5.1 Parametrisation of k -contracting Lurie Networks

Theorem 6.5. *Let $g > 0$ and $k \in [1, n]_{\mathbb{N}}$. Now, let $U_A, U_B, V_C \in \mathcal{O}(n)$, $V_B, U_C \in \mathcal{O}(m)$, $\Sigma_B \in \mathbb{D}_+^{mm}$, $\Sigma_C \in \mathbb{D}_+^{mn}$, $Y_A \in \text{Skew}(n)$, and $G_A \in \mathbb{D}_+^n$. Define*

$$A := \frac{1}{2}U_A\Sigma_A U_A^\top + \frac{1}{2}Y_A \quad \Sigma_A := -\sqrt{\frac{4z_k}{k}}I_n - G_A \quad (6.28a)$$

$$B := U_B\Sigma_B V_B^\top \quad C := U_C\Sigma_C V_C^\top \quad (6.28b)$$

then $(A, B, C) \in \Omega_2(g, k)$.

Proof: The singular value decomposition was leveraged to expose the singular values of B, C , as in (6.28b). This required the matrices U_B, U_C, V_B, V_C to be orthogonal. One can immediately use the unconstrained parametrisation of the orthogonal class from [143] to express these matrices as unconstrained symmetric matrices. The matrices $\Sigma_B \in \mathbb{D}_+^{mm}$, $\Sigma_C \in \mathbb{D}_+^{mn}$ respectively contain the singular values of B and C . These positive diagonal matrices are also treated as unconstrained sets since any matrix can be obtained by taking the absolute value of an unconstrained diagonal matrix with the same shape.

To verify Theorem 6.2, one can combine $\alpha_k < 0$ and (6.8) into a single inequality representing the intersection of the two sets.

$$\sum_{i=1}^k \lambda_i(A + A^\top) < -2k\sqrt{\frac{z_k}{k}} \quad \text{where } z_k := g^2 \sum_{i=1}^k \sigma_i^2(B)\sigma_i^2(C) \quad (6.29)$$

By definition of B and C , the right hand side is a function of the hyperparameters g, k and elements of the parameters Σ_B, Σ_C , so can be easily computed using sort and sum functions.

To impose this constraint directly on the eigenvalues of the symmetric component of A , the A matrix is expressed as a sum of symmetric and skew-symmetric matrices. The skew-symmetric matrix, Y_A , is unconstrained, so no further work is needed. The symmetric matrix is then expressed by its eigenvalue decomposition to obtain (6.28a). This expression allowed the constraint (6.29) to be placed directly on the elements of the diagonal matrix, Σ_A .

Defining Σ_A as in (6.28a) ensures

$$\lambda_i(A + A^\top) < -2\sqrt{\frac{z_k}{k}} \quad \text{for all } i \in [1, n]_{\mathbb{N}} \quad (6.30)$$

which guarantees both conditions of Theorem 6.2 are satisfied. \square

Conservatism is introduced in Theorem 6.5 as all the diagonal elements of Σ_A must be negative. Theorem 6.6 is established to address this issue.

Theorem 6.6. *Let $g > 0$ and $k \in [1, n]_{\mathbb{N}}$. Now, let $U_A, U_B, V_C \in \mathcal{O}(n)$, $V_B, U_C \in \mathcal{O}(m)$, $\Sigma_B \in \mathbb{D}_+^{mm}$, $\Sigma_C \in \mathbb{D}_+^{mm}$, $Y_A \in \text{Skew}(n)$, $\Sigma_{A1} \in \mathbb{D}^{k-1}$, $G_{A2} > 0$, and $G_{A3} \in \mathbb{D}_+^{n-k}$. Define*

$$A := \frac{1}{2}U_A\Sigma_A U_A^\top + \frac{1}{2}Y_A \quad B := U_B\Sigma_B V_B^\top \quad C := U_C\Sigma_C V_C^\top \quad (6.31a)$$

$$\Sigma_A := \text{blockdiag}(\Sigma_{A1}, \Sigma_{A2}, \Sigma_{A3}) \quad \Sigma_{A1} \in \mathbb{D}^{k-1} \quad (6.31b)$$

$$\Sigma_{A2} := -\sqrt{4kz_k} - \sum_i^{k-1} (\Sigma_{A1})_{ii} - G_{A2} \quad \Sigma_{A3} := \min(\Sigma_{A1}, \Sigma_{A2})I_{n-k} - G_{A3} \quad (6.31c)$$

then $(A, B, C) \in \Omega_2(g, k)$.

Proof: The proof follows the same steps as Theorem 6.6, up to the definition of Σ_A . For this reason, the proof starts from this point.

Defining Σ_A , as in Theorem 6.6, guarantees both conditions of Theorem 6.2 are satisfied via (6.29). In this case, the definition of Σ_A is split into one unconstrained block for the first $k - 1$ eigenvalues (Σ_{A1}), a block matrix for $\lambda_k(A + A^\top)$ which ensures (6.29) holds (Σ_{A2}), and finally a block for the remaining eigenvalues which must be defined to ensure the k eigenvalues involved in (6.29) are the largest (Σ_{A3}). \square

Remark 6.4. *In both results, a mapping between the sets $\text{Skew}(\cdot)$ and $\mathcal{O}(\cdot)$ was exploited to express the orthogonal matrices in terms of unconstrained skew-symmetric matrices [143]. The remaining variables are unconstrained or simply require positive elements, which can be obtained by taking the absolute value of unconstrained variables. $\square\square$*

In Theorem 6.5 and Theorem 6.6, the B and C matrices are unconstrained since they are simply expressed by their singular value decomposition. The variables representing the singular values of B and C directly upper bound α_k . The only source of conservatism in the parametrisations is introduced through the definition of Σ_A . In Theorem 6.5, the constraint on α_k is enforced by a uniform negative constraint on the diagonal elements of Σ_A . If this assumption is true, this significantly speeds up the learning process; however, it can be prohibitive if not. Theorem 6.6 allows the largest $(k - 1)$ eigenvalues to be unconstrained (Σ_{A1}), meaning non-Hurwitz A matrices are encapsulated in the parametrisation. The constraint on α_k is implemented by the Σ_{A2} block and the Σ_{A3} block ensures the remaining eigenvalues are less than the other k .

As k is a hyperparameter of the model, it is useful to note the following result which allows the model to be applied in situations where k is unknown.

Proposition 6.1. For $k \in \{2, 3\}$, the sets $\Omega_2(g, k-1)$ and $\Omega_2(g, k)$, both defined by (6.26), intersect when the following inequality is satisfied

$$g^2 \sigma_k^2(B) \sigma_k^2(C) < \frac{1}{4k} \lambda_k^2 (A + A^\top) + \frac{k-1}{k} \alpha_{k-1} \lambda_k (A + A^\top) - \frac{k-1}{k} \alpha_{k-1}^2 \quad (6.32)$$

Proof: First note the following useful relationships, where $\lambda_k := \lambda_k(A + A^\top)$.

$$\alpha_k = \frac{1}{2k} \lambda_k + \frac{k-1}{k} \alpha_{k-1} \quad z_k = g^2 \sigma_k^2(B) \sigma_k^2(C) + z_{k-1} \quad (6.33)$$

From (6.33) it is clear that $\alpha_k < 0$ when

$$\lambda_k < -2(k-1) \alpha_{k-1} \quad (6.34)$$

A subset of these solutions occurs when the first condition of $\Omega_2(g, k-1)$ holds; that is $\alpha_{k-1} < 0$. From (6.33) it is also clear that $z_k < \alpha_k^2 k$ is true when

$$z_{k-1} < \alpha_k^2 k - g^2 \sigma_k^2(B) \sigma_k^2(C) \quad (6.35)$$

If $\alpha_{k-1}^2 (k-1) < \alpha_k^2 k - g^2 \sigma_k^2(B) \sigma_k^2(C)$ then a subset of these solutions occurs when the second condition of $\Omega_2(g, k-1)$ holds, i.e., $z_{k-1} < \alpha_{k-1}^2 (k-1)$. After rearranging, this is equivalent to

$$g^2 \sigma_k^2(B) \sigma_k^2(C) < \alpha_k^2 k - \alpha_{k-1}^2 (k-1) \quad (6.36)$$

The right hand side can equivalently be expressed by a quadratic equation in λ_k

$$\begin{aligned} \alpha_k^2 k - \alpha_{k-1}^2 (k-1) &= \left(\frac{1}{2k} \lambda_k + \frac{k-1}{k} \alpha_{k-1} \right)^2 k - \alpha_{k-1}^2 (k-1) \\ &= \left(\frac{1}{4k^2} \lambda_k^2 + \frac{k-1}{k^2} \alpha_{k-1} \lambda_k + \frac{(k-1)^2}{k^2} \alpha_{k-1}^2 \right) k - \alpha_{k-1}^2 (k-1) \\ &= \frac{1}{4k} \lambda_k^2 + \frac{k-1}{k} \alpha_{k-1} \lambda_k + \frac{(k-1)^2}{k} \alpha_{k-1}^2 - \alpha_{k-1}^2 (k-1) \\ &= \frac{1}{4k} \lambda_k^2 + \frac{k-1}{k} \alpha_{k-1} \lambda_k - \frac{k-1}{k} \alpha_{k-1}^2 \\ &:= Q_k(\lambda_k, \alpha_{k-1}) \end{aligned} \quad (6.37)$$

Therefore, $\Omega_2(g, k-1)$ intersects with $\Omega_2(g, k)$ if (6.32) holds.

To confirm practical solutions exist to (6.32), one must check if negative values of λ_k (a requirement of Theorem 6.5 and Theorem 6.6) ensure $Q_k(\lambda_k; \alpha_{k-1}) > 0$, when $\alpha_{k-1} < 0$. Finding the roots of $Q_k(\lambda_k, \alpha_{k-1})$ confirms what intervals of λ_k ensure $Q_k(\lambda_k, \alpha_{k-1}) > 0$. By the quadratic formula, the general solution is given by

$$\begin{aligned}\lambda_k^r &= 2(1-k)\alpha_{k-1} \pm 2k\sqrt{\frac{(k-1)^2}{k^2}\alpha_{k-1}^2 + \frac{k-1}{k^2}\alpha_{k-1}^2} \\ &= 2(1-k)\alpha_{k-1} \pm 2\sqrt{k(k-1)}\alpha_{k-1} \\ &= 2\left(1-k \pm \sqrt{k(k-1)}\right)\alpha_{k-1}\end{aligned}\tag{6.38}$$

The roots specifically for $k = 2$ and $k = 3$ are

$$\lambda_2^r = (-2 \pm 2\sqrt{2})\alpha_1 \qquad \lambda_3^r = (-4 \pm 2\sqrt{6})\alpha_2\tag{6.39}$$

Given $\alpha_{k-1} < 0$, the intervals $\lambda_2 \in [-\infty, (-2 + 2\sqrt{2})\alpha_1]$ and $\lambda_3 \in [-\infty, (-4 + 2\sqrt{6})\alpha_2]$ were tested to see if the corresponding co-domains of $Q_k(\lambda_k, \alpha_{k-1})$ were positive. These intervals ensure $\lambda_k < 0$ (a requirement for $\alpha_k < 0$ and $\alpha_{k-1} < 0$).

Using $\lambda_2 = \alpha_1$ and $\lambda_3 = \alpha_2$ as the test points, the quadratic inequalities become

$$\begin{aligned}Q_2(\alpha_1; \alpha_1) &= \frac{1}{8}\alpha_1^2 + \frac{1}{2}\alpha_1^2 - \frac{1}{2}\alpha_1^2 \\ &= \frac{1}{8}\alpha_1^2 > 0\end{aligned}\tag{6.40}$$

$$\begin{aligned}Q_3(\alpha_2; \alpha_2) &= \frac{1}{12}\alpha_2^2 + \frac{2}{3}\alpha_2^2 - \frac{2}{3}\alpha_2^2 \\ &= \frac{1}{12}\alpha_2^2 > 0\end{aligned}\tag{6.41}$$

As a result, for $k \in \{2, 3\}$, the sets $\Omega_2(g, k-1)$ and $\Omega_2(g, k)$ intersect if $\lambda_k \in [-\infty, 2(1-k + \sqrt{k(k-1)})\alpha_{k-1}]$ and $g^2\sigma_k^2(B)\sigma_k^2(C) < Q_k(\lambda_k, \alpha_{k-1})$. If $(A, B, C) \in \Omega_2(g, 2)$ then λ_2 must be in this interval since $\lambda_2 \leq \lambda_1$. By the same logic, if $(A, B, C) \in \Omega_2(g, 3)$ then λ_3 must be in this interval since $\lambda_3 \leq \lambda_2 \leq \lambda_1$.

Hence, the requirements for intersection can be simplified to just (6.32). \square

Remark 6.5. Proposition 6.1 is expected to hold for many realistic systems since (6.32) just requires $\lambda_k(A + A^\top)$ to be sufficiently negative to counteract the product of the growth of B and C in the k^{th} direction. Hence, if the best value of k is unknown for a given application, then setting $k = 3$ allows the model to search over many of the stable l -contracting systems parametrised by Theorem 6.5 or Theorem 6.6, where $l \in [1, k]_{\mathbb{N}}$. $\square\square$

6.5.2 Parametrisation of k -contracting Graph Lurie Networks

The next result provides an unconstrained parametrisation of the set $\Omega_4(k, P)$ in (6.27). Coupling this with the earlier parametrisations allows the k -contracting GLN to be trained with gradient based optimisers.

Theorem 6.7. Given $k \in [1, n]_{\mathbb{N}}$, $G_{L1}, G_{L2} \in \mathbb{R}^{qn \times qn}$, $\Theta = \text{blockdiag}(\Theta_1, \dots, \Theta_q)$ where $\Theta_j \in \mathbb{S}_+^n$ for $j \in [1, q]$, $P = \Theta^\top \Theta$ and define

$$L := G_{L1} - P^{-1}G_{L1}^\top P + \Theta^{-1}G_{L2}\Theta \quad (6.42)$$

where $(G_{L2} + G_{L2}^\top)^{[k]} \preceq 0$, then $L \in \Omega_4(k, P)$.

Proof: This proof aims to show that Theorem 6.4 is satisfied when L is defined as in (6.42). First, multiply (6.21) on the left and right by $\Theta^{-(k)}$. Sequentially applying Fact 2.6, Fact 2.9 and Fact 2.8 results in

$$(\Theta L \Theta^{-1} + \Theta^{-1} L^\top \Theta)^{[k]} \preceq 0 \quad (6.43)$$

Subbing in the definition of L from (6.42) and recalling that $P = \Theta^\top \Theta$ leads to

$$(G_{L2} + G_{L2}^\top)^{[k]} \preceq 0 \quad (6.44)$$

which is assumed to hold. \square

Theorem 6.7 provides flexibility in the choice of G_{L1}, G_{L2} to define different graph structures, where the constraint on G_{L2} is equally a constraint on the sum of its k -largest eigenvalues (Fact 2.7). This constraint could be parametrised in a similar manner to the symmetric component of A in Theorem 6.5 and Theorem 6.6; in which case, Theorem 6.7 is satisfied and the graph has a directed all-to-all structure, including self-loops.

Remark 6.6. If $G_{L2} = 0$, along with the main diagonal blocks of G_{L1} , then (6.42) is a boundary condition of (6.21). Thus, Theorem 6.7 is satisfied and the self-loops (a node coupled to itself via a graph connection) are removed from the graph structure. $\square\square$

6.5.3 Complexity Comparison

One consideration when constructing k -contracting Lurie networks is the number of additional parameters required for the parametrisation. A Lurie network, as defined in (6.2), has a total parameter count $N_L = n^2 + 2nm + n + m$; whereas a k -contracting Lurie network, constructed according to Theorem 6.5 or Theorem 6.6, has a total parameter count $N_K \leq 2n^2 + m^2$. The number of parameters only increases significantly when m is large compared to n ; however, throughout the literature, many special cases of a Lurie network set $n = m$ (Section 6.3.1); in which case, N_K becomes marginally smaller than N_L . Furthermore, the all-to-all graph coupling term has $N_C = 2(qn)^2$ parameters or $N_C = qn^2(q - 1)$ if the self-loops are removed according to Remark 6.6. This analysis highlights that ensuring the Lurie network and GLN are k -contracting comes at a minimal computational expense.

6.6 Empirical Evaluation: Dynamical Systems

In Section 6.1, it was highlighted that a range of convergent dynamics are prevalent in many dynamical systems and neural processes. In this section, the impact of the additional expressivity of the Lurie network and the k -contracting parametrization were tested on a range of dynamical systems. The prediction accuracy, generalisation and robustness of the proposed models were used to compare performance.

6.6.1 Data

Three time-invariant dynamical systems were considered: (i) an opinion dynamics model of a social network where all opinions agree and thus converge to a unique equilibrium point; (ii) a Hopfield network of associative memory with two stable equilibrium points and one unstable; (iii) a generic simple attractor which could be used to model the stored patterns in working memory. The nonlinearity used to simulate the opinion dynamics and Hopfield network satisfied the slope-restriction assumption of the Lurie network. On the other hand, the nonlinearity used to simulate the simple attractor did not satisfy this property; thus, this dataset was useful for testing the approximation capabilities of the Lurie network.

Each system had $n = 3$ states allowing for visualisation of the ground truth and predictions. For each dynamical system, a dataset including 1k trajectories, sampled every 0.01s over a 20s interval existed. The test sets were formed by holding out 100 trajectories. The input to each model was the initial condition sampled from a uniform distribution with the domain $(-1, +1)^3$ for the opinion/Hopfield datasets and $(-3, +3)^3$ for the simple attractor. The full trajectory was then used as the target to

train the model. An illustration of these datasets can be seen in Figure 6.1. The data was synthetically generated by numerically integrating over the analytical models of each dynamical system. The integration was performed using the Euler method with step size $\delta = 1 \times 10^{-2}$.

To test the out of distribution (OOD) generalisation and robustness, two additional datasets were generated for each of the opinion, Hopfield and attractor tasks. These differ from the training datasets in the following ways: (i) they include 100 trajectories over a 30s interval; (ii) the initial conditions were sampled from a uniform distribution over the intervals $1 < |x_i(0)| < 4$ for the opinion/Hopfield datasets and $3 < |x_i(0)| < 6$ for the simple attractor, where $i \in \{1, 2, 3\}$. These trajectories were 10s longer than the training data with initial conditions also sampled outside the training distribution. To test the robustness, noise sampled from the standard normal distribution was added to the initial conditions of these datasets, before generating the trajectories.

For the second set of experiments, two 30-dimensional dynamical systems were studied. The first was a graph-coupled (GC) Hopfield network, formed by connecting 10 previously described Hopfield networks through a graph coupling matrix. To ensure the convergence property was preserved, the matrix was expressed by (6.42) with G_{L1} sampled from a uniform distribution and $G_{L2} = 0$. The second system was a GC attractor, constructed in the same way. The datasets were generated as described in the previous paragraphs; however, they each included 30k trajectories.

6.6.1.1 Opinion Dynamics

This model was presented in [134]. It has the following state space equations

$$\dot{x} = -1.5I_3 + 0.5\Phi(Cx) + b \quad (6.45)$$

where

$$C = \begin{bmatrix} +1 & -1 & 0 \\ -1 & +1 & -1 \\ 0 & -1 & +1 \end{bmatrix} \quad b = \begin{bmatrix} +0.2 \\ \pm 0.0 \\ -0.2 \end{bmatrix} \quad (6.46)$$

and the tanh function is applied element-wise. The model is 1-contracting and has a unique equilibrium point at b .

6.6.1.2 Hopfield Network

This model is a variation of the Hopfield network presented in [134]. It has the following state space equations

$$\dot{x} = -2.5I_3 + B\Phi(x) \quad (6.47)$$

where

$$B = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix} \quad (6.48)$$

and the tanh function is also applied element-wise. The model is 2-contracting and has two stable equilibrium points: $e_1 = [0.79, 0.79, 0.79]^\top$, $e_2 = -e_1$ and an unstable equilibrium point $e_3 = 0$.

6.6.1.3 Simple Attractor

This model was presented in [137]. It has the following state space equations

$$\dot{x} = Ax + B\Phi(Cx) \quad (6.49)$$

where

$$A = \begin{bmatrix} 0 & 1 & -2 \\ -1 & 0 & -1 \\ 0.5 & 0 & -0.5 \end{bmatrix} \quad B = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ -0.5 & 0 & 0 \end{bmatrix} \quad C = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix} \quad (6.50)$$

and $\phi(z) = z^3$ is the nonlinearity applied element-wise. This function is not slope-restricted, so the simple attractor does not satisfy the assumptions of the Lurie network. The model is 3-contracting and has several attractor states.

6.6.2 Training

The default training settings common to the isolated (opinion, Hopfield and attractor) and GC datasets are presented in Table 6.1. The Stochastic Gradient Descent (SGD) algorithm, detailed in Section 2.1.4, was used to optimise the model parameters. The only parameters which varied between models were the learning rate (LR) and epoch which it was reduced. Deviations from the default settings are detailed in Table 6.2.

TABLE 6.1: Default training settings for the isolated and GC datasets.

Parameter	Isolated	GC
Batches	10	15
Batch Size	100	2000
Test Split	0.1	$\frac{1}{15}$
Epochs	100	100
Criterion	MSE	MSE
Optimiser	SGD	SGD
LR	1×10^{-2}	5×10^{-3}

These values were chosen based on observations during training; no hyperparameter sweep was performed.

When training the models on the isolated datasets, a single T4 GPU (accessed through Google Colab) was used. A single A100 GPU was used for training the models on the GC datasets. All code was implemented in PyTorch and can be found at <https://github.com/CR-Richardson/LurieNetwork>.

6.6.3 k -contracting Lurie Networks

In this section, the k -contracting Lurie network is compared against five other continuous time models: (i) the unconstrained Lurie network, for testing the importance of the k -contraction constraints; (ii) an unconstrained neural ODE [56] with two hidden layers, each comprised of 20 neurons and ReLU activations; (iii) the three constrained continuous-time RNNs detailed in Section 6.3.1, where the SVD Combo network was comprised of a single node. The numerical integration was performed using the Euler method with step size $\delta = 1 \times 10^{-2}$. Besides the neural ODE, each model used tanh activations. Both activation function choices were slope-restricted with $g = 1$.

For the opinion and Hopfield datasets, the k -contracting Lurie network was constructed according to Theorem 6.5 whereas Theorem 6.6 was used for the simple attractor. The hyperparameter k was set to $k = 1$ for the opinion dynamics, $k = 2$ for the multi-stable Hopfield network and $k = 3$ for the simple attractor. Similar results were obtained when setting $k = 3$ for all examples.

Table 6.3 compares the average (over 3 runs) and best mean squared error (MSE) on the test set of each isolated task. The k -contracting Lurie network achieved the best MSE on two out of three examples. The importance of the k -contraction conditions is particularly clear when comparing the mean and standard deviation with that of the unconstrained Lurie network. These conditions clearly reduce the search space to a tractable region to optimise over as the MSE of the unconstrained Lurie network is at least an order of magnitude worse than its k -contracting counterpart. The other models perform as one would expect: (i) the neural ODE demonstrates strong accuracy

TABLE 6.2: Deviations from the default settings for the isolated and GC datasets.

Dataset	Model	LR	LR Cut @ Epoch
Opinion Dynamics	Lurie Network	5×10^{-3}	-
Opinion Dynamics	Antisymmetric RNN	5×10^{-3}	-
Hopfield Network	k -Lurie Network	5×10^{-3}	-
Hopfield Network	Neural ODE	1×10^{-3}	0.1@75
Hopfield Network	Antisymmetric RNN	5×10^{-3}	-
Simple Attractor	Neural ODE	1×10^{-3}	-
Simple Attractor	Antisymmetric RNN	5×10^{-3}	-
GC Hopfield Network	Lipschitz RNN	5×10^{-3}	0.2 @ 60
GC Hopfield Network	GLN	3×10^{-3}	$\frac{1}{3}$ @ 60
GC Simple Attractor	GLN	1×10^{-2}	-
GC Simple Attractor	k -Lurie Network	1×10^{-2}	-
GC Simple Attractor	Neural ODE	5×10^{-3}	0.5 @ 40

across all tasks; (ii) SVD combo performs well on the opinion dataset, where the 1-contraction assumption is valid, but struggles on the others; (iii) the Lipschitz RNN struggles on the attractor dataset whilst the antisymmetric RNN struggles across the board due to the A matrix being fixed at zero and the eigenvalues of the C matrix being fixed to almost purely imaginary values. Figures 6.2, 6.5, 6.8 show a random sample of trajectories from each test set, along with the associated predictions.

Table 6.4 compares the generalisation and robustness of the models on each task. No new models were trained, instead the best models from Table 6.3 were directly applied to these OOD and noisy datasets (Section 6.6.1). The k -contracting Lurie network performs the best on all of these datasets and for some, it still demonstrates a MSE of an order of magnitude lower than the next best model. The MSE of the unconstrained Lurie network and neural ODE tended to drop off for these datasets whereas the MSE of the constrained models, excluding the antisymmetric RNN, tended to stay fairly consistent when their assumptions were valid. Figures 6.3, 6.6, 6.9 show a random sample of trajectories, initialised OOD, and associated predictions. Figures 6.4, 6.7, 6.10 show the same for a random sample of trajectories, initialised OOD with additive noise. The k -contracting Lurie network predicted the correct long-term behaviour most accurately under all conditions; even when noise was added, the error was predominantly present during the initial transient.

TABLE 6.3: Average and best MSE on isolated test sets.

Model	Opinion	MSE (mean \pm std, best)		
		Hopfield	Attractor	
k -Lurie Network	$(8.0 \pm 3.0, 5.10) \times 10^{-5}$	$(1.5 \pm 1.0, \mathbf{0.26}) \times 10^{-2}$	$(3.5 \pm 1.0, \mathbf{1.70}) \times 10^{-3}$	
Lurie Network	$(3.7 \pm 4.0, 0.53) \times 10^{-3}$	$(3.6 \pm 2.0, 0.39) \times 10^{-1}$	$(5.1 \pm 5.0, 0.57) \times 10^{-1}$	
Neural ODE	$(2.0 \pm 2.0, \mathbf{0.43}) \times 10^{-4}$	$(2.5 \pm 1.0, 1.50) \times 10^{-2}$	$(2.0 \pm 1.0, 1.00) \times 10^{-2}$	
Lipschitz RNN	$(3.9 \pm 3.0, 0.88) \times 10^{-2}$	$(2.9 \pm 2.0, 0.30) \times 10^{-1}$	$1.48 \pm 2.0, 1.10 \times 10^{-2}$	
SVD Combo	$(8.6 \pm 3.0, 5.70) \times 10^{-4}$	$(2.9 \pm 0.6, 2.10) \times 10^{-1}$	$3.12 \pm 0.5, 2.74$	
Antisym. RNN	$(30.0 \pm 0.2, 28.0) \times 10^{-2}$	$(4.3 \pm 0.2, 4.11) \times 10^{-1}$	$6.93 \pm 0.2, 6.74$	

TABLE 6.4: MSE on OOD and noisy isolated test sets.

Model	MSE (OOD)			MSE (OOD + noisy)		
	Opinion	Hopfield	Attractor	Opinion	Hopfield	Attractor
k -Lurie Network	$\mathbf{2.9} \times 10^{-3}$	$\mathbf{5.6} \times 10^{-2}$	$\mathbf{2.3} \times 10^{-1}$	$\mathbf{2.0} \times 10^{-2}$	$\mathbf{3.2} \times 10^{-1}$	$\mathbf{1.28}$
Lurie Network	6.4×10^{-2}	1.8×10^{-1}	5.96	2.7×10^{-1}	4.4×10^{-1}	6.79
Neural ODE	1.2×10^{-2}	1.09	2.31	3.9×10^{-2}	1.63	4.76
Lipschitz RNN	2.3×10^{-1}	7.3×10^{-1}	6.2×10^{-1}	2.9×10^{-1}	9.7×10^{-1}	1.71
SVD Combo	1.0×10^{-2}	2.38	20.9	3.3×10^{-2}	7.97	30.30
Antisym. RNN	6.43	5.25	52.1	7.29	6.33	52.9

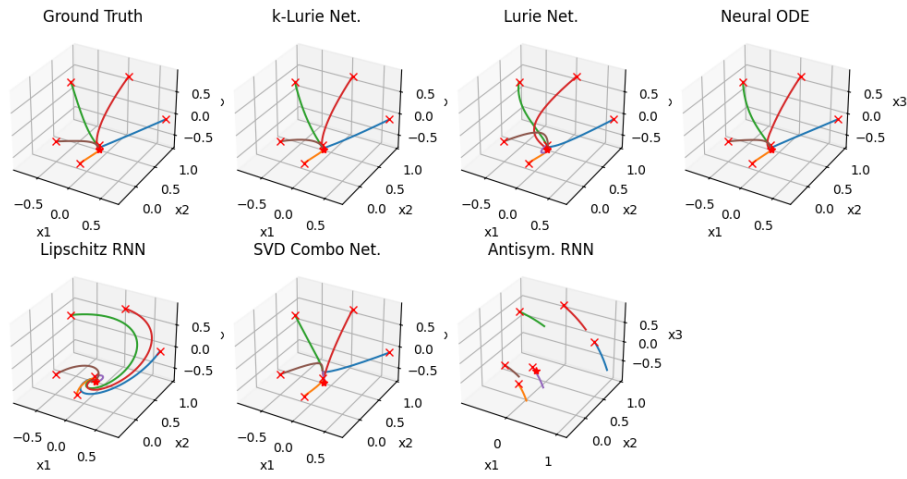


FIGURE 6.2: Ground truth trajectories and the associated predictions of each model for the opinion dynamics test set.

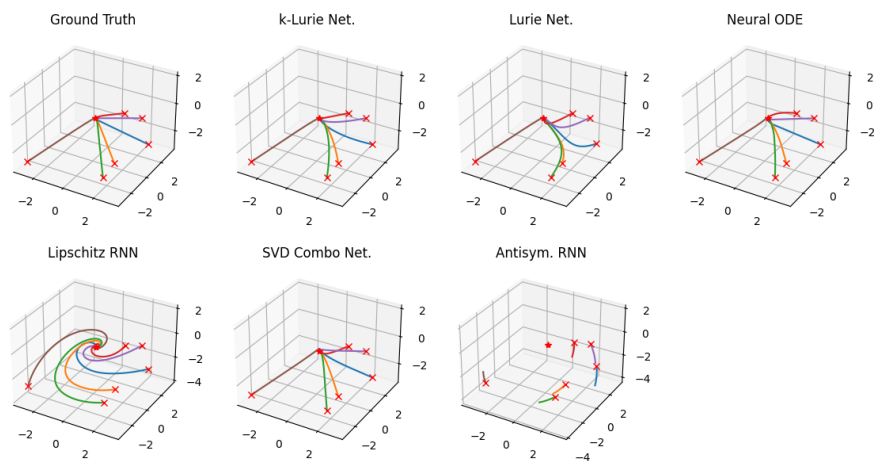


FIGURE 6.3: Ground truth trajectories and the associated predictions of each model for the OOD opinion dynamics test set.

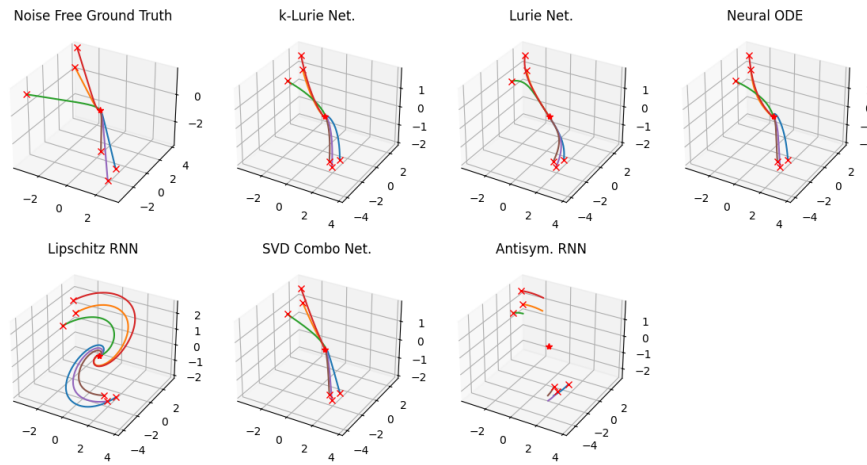


FIGURE 6.4: Ground truth trajectories and the associated predictions of each model for the noisy OOD opinion dynamics test set.

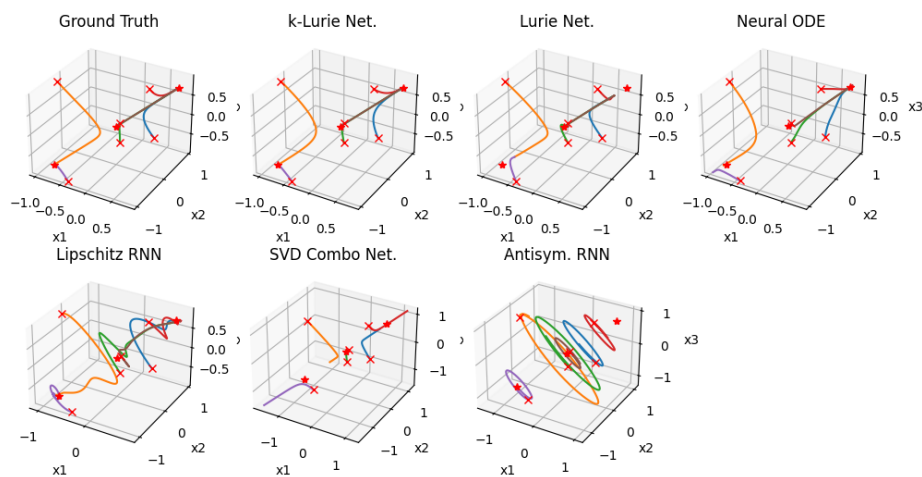


FIGURE 6.5: Ground truth trajectories and the associated predictions of each model for the Hopfield network test set.

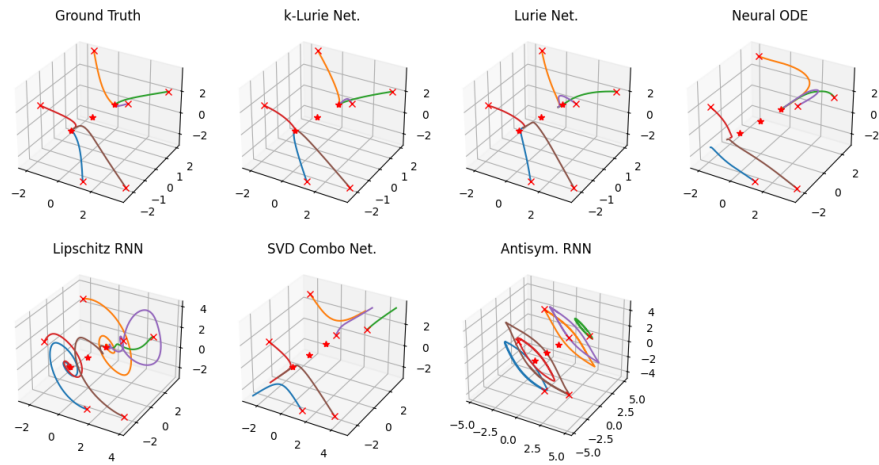


FIGURE 6.6: Ground truth trajectories and the associated predictions of each model for the OOD Hopfield network test set.

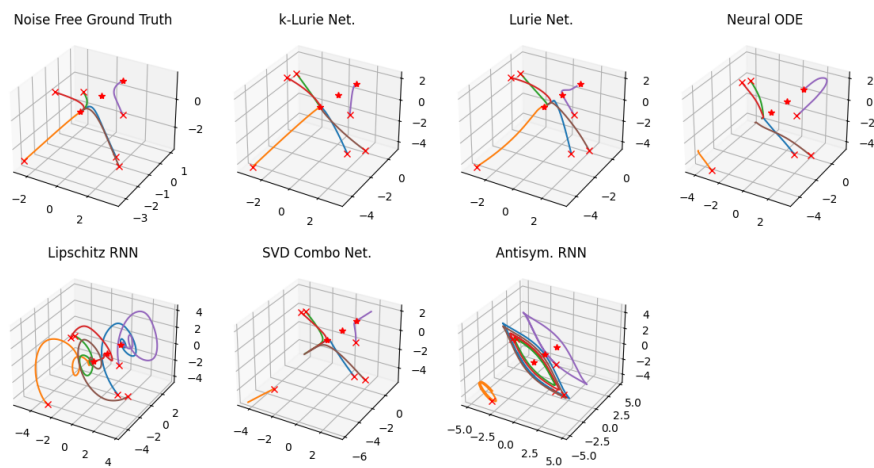


FIGURE 6.7: Ground truth trajectories and the associated predictions of each model for the noisy OOD Hopfield network test set.

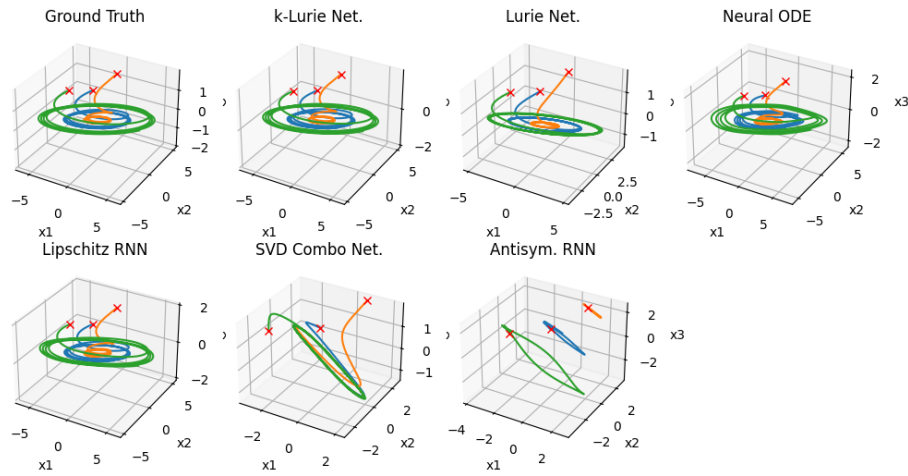


FIGURE 6.8: Ground truth trajectories and the associated predictions of each model for the simple attractor test set.

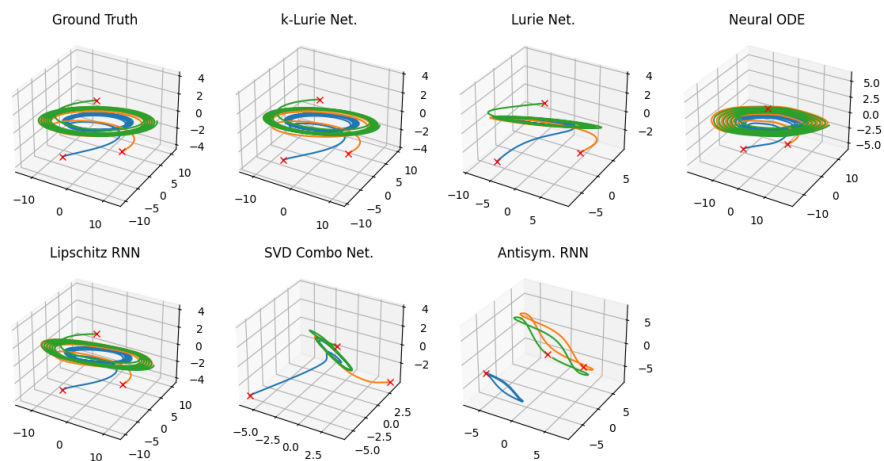


FIGURE 6.9: Ground truth trajectories and the associated predictions of each model for the OOD simple attractor test set.

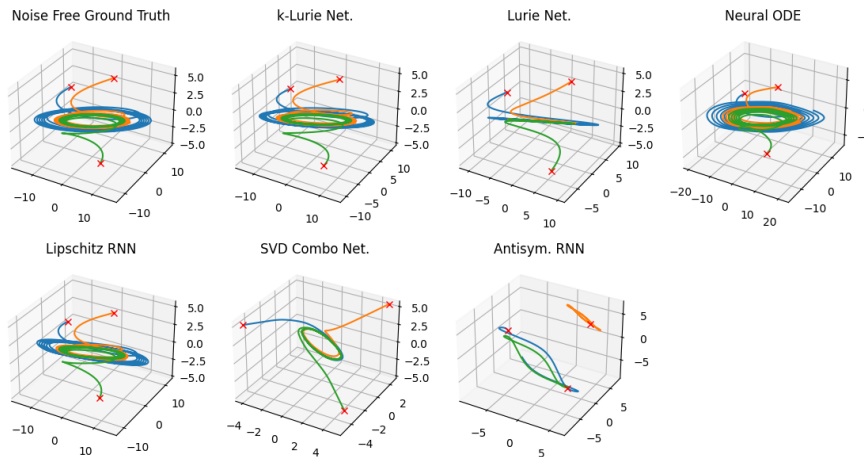


FIGURE 6.10: Ground truth trajectories and the associated predictions of each model for the noisy OOD simple attractor test set.

6.6.4 k -contracting Graph Lurie Networks

This section repeats the same experiments as the previous section, but for two 30-dimensional graph-coupled (GC) dynamical systems: the GC Hopfield network and the GC simple attractor (Section 6.6.1). The state of these datasets is significantly larger than those used in other dynamical systems datasets such as: (i) the LASA dataset [144] where the 2-d trajectories are typically stacked to form 4-d or 8-d trajectories; (ii) simulated datasets of the 2, 4 or 8 link pendulums which, respectively, have 4, 8 or 16 dimension trajectories.

For both datasets, the GLN was constructed according to Theorem 6.7 and Remark 6.6 where $n = m = 3$ and $q = 10$. The individual Lurie networks were constructed according to Theorem 6.5 for the GC Hopfield network, with $k = 2$, and Theorem 6.6 for the GC attractor, with $k = 3$. The neural ODE was formed using two layers with 100 neurons and ReLU activations. The SVD combo leveraged a similar graph structure to the one used in this paper. The other models were the same as in the previous section but with a (qn) -dimensional state.

Table 6.5 shows the GLN had a lower MSE than all other models by a factor of 10. Comparing the GLN, k -contracting Lurie network and the unconstrained Lurie network highlights the improvements due to the k -contraction conditions and the graph structure. Table 6.6 also indicates that the GLN generalised and remained robust to noise, even in these high-dimensional systems. The same can be said for the k -contracting Lurie network which achieved the second lowest MSE on the generalisation and robustness tests. The inherent graph structure of the SVD combo may be the reason

behind its improved ranking in the GC Hopfield tasks whereas the neural ODE particularly struggled to generalise for the GC attractor. Possible explanations behind the poor performance of the constrained benchmark models are suggested in Section 6.8.

TABLE 6.5: Average and best MSE on GC test sets.

Model	MSE (mean \pm std, best)	
	GC Hopfield	GC Attractor
GLN	0.016 \pm 0.0003, 0.016	0.293 \pm 0.1969, 0.015
k -Lurie Network	0.238 \pm 0.0011, 0.237	0.737 \pm 0.0108, 0.723
Lurie Network	2.537 \pm 1.3327, 1.157	291.8 \pm 191.84, 21.83
Neural ODE	0.138 \pm 0.0291, 0.114	3.445 \pm 0.3626, 2.942
Lipschitz RNN	0.124 \pm 0.0161, 0.105	0.658 \pm 0.1604, 0.433
SVD Combo	0.339 \pm 0.1363, 0.229	3.024 \pm 1.1240, 1.435
Antisym. RNN	0.448 \pm 0.0023, 0.444	6.386 \pm 0.0066, 6.380

TABLE 6.6: MSE on OOD and noisy GC test sets.

Model	MSE (OOD)		MSE (OOD + noisy)	
	GC Hopfield	GC Attractor	GC Hopfield	GC Attractor
GLN	0.08	1.67	0.26	2.85
k -Lurie Network	0.77	6.10	1.05	6.90
Lurie Network	20350	5638	25481	6368
Neural ODE	2.12	24.93	2.76	25.11
Lipschitz RNN	3.58	6.17	4.25	7.84
SVD Combo	0.84	11.40	1.07	12.30
Antisym. RNN	5.32	50.68	6.21	52.10

6.7 Empirical Evaluation: Fashion MNIST

It was emphasised in Section 6.1 that convergent dynamics in the brain play a role in forming representations. This was investigated in this section by applying the Lurie network to the Fashion-MNIST (FMNIST) classification task [145].

6.7.1 Data

The FMNIST dataset contains 28×28 grayscale images of 70k fashion products, from 10 different categories. A sample of these images is displayed in Figure 6.11. For all experiments, the images were normalised and zero-centred. The data was split into a training set with 60k images and a test set of 10k images. A batch size of 250 was used throughout (Section 2.1.4).

6.7.2 Models and Training

Two variations of the Lurie network and k -contracting Lurie network were studied. In the first variation, the transformed images were flattened and directly passed into the models through the initial condition. Euler integration was used to approximate the trajectory of (6.2) and the final state was mapped through a linear layer and followed by a softmax layer to obtain the categorical predictions. The k -contraction parameter was set to its highest value, $k = 3$, to permit training over the three convergent behaviours illustrated in Figure 6.1.

In the second variation, the images were pre-processed by a small CNN (55,744 parameters) to artificially replicate a biological neural processing system. The output of this was then passed into the Lurie network and k -contracting Lurie network through the initial conditions. The remainder of the setup was the same as the first variation. Specific details of the architectures are included in Table 6.7 and Table 6.8, where each layer of the CNN used ReLU activation functions. Training settings are detailed in Table 6.9.

Other benchmark models were either highest ranked or biologically-inspired from the recently deprecated comparison site <https://paperswithcode.com/sota/image-classification-on-fashion-mnist>.

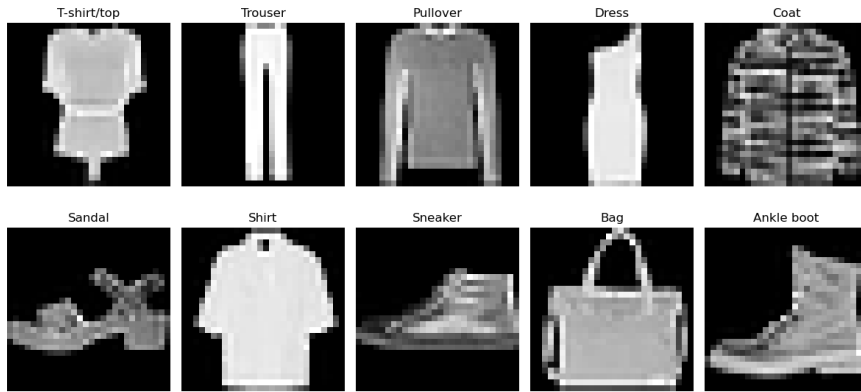


FIGURE 6.11: Random sample of images from each class of FMNIST.

TABLE 6.7: k -Lurie network and Lurie network settings.

Parameter	Description	Without CNN	With CNN
δ	Euler integration step size	1×10^{-2}	1×10^{-2}
N	Number of Euler integration steps	100	100
$\Phi(\cdot)$	Activation function	tanh	tanh
g	Upper bound on slope of activation	1	1
n	dimension of x	784	576
m	dimension of y	784	576
k	k -contraction parameter	3	3

TABLE 6.8: CNN settings.

Layer	Convolution	Downsampling / Reshaping
1	2D convolution (1 input, 32 outputs, 3×3 kernel)	2D max pool (2×2 kernel)
2	2D convolution (32 inputs, 64 outputs, 3×3 kernel)	2D max pool (2×2 kernel)
3	2D convolution (64 inputs, 64 outputs, 3×3 kernel)	Flatten (1D output)

TABLE 6.9: Training settings used by the Lurie network and k -contracting Lurie network with and without the CNN.

Parameter	Lurie Network	k -Lurie Network
Loss	Cross Entropy	Cross Entropy
Optimiser	SGD	SGD
Weight Decay	1×10^{-5}	1×10^{-5}
Epochs	50	50
LR	1×10^{-3}	1×10^{-2}
LR Cut @ Epoch	0.5 @ 35	0.5 @ 35

TABLE 6.10: Classification accuracy on FMNIST test set.

Model	Parameter Count	Accuracy (%)
LR-Net	1,028,234	95.03
Inception v3	23,851,784	94.44
CNN + Wilson-Cowan RNN	5,179,521	91.35
Wilson-Cowan RNN	-	88.39
CNN + k -Lurie Network	1,057,994	91.33
CNN + Lurie Network	1,057,994	90.95
Lurie Network	1,853,386	89.64
k -Lurie Network	1,853,386	85.78

6.7.3 Results

Table 6.10 highlights that the k -contraction constraints result in only a 3.86% drop in classification accuracy compared to the Lurie network. This suggests the k -contraction constraints do not limit the expressivity of the Lurie network too significantly, even with the scalar metric limitation. Figure 6.12 illustrates the application of a t-SNE projection to the state of the k -contracting Lurie network at the final time step. The figure highlights how the k -contracting Lurie network is able to untangle the classes with a degree of interpretability by locating similar classes, such as footwear, closer together in the latent space. The Lurie network also achieves the highest accuracy on this task for models which do not involve a CNN.

When leveraging a small CNN to pre-process the images, the k -contracting Lurie network achieved superior performance. This accuracy is similar to the CNN + Wilson-Cowan RNN [146] which has nearly $5\times$ the number of parameters, only 12% of which are attributed to the Wilson-Cowan RNN. When combined with the CNN, the k -contracting Lurie network outperformed the Lurie network. This is in contrast to without the CNN, where one possible explanation is that the CNN compensates for the reduced expressivity, whilst the k -contraction constraints help to effectively process the information over the 100 Euler integration steps (Table 6.7).

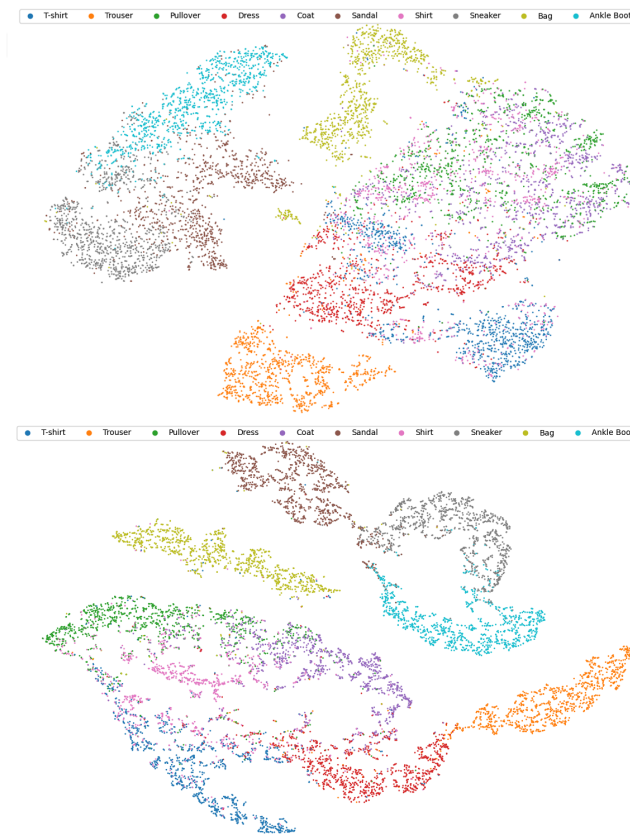


FIGURE 6.12: t-SNE plots of FMNIST test set (top) and output of k -contracting Lurie network with test set as initial conditions (bottom).

6.8 Related Work

Several constrained continuous-time RNN models exist in the ML literature. The model structure of three notable examples were presented in Section 6.3.1 as they happen to be special cases of a Lurie network, when modelling time-invariant systems. The antisymmetric RNN [30] and the Lipschitz RNN [31] were designed to address the exploding and vanishing gradient problem [147]. The antisymmetric RNN did so by parametrising the RNN such that the real eigenvalues of the Jacobian were zero. It achieved this by setting $A = 0$ and restricting C to being skew-symmetric. Whilst this does prevent the gradients from exploding and vanishing, it restricts the dynamics which the model can learn to purely oscillatory behaviour. The Lipschitz RNN has a more relaxed parametrisation. This model constructs the A and C matrices such that they are both convex combinations of symmetric and skew-symmetric matrices. However, the weight of the symmetric matrix can only vary between 0 and 0.5, whereas the weight of the skew-symmetric matrix can vary between 0.5 and 1. Again, this addresses the vanishing and exploding gradient problem, but the model cannot encode dynamics which are predominately decaying or growing. Finally, the RNN proposed in [32] has biological motivations and encodes 1-contracting dynamics. This implies that all possible trajectories will exponentially converge, making the model robust to input disturbances.

Whilst the models above address a variety of problems, each model is quite limited in the range of dynamics they can learn, which is a problem when trying to design a generalised model for learning time-invariant systems. With respect to (6.2), the Lurie network has more flexibility than all of these models. Firstly, it includes all three weight matrices (A, B, C) and biases (b_x, b_y), whereas the models mentioned above fix at least one of these parameters. Secondly, the constraints imposed, and the corresponding parametrisations, allow the model to learn a variety of dynamics including, but not limited to, those mentioned above. The only limitation is that the dynamics must converge in some way. This includes certain types of chaotic behaviour, such as that demonstrated by Thomas' cyclically symmetric attractor [148]. At the edge of the chaotic regime, the trajectories converge to a strange attractor, which can be modelled by a 3-contracting Lurie network.

The Lurie network is also related to a class of feed-forward models named *implicit or equilibrium networks*. These models use an implicit equation to express the relationship between the model output, layer outputs and model input in a compact vectorised form [149]. Like the Lurie network, these models can be represented by the interconnection of a linear time-invariant system and a nonlinearity. This makes analysis tools from control theory, such as Lipschitz bounds, applicable to these models [150]. An additional connection is that the solution to the implicit equations correspond to equilibrium points of a Lurie system [151].

As mentioned in the introduction, the k -contraction constraints used in this chapter have an interesting connection to some properties observed in biological learning systems. A 2-contracting model can replicate the behaviour of associative memory, where every stored pattern corresponds to an equilibrium [132]. Furthermore, a 3-contracting model can replicate the dynamics of working memory, where patterns are retained as attractor states [128]. Due to the 2-contraction (3-contraction) constraints, the equilibrium points (attractor states) must all fall on a line (plane). As a result, the conditions developed in this chapter could be of interest to neuroscientists and ML researchers interested in memory storage and retrieval [152; 127; 130].

The relationship between properties guaranteed by k -contraction analysis and those observed in associative and working memory suggest the k -contracting Lurie network possesses a number of appealing properties for an ML model; hence, it may be suitable for a wider class of ML problems beyond system identification for dynamical systems. This proposition is supported by the successful application of the Lurie network on FMNIST and application of the other stability-constrained RNN models (Section 6.3.1) on a wide range of ML tasks. Since the Lurie network is a more structured, time-invariant example of a neural ODE [56], it will also be applicable to a similar array of tasks, such as continuous normalising flows [153]. The only limiting requirement is that the input must be passed in through the initial condition.

6.9 Conclusion

The Lurie network was presented as a novel and unifying time-invariant neural ODE. Absolute stability-like results were presented which verified k -contraction of the Lurie network. These results were incorporated in a principled approach for constructing k -contracting Lurie networks and graph Lurie networks with more expressiveness than existing stability-constrained models, whilst still having convergent guarantees. Empirical results showed the benefit of the graph structure and the k -contraction constraints through improved prediction accuracy, out of distribution generalisation and robustness on a range of dynamical systems datasets, including multi-stable and orbitally stable systems. Further experiments highlighted the models ability to form interpretable representations for other ML tasks.

Chapter 7

Conclusion

7.1 Thesis Summary

This thesis was underpinned by the observation that many dynamical systems involving neural networks (NNs), which appear in machine learning (ML) and control, are special cases of a *Forced Lurie System*. This observation allowed the analysis of these systems to be framed as an absolute stability problem, opening the door to decades of existing research. Building upon this, two key outcomes were achieved in this thesis:

1. Reduced conservatism of absolute stability problems involving NNs.
2. A robust and scalable ML method for the identification of convergent systems.

The contributions of the individual chapters are summarised next.

Chapter 3 established tailored quadratic constraints (QCs) for the repeated ReLU function and exploited these to derive less conservative stability criteria, presented as linear matrix inequalities (LMIs), for Lurie systems with repeated ReLU nonlinearities. Numerical examples showed that the new criteria were tractable for the high-dimensional Hopfield networks, whilst being significantly less conservative than existing criteria.

Chapter 4 derived parallel results to Chapter 3, but for discrete-time Lurie systems with repeated ReLU nonlinearities. This is a more natural setting since NN controllers would typically be implemented digitally. Additionally, local stability analysis showed that provided the system had a unique equilibrium point, then under some mild conditions, it must be globally stable or unstable. Numerical examples not only showed that the new criteria were tractable for high-dimensional systems, but they strike an appealing balance between conservatism and computational complexity.

Chapter 5 extended the results in Chapter 3 to a wider class of NN activation functions, including the ReLU and leaky ReLU functions. This was achieved by first establishing a loop transformation between Lurie systems with repeated (leaky) ReLU nonlinearities and Lurie systems with repeated magnitude nonlinearities. Novel QCs were then established for the repeated magnitude function and were leveraged to obtain novel stability criteria, posed as LMIs, for Lurie systems with repeated magnitude nonlinearities. Similar to Chapter 4, a local stability analysis showed that provided the nonlinearity was positively homogenous and the system had a unique equilibrium point, then under some mild conditions, it must be globally stable or unstable. Numerical examples showed that the criteria developed in this chapter were competitive to those from Chapter 3, whilst being applicable to a larger class of nonlinearities.

Chapter 6 introduced the Lurie network as a unifying time-invariant neural ODE. To capture a wider range of convergent dynamics, imposing stability in the traditional control theoretic sense was too restrictive; instead, k -contraction analysis was used to obtain an absolute stability-like result for Lurie network's with slope-restricted nonlinearities. Based on this result, an unconstrained parametrisation of the Lurie network was established to enable gradient-based training. Both of these results were extended to explicitly model k -contracting graph coupled Lurie networks. Empirical results, for systems with multiple equilibrium points and limit cycles, showed improved robustness and out of distribution generalisation when the k -contraction condition and graph structure were included as inductive biases.

7.2 Future Work

This section outlines several promising directions for future research based on the work in this thesis.

Scalability

The semi-definite programming (SDP) based stability criteria, developed in this thesis, have been applied to small NNs with up to 100 neurons. NNs of this size have demonstrated some practical use, for example in [21]; however, for NNs such as the one used in [1], the criteria need to handle systems in the order of 1,000 neurons. Of course, one simple approach is to leverage more intensive computational resources, but this is expensive. A more interesting theoretical approach would be to study whether any features of the problem can be exploited to speed up the computation of a solution, whilst maintaining the reduced conservatism. The work on chordal sparsity for NN verification [44] could potentially be adopted in this setting.

Robustness

In Chapters 3 - 5, perfect knowledge of the systems under consideration was assumed. In reality, uncertainties and other troublesome components will exist within the system; for example, other nonlinearities, model uncertainty, and time delays may all be present and should be accounted for. This is the principle behind *robust control theory* [34]. As discussed in Section 2.3.3, the integral quadratic constraint (IQC) approach provides a natural way to handle systems with multiple uncertainties. Formulating the results developed in this thesis within the IQC framework would provide a way to gain the benefit of this work in a more realistic setting.

Neural Network Biases

In Chapters 3 - 5, it was assumed the biases of the NN were set to zero, so that the origin was an equilibrium point of the Lurie system. For the same reason, this simplifying assumption was also made in other work [20; 21]. Of course, the biases play an important role in the NNs ability to represent different functions. Consider a single NN layer, $y = \Phi(Wu + b)$, with parameters W, b . Without the bias parameter, b , the neuron can only learn functions that pass through the origin. Thus, the bias shifts the input to the activation function, which may be critical for successful learning [46, Chapter 6]. Developing stability criteria which incorporate biases would permit more flexible NNs to be deployed on safety-critical systems.

Generalised Contraction Metrics

The Lurie network was proposed in Chapter 6 as a ML model with built in convergence guarantees. The main advantage of this model is that it accounts for a wider class of convergent dynamics than existing stability-constrained approaches; for example, it can model systems with multiple equilibrium points and limit cycles. However, the expressiveness of the model is limited to systems which are k -contracting in a scalar metric. Work such as [154; 138] has developed criteria for verifying if a system is k -contracting in a more general metric, finding parametrisations of such results would increase the expressiveness of the Lurie network, giving it more freedom to approximate a wider class of systems.

Outlook

This thesis lays the groundwork for certifiably stable control systems with NN involvement, and has several promising directions for future research. Improving the scalability of SDP-based stability criteria is essential for handling larger NNs, and techniques exploiting problem structure, such as chordal sparsity, offer a promising path forward. Incorporating robustness to uncertainty via frameworks like IQC analysis would make the results more applicable to real-world systems, where perfect knowledge is rarely available. Relaxing the assumption of zero biases would enable the use of more expressive NNs in safety-critical settings. Finally, extending the Lurie network framework

with generalised contraction metrics could significantly broaden the class of dynamical systems that can be modelled with built-in convergence guarantees. Collectively, these directions aim to advance the development of scalable, robust, and theoretically grounded learning and control.

References

- [1] J. Degraeve, F. Felici, J. Buchli, M. Neunert, B. Tracey, F. Carpanese, T. Ewalds, R. Hafner, A. Abdolmaleki, D. de Las Casas, *et al.*, “Magnetic control of tokamak plasmas through deep reinforcement learning,” *Nature*, vol. 602, no. 7897, pp. 414–419, 2022.
- [2] C. R. Richardson, M. C. Turner, and S. R. Gunn, “Strengthened Circle and Popov Criteria for the stability analysis of feedback systems with ReLU neural networks,” *IEEE Control Systems Letters*, 2023.
- [3] C. R. Richardson, M. C. Turner, and S. R. Gunn, “Strengthened Circle and Popov Criteria and the analysis of ReLU neural networks,” in *2024 UKACC 14th International Conference on Control (CONTROL)*, pp. 127–128, IEEE, 2024.
- [4] C. R. Richardson, M. C. Turner, S. R. Gunn, and R. Drummond, “Strengthened stability analysis of discrete-time Lurie systems involving ReLU neural networks,” in *6th Annual Learning for Dynamics and Control Conference*, pp. 209–221, PMLR, 2024.
- [5] C. R. Richardson, M. C. Turner, and S. R. Gunn, “Lurie networks with k-contracting dynamics,” in *New Frontiers in Associative Memory Workshop at ICLR*, 2025.
- [6] C. R. Richardson, M. C. Turner, and S. R. Gunn, “Lurie networks with robust convergent dynamics,” *TMLR: Transactions on Machine Learning Research*, 2025.
- [7] C. R. Richardson, M. C. Turner, and S. R. Gunn, “Analysis of Lurie systems with magnitude nonlinearities and connections to neural network stability analysis,” *IEEE Transactions on Automatic Control*, 2026.
- [8] B. L. Stevens, F. L. Lewis, and E. N. Johnson, *Aircraft control and simulation: dynamics, controls design, and autonomous systems*. John Wiley & Sons, 2015.
- [9] D. E. Seborg, T. F. Edgar, D. A. Mellichamp, and F. J. Doyle III, *Process dynamics and control*. John Wiley & Sons, 2016.

- [10] B. W. Bequette, "A critical assessment of algorithms and challenges in the development of a closed-loop artificial pancreas," *Diabetes technology & therapeutics*, vol. 7, no. 1, pp. 28–47, 2005.
- [11] R. H. Battin, "Space guidance evolution-a personal narrative," *Journal of Guidance, Control, and Dynamics*, vol. 5, no. 2, pp. 97–110, 1982.
- [12] D. A. Mindell, *Digital Apollo: Human and machine in spaceflight*. Mit Press, 2011.
- [13] F. Blanchini, "Set invariance in control," *Automatica*, vol. 35, no. 11, pp. 1747–1767, 1999.
- [14] H. K. Khalil, "Nonlinear systems," *Patience Hall*, vol. 115, 2002.
- [15] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," *arXiv preprint arXiv:1509.02971*, 2015.
- [16] H. Ravichandar, A. S. Polydoros, S. Chernova, and A. Billard, "Recent advances in robot learning from demonstration," *Annual review of control, robotics, and autonomous syst.*, vol. 3, pp. 297–330, 2020.
- [17] H. Yu, S. Park, A. Bayen, S. Moura, and M. Krstic, "Reinforcement learning versus PDE backstepping and PI control for congested freeway traffic," *IEEE Transactions on Control Syst. Tech.*, vol. 30, no. 4, pp. 1595–1611, 2021.
- [18] S. Levine, A. Kumar, G. Tucker, and J. Fu, "Offline reinforcement learning: Tutorial, review, and perspectives on open problems," *arXiv preprint arXiv:2005.01643*, 2020.
- [19] L. Brunke, M. Greeff, A. W. Hall, Z. Yuan, S. Zhou, J. Panerati, and A. P. Schoellig, "Safe learning in robotics: From learning-based control to safe reinforcement learning," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, no. 1, pp. 411–444, 2022.
- [20] H. Yin, P. Seiler, and M. Arcak, "Stability analysis using quadratic constraints for systems with neural network controllers," *IEEE Transactions on Automatic Control*, vol. 67, no. 4, pp. 1980–1987, 2021.
- [21] H. Yin, P. Seiler, M. Jin, and M. Arcak, "Imitation learning with stability and safety guarantees," *IEEE Control Systems Letters*, vol. 6, pp. 409–414, 2021.
- [22] A. Gu, K. Goel, and C. Ré, "Efficiently modelling long sequences with structured state spaces," *arXiv preprint arXiv:2111.00396*, 2021.
- [23] H. Zhou, S. Zhang, J. Peng, S. Zhang, J. Li, H. Xiong, and W. Zhang, "Informer: Beyond efficient transformer for long sequence time-series forecasting," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 35, pp. 11106–11115, 2021.

- [24] S. Lanthaler, T. K. Rusch, and S. Mishra, "Neural oscillators are universal," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [25] K. O'Shea and R. Nash, "An introduction to convolutional neural networks," *arXiv preprint arXiv:1511.08458*, 2015.
- [26] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [27] T. K. Rusch, B. Chamberlain, J. Rowbottom, S. Mishra, and M. Bronstein, "Graph-coupled oscillator networks," in *International Conference on Machine Learning*, pp. 18888–18909, PMLR, 2022.
- [28] N. B. Erichson, M. Muehlebach, and M. W. Mahoney, "Physics-informed autoencoders for Lyapunov-stable fluid flow prediction," *arXiv preprint arXiv:1905.10866*, 2019.
- [29] L. Kozachkov, M. Lundqvist, J.-J. Slotine, and E. K. Miller, "Achieving stable dynamics in neural circuits," *PLoS computational biology*, vol. 16, no. 8, p. e1007659, 2020.
- [30] B. Chang, M. Chen, E. Haber, and E. H. Chi, "Antisymmetricrnn: A dynamical system view on recurrent neural networks," *arXiv preprint arXiv:1902.09689*, 2019.
- [31] N. B. Erichson, O. Azencot, A. Queiruga, L. Hodgkinson, and M. W. Mahoney, "Lipschitz recurrent neural networks," *arXiv preprint arXiv:2006.12070*, 2020.
- [32] L. Kozachkov, M. Ennis, and J.-J. Slotine, "RNNs of RNNs: Recursive construction of stable assemblies of recurrent neural networks," *Advances in Neural Information Processing Systems*, vol. 35, pp. 30512–30527, 2022.
- [33] M. M. Bronstein, J. Bruna, T. Cohen, and P. Veličković, "Geometric deep learning: Grids, groups, graphs, geodesics, and gauges," *arXiv preprint arXiv:2104.13478*, 2021.
- [34] M. Green and D. J. Limebeer, *Linear robust control*. Courier Corporation, 2012.
- [35] M. Vidyasagar, *Nonlinear systems analysis*. SIAM, 2002.
- [36] I. R. Manchester, R. Wang, and N. H. Barbara, "Neural networks in the loop: Learning with stability and robustness guarantees," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 9, 2026.
- [37] M. Anghel, F. Milano, and A. Papachristodoulou, "Algorithmic construction of Lyapunov functions for power system stability analysis," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 60, no. 9, pp. 2533–2546, 2013.

- [38] S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan, *Linear matrix inequalities in system and control theory*. SIAM, 1994.
- [39] S. P. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
- [40] J. F. Sturm, "Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones," *Optimization methods and software*, vol. 11, no. 1-4, pp. 625–653, 1999.
- [41] K.-C. Toh, M. J. Todd, and R. H. Tütüncü, "SDPT3—a MATLAB software package for semidefinite programming, version 1.3," *Optimization methods and software*, vol. 11, no. 1-4, pp. 545–581, 1999.
- [42] M. ApS, "Mosek optimization toolbox for MATLAB," *User's Guide and Reference Manual, Version*, vol. 4, no. 1, p. 116, 2019.
- [43] J. Lofberg, "YALMIP: A toolbox for modelling and optimisation in MATLAB," in *2004 IEEE international conference on robotics and automation (IEEE Cat. No. 04CH37508)*, pp. 284–289, IEEE, 2004.
- [44] M. Newton and A. Papachristodoulou, "Exploiting sparsity for neural network verification," in *Learning for dynamics and control*, pp. 715–727, PMLR, 2021.
- [45] J. Nocedal and S. J. Wright, *Numerical optimization*. Springer, 2006.
- [46] I. Goodfellow, Y. Bengio, A. Courville, and Y. Bengio, *Deep learning*, vol. 1. MIT press Cambridge, 2016.
- [47] E. Bar-Shalom, O. Dalin, and M. Margaliot, "Compound matrices in systems and control theory: a tutorial," *Mathematics of Control, Signals, and Systems*, pp. 1–55, 2023.
- [48] C. Wu, I. Kanevskiy, and M. Margaliot, "k-contraction: Theory and applications," *Automatica*, vol. 136, p. 110048, 2022.
- [49] W. Lohmiller and J.-J. E. Slotine, "On contraction analysis for non-linear systems," *Automatica*, vol. 34, no. 6, pp. 683–696, 1998.
- [50] M. R. Liberzon, "Lur'e problem of absolute stability—a historical essay," *IFAC Proceedings Volumes*, vol. 34, no. 6, pp. 25–28, 2001.
- [51] A. A. Adegbege and W. P. Heath, "Multivariable algebraic loops with complementarity constraints enforcing some KKT conditions," in *2014 52nd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 1033–1039, IEEE, 2014.

- [52] A. A. Adegbege and W. P. Heath, "Multivariable algebraic loops in linear anti-windup implementations," in *2015 23rd Mediterranean Conference on Control and Automation (MED)*, pp. 514–519, IEEE, 2015.
- [53] A. A. Adegbege and W. P. Heath, "A framework for multivariable algebraic loops in linear anti-windup implementations," *Automatica*, vol. 83, pp. 81–90, 2017.
- [54] G. Valmorbidia, R. Drummond, and S. R. Duncan, "Regional analysis of slope-restricted Lurie systems," *IEEE Transactions on Automatic Control*, vol. 64, no. 3, pp. 1201–1208, 2018.
- [55] P. Pauli, D. Gramlich, J. Berberich, and F. Allgöwer, "Linear systems with neural network nonlinearities: Improved stability analysis via acausal Zames-Falb multipliers," in *2021 60th IEEE Conference on Decision and Control (CDC)*, pp. 3611–3618, IEEE, 2021.
- [56] R. T. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, "Neural ordinary differential equations," *Advances in neural information processing systems*, vol. 31, 2018.
- [57] V. A. Yakubovich, "Frequency conditions for the absolute stability and dissipativity of control systems with a single differentiable nonlinearity," in *Soviet Mathematics Doklady*, vol. 6, p. 101, 1965.
- [58] G. Zames and P. L. Falb, "Stability conditions for systems with monotone and slope restricted nonlinearities," *SIAM J. of Control*, vol. 6, no. 1, pp. 89–108, 1968.
- [59] A. Megretski and A. Rantzer, "System analysis via integral quadratic constraints," *IEEE transactions on automatic control*, vol. 42, no. 6, pp. 819–830, 1997.
- [60] U. Jonsson and A. Megretski, "The zames-falb iqc for systems with integrators," *IEEE Transactions on Automatic Control*, vol. 45, no. 3, pp. 560–565, 2002.
- [61] J. Carrasco, M. C. Turner, and W. P. Heath, "Zames–Falb multipliers for absolute stability: From O’Shea’s contribution to convex searches," *European Journal of Control*, vol. 28, pp. 1–19, 2016.
- [62] M. Fetzner and C. W. Scherer, "Zames–falb multipliers for invariance," *IEEE Control Systems Letters*, vol. 1, no. 2, pp. 412–417, 2017.
- [63] M. C. Turner and R. Drummond, "Analysis of systems with slope restricted nonlinearities using externally positive Zames–Falb multipliers," *IEEE Transactions on Automatic Control*, vol. 65, no. 4, pp. 1660–1667, 2019.
- [64] R. Drummond, C. Guiver, and M. C. Turner, "Exponential input-to-state stability for Lur’e systems via integral quadratic constraints and Zames–Falb multipliers," *IMA Journal of Mathematical Control and Information*, p. dnae003, 2024.

- [65] J. Carrasco, W. Heath, and P. Seiler, "Absolute stability methods," in *Encyclopedia of Systems and Control Engineering*, 2025.
- [66] R. J. Caverly and J. R. Forbes, "Lmi properties and applications in systems, stability, and control theory," *arXiv preprint arXiv:1903.08599*, 2019.
- [67] U. Jönsson, *Lecture notes on integral quadratic constraints*. KTH, 2000.
- [68] M. C. Turner, "Zames-falb multipliers: don't panic," *arXiv preprint arXiv:2106.15913*, 2021.
- [69] W. M. Haddad and D. S. Bernstein, "Explicit construction of quadratic Lyapunov functions for the Small gain, Positivity, Circle, and Popov Theorems and their application to robust stability. Part II: Discrete-time theory," *Int. Journal of Robust and Nonlinear Control*, vol. 4, no. 2, pp. 249–265, 1994.
- [70] P. Park, "A revisited Popov criterion for nonlinear Lur'e systems with sector-restrictions," *International Journal of Control*, vol. 68, no. 3, pp. 461–470, 1997.
- [71] W. P. Heath and G. Li, "Lyapunov functions for the multivariable Popov Criterion with indefinite multipliers," *Automatica*, vol. 45, no. 12, pp. 2977–2981, 2009.
- [72] P. Park, "Stability criteria of sector-and slope-restricted Lur'e systems," *IEEE Transactions on Automatic Control*, vol. 47, no. 2, pp. 308–313, 2002.
- [73] J. Park, S. Y. Lee, and P. Park, "A less conservative stability criterion for discrete-time Lur'e systems with sector and slope restrictions," *IEEE Transactions on Automatic Control*, vol. 64, no. 10, pp. 4391–4395, 2019.
- [74] R. Drummond and G. Valmorbida, "Generalised Lyapunov functions for discrete-time Lurie systems with slope-restricted nonlinearities," *IEEE Transactions on Automatic Control*, 2023.
- [75] R. P. O'Shea, "A combined frequency-time domain stability criterion for autonomous continuous systems," *IEEE Transactions on Automatic Control*, vol. 11, no. 3, pp. 477–484, 1966.
- [76] R. P. O'Shea, "An improved frequency-time domain stability criterion for autonomous continuous systems," *IEEE Transactions on Automatic Control*, vol. 12, no. 6, pp. 725–731, 1967.
- [77] X. Chen and J. T. Wen, "Robustness analysis for linear time-invariant systems with structured incrementally sector bounded feedback nonlinearities," 1996.
- [78] M. C. Turner, M. Kerr, and I. Postlethwaite, "On the existence of stable, causal multipliers for systems with slope-restricted nonlinearities," *IEEE Transactions on Automatic Control*, vol. 54, no. 11, pp. 2697–2702, 2009.

- [79] M. Chang, R. Mancera, and M. Safonov, "Computation of zames-falb multipliers revisited," *IEEE Transactions on Automatic Control*, vol. 57, no. 4, pp. 1024–1029, 2011.
- [80] J. Carrasco, M. Maya-Gonzalez, A. Lanzon, and W. P. Heath, "Lmi searches for anticausal and noncausal rational zames–falb multipliers," *Systems & Control Letters*, vol. 70, pp. 17–22, 2014.
- [81] J. Carrasco, W. P. Heath, J. Zhang, N. S. Ahmad, and S. Wang, "Convex searches for discrete-time Zames–Falb multipliers," *IEEE Transactions on Automatic Control*, vol. 65, no. 11, pp. 4538–4553, 2019.
- [82] M. C. Turner and R. Drummond, "Discrete-time systems with slope restricted nonlinearities: Zames–falb multiplier analysis using external positivity," *International Journal of Robust and Nonlinear Control*, vol. 31, no. 6, pp. 2255–2273, 2021.
- [83] M. Fetzner and C. W. Scherer, "Full-block multipliers for repeated, slope restricted scalar nonlinearities," *International Journal of Robust Nonlinear Control*, vol. 27, no. 17, pp. 3376–3411, 2017.
- [84] V. V. Kulkarni and M. G. Safonov, "All multipliers for repeated monotone nonlinearities," *IEEE Transactions on Automatic Control*, vol. 47, no. 7, pp. 1209–1212, 2002.
- [85] F. J. D'Amato, M. A. Rotea, A. Megretski, and U. Jönsson, "New results for analysis of systems with repeated nonlinearities," *Automatica*, vol. 37, no. 5, pp. 739–747, 2001.
- [86] R. Mancera and M. G. Safonov, "All stability multipliers for repeated MIMO nonlinearities," *Systems & Control Letters*, vol. 54, no. 4, pp. 389–397, 2005.
- [87] T. Yuno, S. Nishinaka, R. Saeki, and Y. Ebihara, "On static O'Shea-Zames-Falb multipliers for idempotent nonlinearities," in *2024 IEEE 63rd Conference on Decision and Control (CDC)*, pp. 5870–5875, IEEE, 2024.
- [88] P. Seiler, "Stability analysis with dissipation inequalities and integral quadratic constraints," *IEEE Transactions on Automatic Control*, vol. 60, no. 6, pp. 1704–1709, 2014.
- [89] J. Grönqvist and A. Rantzer, "Integral quadratic constraints for neural networks," in *2022 European Control Conference (ECC)*, pp. 1864–1869, IEEE, 2022.
- [90] U. Jönsson, "Stability analysis with popov multipliers and integral quadratic constraints," *Systems & Control Letters*, vol. 31, no. 2, pp. 85–92, 1997.
- [91] M. C. Turner and R. Drummond, "Analysis of MIMO Lurie systems with slope restricted nonlinearities using concepts of external positivity," in *2019 IEEE 58th Conference on Decision and Control (CDC)*, pp. 163–168, IEEE, 2019.

- [92] J. J. Hopfield, "Neural networks and physical systems with emergent collective computational abilities.," *Proceedings of the national academy of sciences*, vol. 79, no. 8, pp. 2554–2558, 1982.
- [93] G. A. Leonov, D. V. Ponomarenko, and V. B. Smirnova, *Frequency-Domain Methods for Nonlinear Analysis*. Singapore: World Scientific, 1996.
- [94] A. d. L. J. Bertolin, R. C. L. F. Oliveira, G. Valmorbidia, and P. L. D. Peres, "Dynamic output-feedback control of continuous-time Lur'e systems using Zames-Falb multipliers by means of an LMI-based algorithm," *IFAC-PapersOnLine*, vol. 55, no. 25, pp. 109–114, 2022.
- [95] N. Hashemi, J. Ruths, and M. Fazlyab, "Certifying incremental quadratic constraints for neural networks via convex optimization," in *Learning for Dynamics and Control*, pp. 842–853, PMLR, 2021.
- [96] N. Junnarkar, H. Yin, F. Gu, M. Arcak, and P. Seiler, "Synthesis of stabilizing recurrent equilibrium network controllers," in *2022 IEEE 61st Conference on Decision and Control (CDC)*, pp. 7449–7454, IEEE, 2022.
- [97] M. Fazlyab, M. Morari, and G. J. Pappas, "Safety verification and robustness analysis of neural networks via quadratic constraints and semidefinite programming," *IEEE Transactions on Automatic Control*, vol. 67, no. 1, pp. 1–15, 2020.
- [98] M. Fazlyab, M. Morari, and G. J. Pappas, "An introduction to neural network analysis via semidefinite programming," in *2021 60th IEEE Conference on Decision and Control (CDC)*, pp. 6341–6350, IEEE, 2021.
- [99] P. Pauli, N. Funcke, D. Gramlich, M. A. Msalmi, and F. Allgöwer, "Neural network training under semidefinite constraints," in *2022 IEEE 61st Conference on Decision and Control (CDC)*, pp. 2731–2736, IEEE, 2022.
- [100] S. R. Dubey, S. K. Singh, and B. B. Chaudhuri, "Activation functions in deep learning: A comprehensive survey and benchmark," *Neurocomputing*, 2022.
- [101] R. Drummond, M. C. Turner, and S. R. Duncan, "Reduced-order neural network synthesis with robustness guarantees," *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- [102] J. C. Willems, "Dissipative dynamical systems part i: General theory," *Archive for rational mechanics and analysis*, vol. 45, no. 5, pp. 321–351, 1972.
- [103] R. Wrede and M. Spiegel, *Schaum's Outline of Advanced Calculus*. McGraw-Hill Education, 2010.
- [104] R. Drummond, C. Guiver, and M. C. Turner, "Aizerman conjectures for a class of multivariate positive systems," *IEEE Transactions on Automatic Control*, 2022.

- [105] M. C. Turner, M. L. Kerr, and J. Sofrony, "Tractable stability analysis for systems containing repeated scalar slope-restricted nonlinearities," *International Journal of Robust Nonlinear Control*, vol. 25, no. 7, pp. 971–986, 2015.
- [106] M. C. Turner and R. Drummond, "Analysis of MIMO Lurie systems with slope restricted nonlinearities using concepts of external positivity," in *2019 IEEE 58th Conference on Decision and Control (CDC)*, pp. 163–168, 2019.
- [107] P. Gahinet and A. Nemirovski, "The projective method for solving linear matrix inequalities," *Math. programming*, vol. 77, pp. 163–190, 1997.
- [108] P. Park and S. W. Kim, "A revisited Tsympkin criterion for discrete-time nonlinear Lur'e systems with monotonic sector-restrictions," *Automatica*, vol. 34, no. 11, pp. 1417–1420, 1998.
- [109] V. Kapila and W. M. Haddad, "A multivariable extension of the Tsympkin Criterion using a Lyapunov function approach," *IEEE Transactions on Automatic Control*, vol. 41, no. 1, pp. 149–152, 1996.
- [110] N. S. Ahmad, J. Carrasco, and W. P. Heath, "A less conservative LMI condition for stability of discrete-time systems with slope-restricted nonlinearities," *IEEE Transactions on Automatic Control*, vol. 60, no. 6, pp. 1692–1697, 2014.
- [111] H. Wang, Z. Xiong, L. Zhao, and A. Papachristodoulou, "Model-free verification for neural network controlled systems," *arXiv preprint arXiv:2312.08293*, 2023.
- [112] M. Newton and A. Papachristodoulou, "Sparse polynomial optimisation for neural network verification," *Automatica*, vol. 157, p. 111233, 2023.
- [113] H. M. Hedesh and M. Siami, "Ensuring both positivity and stability using sector-bounded nonlinearity for systems with neural network controllers," *arXiv preprint arXiv:2406.12744*, 2024.
- [114] S. V. Noori, B. Hu, G. Dullerud, and P. Seiler, "A complete set of quadratic constraints for repeated ReLU," *arXiv preprint arXiv:2407.06888*, 2024.
- [115] S. V. Noori, B. Hu, G. Dullerud, and P. Seiler, "Stability and performance analysis of discrete-time ReLU recurrent neural networks," *arXiv preprint arXiv:2405.05236*, 2024.
- [116] T. Yuno, K. Fukuchi, and Y. Ebihara, "A Lyapunov-based method of reducing activation functions of recurrent neural networks for stability analysis," *IEEE Control Systems Letters*, 2024.
- [117] P. Pauli, A. Havens, A. Araujo, S. Garg, F. Khorrami, F. Allgöwer, and B. Hu, "Novel quadratic constraints for extending lipsdp beyond slope-restricted activations," *arXiv preprint arXiv:2401.14033*, 2024.

- [118] L. B. Groff, J. M. Gomes da Silva, and G. Valmorbida, "Regional stability of discrete-time linear systems subject to asymmetric input saturation," in *IEEE Conference on Decision and Control (CDC)*, pp. 169–174, 2019.
- [119] L. Cabral, G. Valmorbida, and J. M. Gomes da Silva, "Exponential stability of continuous-time piecewise affine systems," *IEEE Control Systems Letters*, vol. 8, pp. 1649–1654, 2024.
- [120] S. J. Linz and J. Sprott, "Elementary chaotic flow," *Physics Letters A*, vol. 259, no. 3-4, pp. 240–245, 1999.
- [121] B. D. Anderson, "The small-gain theorem, the passivity theorem and their equivalence," *Journal of the Franklin Institute*, vol. 293, no. 2, pp. 105–115, 1972.
- [122] M. Fazlyab, T. Entesari, A. Roy, and R. Chellappa, "Certified robustness via dynamic margin maximization and improved lipschitz regularization," *Advances in Neural Information Processing Systems*, vol. 36, pp. 34451–34464, 2023.
- [123] E. R. Kandel, J. H. Schwartz, T. M. Jessell, S. Siegelbaum, A. J. Hudspeth, S. Mack, et al., *Principles of neural science*, vol. 4. McGraw-hill New York, 2000.
- [124] M. Khona and I. R. Fiete, "Attractor and integrator networks in the brain," *Nature Reviews Neuroscience*, vol. 23, no. 12, pp. 744–766, 2022.
- [125] R. Vogt, M. Puelma Touzel, E. Shlizerman, and G. Lajoie, "On Lyapunov exponents for RNNs: Understanding information propagation using dynamical systems tools," *Frontiers in Applied Mathematics and Statistics*, vol. 8, p. 818799, 2022.
- [126] V. Centorrino, F. Bullo, and G. Russo, "Contraction analysis of Hopfield neural networks with Hebbian learning," in *2022 IEEE 61st Conference on Decision and Control (CDC)*, pp. 622–627, IEEE, 2022.
- [127] J. J. Hopfield, "Neurons with graded response have collective computational properties like those of two-state neurons.," *Proceedings of the national academy of sciences*, vol. 81, no. 10, pp. 3088–3092, 1984.
- [128] L. Kozachkov, J. Tauber, M. Lundqvist, S. L. Brincat, J.-J. Slotine, and E. K. Miller, "Robust and brain-like working memory through short-term synaptic plasticity," *PLOS Computational Biology*, vol. 18, no. 12, p. e1010776, 2022.
- [129] M. Pals, J. H. Macke, and O. Barak, "Trained recurrent neural networks develop phase-locked limit cycles in a working memory task," *PLOS Computational Biology*, vol. 20, no. 2, p. e1011852, 2024.
- [130] D. Krotov and J. Hopfield, "Large associative memory problem in neurobiology and machine learning," *arXiv preprint arXiv:2008.06996*, 2020.

- [131] S. Sharma, S. Chandra, and I. Fiete, "Content addressable memory without catastrophic forgetting by heteroassociation with a fixed scaffold," in *International Conference on Machine Learning*, pp. 19658–19682, PMLR, 2022.
- [132] L. Kozachkov, J.-J. Slotine, and D. Krotov, "Neuron-astrocyte associative memory," *arXiv preprint arXiv:2311.08135*, 2023.
- [133] R. Ofir, J.-J. Slotine, and M. Margaliot, " k -contraction in a generalized Lurie system," *arXiv preprint arXiv:2309.07514*, 2023.
- [134] R. Ofir, A. Ovseevich, and M. Margaliot, "Contraction and k -contraction in Lurie systems with applications to networked systems," *Automatica*, vol. 159, p. 111341, 2024.
- [135] X. Zhang and B. Cui, "Synchronization of Lurie system based on contraction analysis," *Applied Mathematics and Computation*, vol. 223, pp. 180–190, 2013.
- [136] J. S. Muldowney, "Compound matrices and ordinary differential equations," *The Rocky Mountain Journal of Mathematics*, pp. 857–872, 1990.
- [137] A. Cecilia, S. Zoboli, D. Astolfi, U. Serres, and V. Andrieu, "Generalized Lyapunov conditions for k -contraction: analysis and feedback design," 2023.
- [138] S. Zoboli, A. Cecilia, and S. Tarbouriech, "Quadratic abstractions for k -contraction," 2024.
- [139] A. Gu, T. Dao, S. Ermon, A. Rudra, and C. Ré, "Hippo: Recurrent memory with optimal polynomial projections," *Advances in neural information processing systems*, vol. 33, pp. 1474–1487, 2020.
- [140] W. L. Hamilton, R. Ying, and J. Leskovec, "Representation learning on graphs: Methods and applications," *arXiv preprint arXiv:1709.05584*, 2017.
- [141] K. B. Petersen, M. S. Pedersen, *et al.*, "The matrix cookbook," *Technical University of Denmark*, vol. 7, no. 15, p. 510, 2008.
- [142] R. A. Horn and C. R. Johnson, *Topics in matrix analysis*. Cambridge university press, 1994.
- [143] M. Lezcano-Casado and D. Martínez-Rubio, "Cheap orthogonal constraints in neural networks: A simple parametrization of the orthogonal and unitary group," in *International Conference on Machine Learning*, pp. 3794–3803, PMLR, 2019.
- [144] A. Lemme, Y. Meirovitch, M. Khansari-Zadeh, T. Flash, A. Billard, and J. J. Steil, "Open-source benchmarking for learned reaching motion generation in robotics," *Paladyn, Journal of Behavioral Robotics*, vol. 6, no. 1, 2015.

- [145] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms," *arXiv preprint arXiv:1708.07747*, 2017.
- [146] R. Marino, L. Buffoni, L. Chicchi, F. Di Patti, D. Febbe, L. Giambagli, and D. Fanelli, "Learning in Wilson-Cowan model for metapopulation," *arXiv preprint arXiv:2406.16453*, 2024.
- [147] R. Pascanu, T. Mikolov, and Y. Bengio, "On the difficulty of training recurrent neural networks," in *International conference on machine learning*, pp. 1310–1318, Pmlr, 2013.
- [148] R. Thomas, "Deterministic chaos seen in terms of feedback circuits: Analysis, synthesis, labyrinth chaos," *International Journal of Bifurcation and Chaos*, vol. 9, no. 10, pp. 1889–1905, 1999.
- [149] L. El Ghaoui, F. Gu, B. Travacca, A. Askari, and A. Tsai, "Implicit deep learning," *SIAM Journal on Mathematics of Data Science*, vol. 3, no. 3, pp. 930–958, 2021.
- [150] M. Fazlyab, A. Robey, H. Hassani, M. Morari, and G. Pappas, "Efficient and accurate estimation of lipschitz constants for deep neural networks," *Advances in neural information processing systems*, vol. 32, 2019.
- [151] M. Revay, R. Wang, and I. R. Manchester, "Lipschitz bounded equilibrium networks," *arXiv preprint arXiv:2010.01732*, 2020.
- [152] H. Ramsauer, B. Schöfl, J. Lehner, P. Seidl, M. Widrich, T. Adler, L. Gruber, M. Holzleitner, M. Pavlović, G. K. Sandve, *et al.*, "Hopfield networks is all you need," *arXiv preprint arXiv:2008.02217*, 2020.
- [153] P. Kidger, "On neural differential equations," *arXiv preprint arXiv:2202.02435*, 2022.
- [154] S. Zoboli and A. Cecilia, "k-contraction analysis for discrete-time systems," in *4th IFAC Conference of Modelling, Identification and Control of Nonlinear Systems*, 2024.

Appendix A

Theorem A.1 (Theorem 1 [134]). *Consider the Lurie network (6.2) with $b_x = b_y = 0$. Fix $k \in [1, n]_{\mathbb{N}}$. Suppose that there exists $\eta_1, \eta_2 \in \mathbb{R}$ and $P \in \mathbb{R}^{n \times n}$, where $P = \Theta \Theta$ with $\Theta \succ 0$, such that*

$$P^{(k)} A^{[k]} + (A^{[k]})^\top P^{(k)} + \eta_1 P^{(k)} + \Theta^{(k)} \left((\Theta B B^\top \Theta)^{[k]} + (\Theta^{-1} C^\top C \Theta^{-1})^{[k]} \right) \Theta^{(k)} \preceq 0 \quad (\text{A.1})$$

and, furthermore, at least one of the following two conditions hold:

$$\sum_{i=1}^k \lambda_i \left(\Theta^{-1} C^\top (J_\Phi^\top(y) J_\Phi(y) - I) C \Theta^{-1} \right) \leq -\eta_2 \quad (\text{A.2})$$

or

$$\sum_{i=1}^k \lambda_i \left(\Theta B (J_\Phi(y) J_\Phi^\top(y) - I) B^\top \Theta \right) \leq -\eta_2 \quad (\text{A.3})$$

for all $y \in \mathbb{R}^m$. Then the Jacobian of the Lurie network satisfies

$$\mu_{2, \Theta^{[k]}}(J_{\dot{x}}^{[k]}(x)) \leq -\frac{1}{2}(\eta_1 + \eta_2) \quad (\text{A.4})$$

for all $x \in \mathbb{R}^n$. In particular, if $\eta_1 + \eta_2 > 0$, then the Lurie network (6.2) with $b_x = b_y = 0$ is k -contractive with rate $(\eta_1 + \eta_2)/2$ with respect to the norm $\|z\|_{2, \Theta^{(k)}} = \|\Theta^{(k)} z\|_2$.