

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: Andrei C. Rusu (2024) "An inquiry into diffusion processes over interaction networks", University of Southampton, School of Electronics and Computer Science, PhD Thesis, pagination.

Data: Andrei C. Rusu (2021) Epidemic simulation data for the PLoS One study "Modelling digital and manual contact tracing for COVID-19. Are low uptakes and missed contacts deal-breakers?". <https://doi.org/10.6084/m9.figshare.14101946>. [dataset]

UNIVERSITY OF SOUTHAMPTON

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

**An inquiry into diffusion processes over
interaction networks**

DOI: 10.5258/SOTON/T0065

by

Andrei C. Rusu

[BSc Computer Science]

ORCID: 0000-0002-6053-1685

*A thesis for the degree of
Doctor of Philosophy*

11th January 2024

University of Southampton

Abstract

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

Doctor of Philosophy

An inquiry into diffusion processes over interaction networks

by Andrei C. Rusu

This thesis aims to develop a comprehensive framework for modelling and controlling diffusion processes over interaction networks, striving to inform and improve public health policies against viral epidemics. Our work introduces four main contributions: (1) a new modelling technique that captures the heterogeneity and uncertainty of contact patterns and evaluates the impact of different testing and tracing strategies, which can be utilized in conjunction with any compartmental formulation to study complex spreading dynamics. Using this technique, we introduce and simulate a novel epidemiological model, SEIR-T, showing that contact tracing in a COVID-19 epidemic can be effective despite suboptimal digital uptakes or pervasive interview inefficiencies; (2) a versatile and cost-effective approach to optimizing the allocation of testing, tracing and vaccination resources based on the network structure and epidemic dynamics, which ranks individuals based on their role in the network and the epidemic state, being adaptable to the budget and risk preferences of regional policy makers, while still breaking high-risk transmission chains; (3) a reinforcement learning-based agent, underpinned by a highly transferable graph neural architecture, that can find optimal epidemic control policies from simulation data, outperforming standard heuristic approaches by up to 15% in the containment rate, while far surpassing more standard random samplers by margins of 50% or more; and (4) a range of visualization tools that can aid in understanding and communicating the effects of public health interventions to policy makers and the populace, which include prediction explanation and state visualization techniques for scrutinizing the learning-based policies introduced, and other tools the authorities can use to assess the cost-benefit trade-off of enacting different combinations of interventions. The simulation-control framework we introduce is particularly flexible and can effectually model the spread of various pathogens or analogous diffusion processes, such as information dissemination. Similarly, the learned epidemic policies are versatile and easily transferable to a wide range of diffusion scenarios and network structures.

Contents

List of Figures	ix
List of Tables	xv
Declaration of Authorship	xvii
Acknowledgements	xix
Definitions and Abbreviations	xxiii
1 Introduction	1
1.1 Diffusion processes: Overview	1
1.2 Findings and contributions	5
1.3 Modelling the diffusion in pathogen epidemics	7
1.4 Going beyond pathogen diffusion	11
1.4.1 Information diffusion problem	12
1.4.2 Modelling the diffusion of information	13
2 Graphs: Generation and Modelling	17
2.1 Introduction	17
2.2 Notations	18
2.3 Network generation	19
2.4 Network heuristics	22
2.5 Machine learning with network data	24
2.6 Kernel methods and graph kernels	24
2.7 Graph neural networks	26
2.8 Explaining graph neural networks	30
2.9 Applications of graph-based ML to diffusion processes	31
2.9.1 Identifying patient zero	31
2.9.2 Controlling the diffusion dynamics with GNNs and RL	32
2.9.3 Further related work in graph-based ML for diffusion processes	34
3 Modelling Contact Tracing	37
3.1 Problem overview	37
3.2 Related work	39
3.3 Methodology	40
3.3.1 Compartmental model outline	40
3.3.2 Network propagation mechanism	40

3.3.3	Simulation overview	44
3.3.4	Metrics under consideration	47
3.4	Results and discussion	47
3.4.1	Variation induced by population size	47
3.4.2	Initial exploration of tracing overlap in larger populations	49
3.4.3	Hospitalizations and deaths when improving tracing overlaps	52
3.4.4	Effects of average degree and app uptake	54
3.4.5	Combining digital tracing with an imperfect manual process	56
3.4.6	Contact tracing efficiency in a real social network	60
3.5	Case study conclusion	66
4	<i>EpiCURB</i>: Epidemic Control Using RL for Budget allocation	69
4.1	Introduction	69
4.2	Related work and Background	71
4.2.1	Epidemic modelling	71
4.2.2	Graph neural networks	71
4.2.3	Reinforcement learning	72
4.2.4	Combinatorial action spaces	74
4.2.5	Influencing graph dynamics	75
4.3	Methodology	76
4.3.1	Simulating epidemics	76
4.3.2	Network configurations	77
4.3.3	Control setup	78
4.3.4	Baseline agents	79
4.3.5	Learning agents	80
4.3.5.1	Overview	80
4.3.5.2	Reinforcement learning routines	82
4.3.6	Simulation-control framework pseudocode	84
4.4	Results and discussion	88
4.4.1	Initial comparison of learning-based variants	88
4.4.2	Prioritizing testing in static graphs	88
4.4.3	Prioritizing testing in dynamic graphs	90
4.4.4	Testing performance analysis	91
4.4.5	Targeted test and trace programmes	93
4.4.6	Agents interacting with different spreading dynamics	95
4.5	Targeting vaccination campaigns	96
4.6	Execution time analysis for test prioritization	97
4.7	Conclusion and future work	98
5	Conclusions and Future Perspectives	101
5.1	Summary of principal contributions	101
5.2	Broader impact	103
5.2.1	Impact on pathogen epidemiology	103
5.2.2	Impact on information diffusion	104
5.3	Limitations and future perspectives	105
5.3.1	Future work needed to address current limitations	105

5.3.2	Future directions inspired by preliminary findings and related works	108
Appendix A Additional Review Material on Graph-based Machine Learning		111
Appendix A.1	Graph-level kernels	111
Appendix A.2	Alternative approaches for building GNNs	112
Appendix A.3	More graph-based ML applications to diffusion processes . .	114
Appendix A.3.1	Modelling the region-to-region spread of COVID-19 .	114
Appendix A.3.2	Temporal graph kernels for diffusion processes	116
Appendix B Modelling Contact Tracing: Supporting Information		119
Appendix B.1	Open-source model and data	119
Appendix B.2	Simulation statistics	119
Appendix B.3	Further analysis of noteworthy trends	120
Appendix C Informing Policy using <i>EpiCURB</i>		123
Appendix C.1	Introduction	123
Appendix C.2	Policy choices	124
Appendix C.3	Methodology	128
Appendix C.4	Cost-based balancing of public health interventions	128
Appendix C.5	Inspecting the latent space of learning-based agents	131
Appendix C.6	Explaining the decisions taken by the GNN ranker	132
Appendix C.7	Conclusion and future work	134
Appendix D <i>EpiCURB</i>: Supporting Material		135
Appendix D.1	Open-source model	135
Appendix D.2	Markov chain Monte Carlo epidemic model fitting	135
Appendix E Tools and Preliminary Experiments		139
Appendix E.1	Tools	139
Appendix E.2	Employed or explored network datasets	139
Appendix E.3	Preliminary benchmarking of graph-based ML methods . . .	140
Bibliography		143

List of Figures

- 1.1 **The SEIZ information diffusion model.** Diagram of possible transitions is given on the left, while the characteristic ODE equations are presented on the right. The possible system states are: S (Susceptible), E (exposed), I (infectious), and Z (skeptical). Source: [Jin et al. \(2013\)](#). 15
- 2.1 **Graph neural networks intuition.** Every node in the input graph determines a separate compute graph. The squares represent neural networks with shared weights W_l and biases towards the current representation B_l per each layer l . In this context, layers refer to hops within the network, i.e. for each vertex v , layer 0 are its node features x_v , layer 1 aggregates information from its immediate neighborhood (features from adjacent nodes and their common edges) according to W_1 and B_1 etc. Source: [Leskovec \(2021\)](#). 27
- 2.2 **Visualization of the RL+GNN approach.** In this toy example, the ranking module receives as input a sequence G_1, \dots, G_5 , constructs a temporal multi-graph from this, and feeds it to the actor ranking module. The latter outputs a probability distribution over actions, which ultimately gets sampled by the actor. The blue node is chosen in this scenario and then gets isolated, the resulting new state being aggregated to the original input to obtain the next timestamp $t = 6$. At the same time, the ranking module's parameters are updated according to a PPO loss involving the reward of this new state (based on the RL criterion above) and the critic's feedback. Source: [Meirom et al. \(2021\)](#). 33
- 3.1 **The SEIR-T Compartmental Model for COVID-19.** Each node in the interaction network has 2 allocated variables: an infection state and a tracing status. The infection states from top to bottom are: S - susceptible; E - exposed but not infectious; I_p - infectious, presymptomatic; I_a - infectious, asymptomatic; I_s - infectious, symptomatic; H - hospitalized; R - recovered/removed; D - dead. At any point in time, a node's tracing status can either be T (traced and isolated) or N (not traced/isolated or non-compliant). Each state transition has a certain time-dependent probability $p_{s_1 \rightarrow s_2}$; the edge labels here represent both $\frac{p_{s_1 \rightarrow s_2}}{\Delta t}$, and the λ rate of the corresponding exponential to sample from in the continuous-time simulations. 41
- 3.2 **Final state of an epidemic simulation over a dual topology.** Infection spreads with respect to the neighborhoods of the first network (here a SF graph); the second network represents a digital tracing view at uptake $r = 0.5$ 41

- 3.3 **Final state of an epidemic simulation over a *triad* topology.** Infection spreads with respect to the neighborhoods of the first network (here a SW graph); the second network represents a digital tracing view at uptake $r = 0.5$, while the third a manual tracing view with overlap $\Gamma = 0.5$ 42
- 3.4 **Uncertainty of simulation results w.r.t. the infection peak.** Averages from 80-100 runs involving different population sizes, $K = 10$, $\tau_t = \tau_r = 0.1$, $p_a = 0.2$. On top, boxplots with quartiles represented via whiskers, medians via orange lines, and averages via green triangles; standard deviations given below. The left-hand side shows absolute values, whereas on the right all variation levels are scaled down by N 48
- 3.5 **ER network – Peak suppression (left) and the time of peak (right) at various tracing network overlaps.** Values are averaged over 105 runs, representing results for $N = 10000$, $K = 10$, $p_a = 0.2$. The suppression is calculated by subtracting the average maximal infected point given by each parameter configuration from the average point obtained with no tracing ($\tau_t = 0$, colored in black). Excepting $\tau_r = 0.001$ and $\tau_t = 0.01$, which produce inconclusive and noisy results, the effectiveness of an epidemic containment strategy expectedly scales with the testing and the tracing rates. 50
- 3.6 **ER network – Epidemic evolution over time given a less efficient (left) and a more effective (right) testing regimes.** Results averaged over 105 simulations, obtained for $N = 10000$, $K = 10$, $p_a = 0.2$. As the contact tracing rate increases, the accuracy of the network given by Γ becomes more important for ‘flattening’ the curves. The case with no contact tracing ($\tau_t = 0$) is colored in black. 51
- 3.7 **ER network – Heatmaps of achieved peak suppression for various testing and tracing rates.** $N = 10000$, $K = 10$, $p_a = 0.2$. Averaged over 105 runs. 52
- 3.8 **SF network – Hospitalizations and total deaths over time.** Setting $\tau_r = 0.1$, $N = 10000$ and $K = 10$. Each row represents a different community: adults only, seniors only, and mixed (adjusted rates based on UK demographics). The 95% confidence intervals resulted from 105 runs are displayed around each line. The case with no contact tracing ($\tau_t = 0$) is colored in black. 53
- 3.9 **ER network – Uptake rate r against peak suppression.** Suppression is difference in peak to no tracing, i.e. $\tau_t = 0$. $N = 1000$, $p_a = 0.2$, $\eta = .001$. $K = 10$ given on the left, $K = 20$ on the right. The case with no tracing ($\tau_t = 0$) is colored in black. Lines are plotted with the 95% confidence intervals from 750 runs. 54
- 3.10 **ER network – Epidemic evolution over time for $\tau_r = 0.1$ $N = 1000$ and $K = 20$.** Results averaged over 750 runs. No tracing ($\tau_t = 0$) is given in black. 55

- 3.11 **HK network – Uptake rate r against peak suppression.** The results here correspond to networks generated with parameters $N = 1000$, $m_{HK} = 10$, $p_{\Delta} = 0.2$, $p_a = 0.2$, $\eta = .001$. On the left, we present a scenario in which only digital tracing is conducted, whereas the next 3 columns represent simulations with a combination of digital tracing on a second network, and manual tracing over a third network with various overlaps: 0.1, 0.55, 1. The 95% confidence intervals are also displayed. The case with no tracing ($\tau_t = 0$) is colored in black. 57
- 3.12 **HK network – Uptake rate r against the effective reproduction number.** Results here correspond to networks generated with $N = 1000$, $m_{HK} = 10$, $p_{\Delta} = 0.2$, $p_a = 0.2$, $\eta = .001$. On the left, we present scenarios in which only digital tracing is conducted, whereas the next columns represent simulations with a combination of digital tracing on a second network, and manual tracing over a third network with various overlaps: 0.1, 0.55, 1. The case with no tracing ($\tau_t = 0$) is colored in black. 58
- 3.13 **HK network – Contour plots of R_e based on the level of manual tracing overlap Γ and digital tracing uptake r .** The results here are for a networks generated with parameters $N = 1000$, $m_{HK} = 10$, $p_{\Delta} = 0.2$, $p_a = 0.2$, $\eta = .001$. Each line represents a different testing level τ_r , while the columns correspond to a moderate (left) and efficient (right) level of tracing effort and isolation compliance given by τ_t 59
- 3.14 **Social Evolution – Uptake rate r against peak suppression.** The results here correspond to the real Social Evolution network, dynamic over the studied period of 31 weeks, $p_a = 0.2$, $\eta = .001$. On the left, we have a scenario in which only digital tracing is conducted, whereas the next 3 columns represent simulations with a combination of digital tracing on a second network, and manual tracing over a third network with various overlaps: 0.1, 0.55, 1. The 95% confidence intervals are displayed. No tracing case is colored in black. 61
- 3.15 **Social Evolution – Uptake rate r against the effective reproduction number.** The results here correspond to the data-driven Social Evolution network, $p_a = 0.2$, $\eta = .001$. On the left, we have a scenario in which only digital tracing is conducted, whereas the next columns represent simulations with a combination of digital tracing on a second network, and manual tracing over a third network with various overlaps: 0.1, 0.55, 1. No tracing case is colored in black. 62
- 3.16 **Social Evolution – Contour plots of R_e based on the level of manual tracing overlap Γ and digital tracing uptake r .** The results here correspond to the real Social Evolution network, dynamic over the studied period of 31 weeks, $p_a = 0.2$, $\eta = .001$. Each line represents a different testing rate τ_r , while the columns showcase a less efficient (left), a moderate (center-left), an efficient (center-right) and a very efficient (right) level of tracing effort and isolation compliance τ_t 63
- 3.17 **Social Evolution – Uptake rate r against peak suppression, for different p_a values.** The results here correspond to the real Social Evolution network, dynamic over the studied period of 31 weeks, $\eta = .001$, and either $p_a = 0.2$ (on the left of each pair) or $p_a = 0.5$ (on the right of each pair). The left-quadrant pairs represent a triad network scenario with manual overlap $\Gamma = 0.1$, while the right quadrant shows $\Gamma = 0.5$. 95% confidence intervals are displayed. No tracing case is colored in black. . 64

3.18 **Social Evolution – Contour plots of R_e based on the level of manual tracing overlap Γ and digital tracing uptake r , for different p_a values.** The results here correspond to the data-driven Social Evolution network when $\eta = .001$, with each pair of charts describing $p_a = 0.2$ on the left and $p_a = 0.5$ on the right. Each line represents a different testing level τ_r , while the columns showcase combinations of one p_a value together with a level of tracing effort given by τ_t 65

4.1 **Block diagram of our simulation-control framework.** The *Agent* is passed as a parameter to the *Simulator* instance, which is responsible for delegating calls to the *control(.)* method of the *Agent* once enough time has passed and other preset conditions are met (e.g. minimum infection threshold). The *control(.)* function performs preliminary node filterings, then calls *control_test(.)*, *control_trace(.)* and *control_vax(.)*. These methods are responsible for producing node rankings for testing, tracing and vaccinating, respectively, exhibiting a different logic in each *Agent* subtype. Every *MeasureAgent* decides the final ranking via a node-level score. Both the *SLAgent* and *RLAgent* feature pluggable and trainable weights, sharing an identical evaluation mode. Combinations of agent behaviors for each of the three tasks can be selected by utilizing a *MixAgent*. Further details about each type of agent can be consulted in Sections 4.3.4 and 4.3.5. 78

4.2 **Neural architecture for the learning-based control agents.** At each time increment t , the graph node features $x_t = f_v(t)$ and connections get passed to the Information and Diffusion modules, producing the outputs i_t and d_t , respectively. The Information component has 3 GIN layers and a final FA layer, interconnected by Dropout, ReLU and LayerNorm. The Diffusion module has a single GATv2 layer, followed by Dropout, ReLU and LayerNorm. i_t , d_t , and the node hidden states from the previous timestamp h_{t-1} are utilized as inputs for the MLP that outputs the new standard-scaled hidden states h_t . By using x_t , h_t , and h_{t-1} , two parallel MLPs produce the per-node action-values of the actor $m_t = Q(s_t, a_t)$, and the state value of the critic $v_t = V(s_t)$. The latter is obtained by applying a *MaxPool* per rows before passing the inputs to the afferent MLP. The *RLAgent* uses the m_t output to rank and sample nodes, and the v_t output to receive feedback, both being key components of the PPO objective. The *SLAgent* feeds m_t directly into a cross-entropy loss, using the infection indicators $\mathbb{I}_v(t)$ as labels. MLP sketch source: <https://www.gabormelli.com/RKB/File:2NNw.png>. 81

4.3 **Number of infected during training (left) and fraction kept healthy during evaluation (right) of learning- and heuristic-based agents.** This is reported for dual Barabási-Albert networks with $N = 1000$ and mean degree ≈ 3 . Test budget set to $k = 2$. The left panel shows results for a single network for consistency, while the right panel provides boxplots from 7 different networks and 7 infection seeds, marking the Q1, Q2 and Q3 quartiles, the means, and the extremity points without outliers from these runs. All agents have recollection of negative test results. Since the training behavior of the SL and RL agents tends to fluctuate across episodes and configurations, we plot the average curves and standard deviations from multiple seeded runs. 88

4.4	Infection control performance on different static network architectures and sizes, with a budget of $k = 2$. Uncertainties are shown as boxplots.	89
4.5	Infection control performance on different static network architectures with varying budgets. The uncertainties are shown as boxplots.	90
4.6	Infection control performance on dynamic network architectures. Suffix ‘nor’ shows variants with no recollection.	91
4.7	Visualization of the spread for the degree w/ R and the RL agents. This corresponds to a stochastic-block network Holland et al. (1983) with three communities. Susceptibles are green (light), infectious red (darker), and detected blue (the darkest). Initially, the two policies target the same nodes, but later on the RL agent preferentially tests the bridges.	92
4.8	Detection-infection time lag histograms. Results are summed up across 49 simulated epidemics, and displayed per each recovery rate γ	93
4.9	Cumulative detection within X days since exposure. Second row displays results normalized by the average total number of detected.	94
4.10	Averaged epidemic curves and standard deviations during test and trace control. These are for 5000 nodes dual Barabási-Albert graphs featuring a mean degree of approximately 3, with a daily testing budget of $k = 1\%$ and no tracing on the left, and $k = 10$ with a limit of $k_c = 25$ traced contacts on the right. Results displayed for two RL agents: one trained for 50, and the other for 200 episodes.	94
4.11	Infection control performance when prioritizing vaccination. First 7 agents utilize k_c , while the rest use k_v . The total budget of vaccines is capped at 100.	97
4.12	Attack rate when varying the vaccination fraction. The attack rate is the percentage of nodes getting infected from a given population, representing the complement of the fraction kept healthy analyzed in the rest of this chapter. The differences between the heuristic policies become smaller when testing with an RL agent (second column), as opposed to conducting no testing (first column).	97
Appendix A.1 Illustration of different discretization schemes for Eq A.6.		
	From left to right, the schemes are explicit Euler, 4th order Runge-Kutta, and implicit Euler. Source: Chamberlain et al. (2021)	113
Appendix A.2 Spatio-temporal input graph. White node from Day 10 gets an aggregated representation from its neighbors at the current timestamp, as well as its images from previous days (i.e. red-filled vertices). Source: Kapoor et al. (2020)		
		115
Appendix B.1 Tracing false positives. The amount of susceptibles being incorrectly traced and isolated. Results here correspond to the last experiment, with $N = 1000$, Holme-Kim graph topology and 10% initial infected.		
		120
Appendix B.2 Tracing false negatives. Number of <i>infectious</i> people not traced. Results correspond to a Holme-Kim graph topology with $N = 1000$ and 10% initial infected.		
		121
Appendix B.3 The overlap influencing the efficacy of contact tracing rates. Results for $N = 10000$, random graph topology with mean degree $K = 10$. τ_r fixed at 0.04.		
		121

Appendix B.4	Uptake influencing the efficacy of contact tracing rates. Results for $N = 1000$, random graph topology with mean degree $K = 10$. τ_r fixed at 0.05.	122
Appendix B.5	Reproduction number R_e vs. uptake r vs. overlap Γ. Results correspond to a Holme-Kim graph topology with $N = 1000$ and 10% initial infected.	122
Appendix C.1	Contour plots of the containment levels (i.e. percentage of population kept healthy) achieved by combinations of different public health interventions. This figure illustrates how policy makers can assess and allocate resources to each measure, according to the desired level of containment. The plots on the first row have marked the reference points mentioned in text, with arrows delineating the space of effort-increase actions needed to reach the next infection threshold. The Digital-Manual tracing plot also highlights the most important region of the space, where uptakes are between 20-50%, overlaps are between 10-60%, and 10-15% increases in each direction leads to $\approx 8\%$ less infections. Contours are obtained by averaging 50-100 simulations, and applying a Gaussian filter to remove remnant stochastic noise.	129
Appendix C.2	t-SNE plots of the node hidden states h_t and dendrograms of their hierarchical clustering into 7 groups. The first column shows a GNN ranker after training, while the second column shows the same model with initial parameter values. We plot multiple day intervals for a network of $N = 700$ nodes, fixing the daily testing budget to $k = 5$. The latent space of the trained ranker is informative, often grouping the new infected (red) close to the recent negatives (green), while pushing most detected individuals (blue) to specific regions. In contrast, the untrained model is unable to cluster the nodes in a meaningful manner, with infections constantly soaring.	131
Appendix C.3	Feature importances β_{ip} in the earlier stages. (a) Initially, the agent does not possess enough information about the epidemic state, and therefore mostly focuses on the centrality features. (b) After a while, the agent starts attending to the epidemic state features as well, such as the previously untested or the positive flags, and the number of infected neighbors. Top row displays each node's feature values, while neighborhood averages are shown underneath.	133
Appendix C.4	Feature importances β_{ip} in the final stages. (a) The epidemic state features continue to be important predictors of a node's rank in the final segment of the simulation, while the degree centrality remains the top predictor. (b) In the final days, both the degree and the eigenvector centrality become again the most important features for ranking vertices, with the epidemic state information either smoothing out or becoming highly correlated with the latter.	133
Appendix D.1	Mortality curve fitting using MCMC. Data represents COVID-19 deaths recorded in 4 US states. Simulations performed via the IBMF model of Section 4.3.1.	136

List of Tables

3.1	Compartmental model parameters used in our simulations.	42
3.2	Network-generation parameters.	43
4.1	Legend for the control framework pseudocode.	85
4.2	Fraction kept healthy with budget $k = 1\%$ and different recovery rates. Average over 5 seeded runs for each of the considered 5 realizations of dual Barabási-Albert networks with $N = 1000$ nodes and a mean degree of approximately 3. ‘w/ R’ denotes agents with recollection of recent negative test results.	89
4.3	Fraction kept healthy with budget $k = 1\%$ and different population sizes. Average over 5 seeded runs for each of the considered 5 realizations of dual Barabási-Albert networks with a mean degree of approximately 3. ‘w/ R’ denotes agents with recollection of recent negative test results. Here, a single model is trained for 50 episodes on a network of size 1000, but its policy is able to generalize to appreciably larger graphs.	90
4.4	Fraction kept healthy for a 1000-node network built from real tracing statistics. Results are averaged over 5 runs for each of the 5 realizations of a configuration model using these statistics.	91
4.5	Percentage of nodes kept healthy for graphs of size 1000 and an approximate mean degree of 3 (or $m_{HK} = 3$ for Holme-Kim), with budgets set to $k = 2$, $k_c = 5$. Averages over 5 runs for each of the considered 5 realizations of the following: dual Barabási-Albert with $m_1 = 5$, $m_2 = 1$ (BA 5-1) and $m_1 = 10$, $m_2 = 1$ (BA 10-1), Holme-Kim (HK), and Erdős-Rényi (ER).	95
4.6	Percentage of nodes kept healthy when controlling epidemics over a dynamic real interaction network. Graph consists of 74 vertices, and is reconstructed from the Social Evolution dataset Madan et al. (2012). Averages over 5 runs for each of the considered 5 infection seeds. Test budget is $k = 2$	95
4.7	Fraction kept healthy for 2000 nodes and an average degree of approximately 3. Results represent averages over 5 runs for each of the considered 5 instances of a dual Barabási-Albert model ($m_1 = 10$, $m_2 = 1$). Testing budget is $k = 2$ and no contact tracing is conducted.	96
4.8	Fraction kept healthy for 2000 nodes and an average degree of approximately 3. Results represent averages over 5 runs for each of the considered 5 instances of a dual Barabási-Albert model ($m_1 = 10$, $m_2 = 1$). Budgets are $k = 2$ and $k_c = 10$	96
4.9	Average wall clock time per epidemic during evaluation. Configuration: dual Barabási-Albert networks of 2000 nodes, an average degree ≈ 3 , and a daily testing budget of $k = 2$	98

Appendix E.1 **Test accuracies on node classification.** 141

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: [Rusu et al. \(2021\)](#), [Rusu et al. \(2022\)](#), and [Rusu et al. \(2023\)](#).

Signed:.....

Date:.....

Acknowledgements

We would like to thank both Dr Katayoun Farrahi and Prof Mahesan Niranjan for painstakingly supervising and supporting this work. Additionally, we would like to extend our gratitude towards Dr Remi Emonet, co-author of our first paper titled 'Modelling digital and manual contact tracing for COVID-19', for his lucrative ideas and exceptional responsiveness, and towards Dr Eli Meirom, for sharing with us valuable insights about the epidemic control framework that his research group has developed. Finally, we acknowledge the use of the IRIDIS High Performance Computing Facility, and its support services at the University of Southampton, in the completion of this work.

To my beloved Family ...

Definitions and Abbreviations

G	Graph
N	Graph size
θ	Model parameters (vector form)
Θ	Layer parameters (matrix form, $vec(\Theta) \subseteq \theta$)
ER	Erdős–Rényi random network (also the generation model)
SF	Scale-free network (also the generation model)
SW	Small-world network (also the generation model)
BA	Barabási-Albert network (also the generation model)
HK	Holme-Kim network (also the generation model)
ML	Machine Learning
RL	Reinforcement Learning
SVM	Support Vector Machine
NN	Neural Network
MLP	Multi-Layer Perceptron
LSTM	Long-Short Term Memory network
CNN	Convolutional Neural Network
GNN	Graph Neural Network
GCN	Graph Convolutional Network
GAT	Graph Attention Network
GIN	Graph Isomorphism Network
WL	Weisfeiler-Lehman
MDP	Markov Decision Process
POMDP	Partially Observable Markov Decision Process
MC	Monte Carlo
MCMC	Markov Chain Monte Carlo
TD	Temporal Difference
API	Application Programming Interface

Chapter 1

Introduction

1.1 Diffusion processes: Overview

The concept of *diffusion* is paramount to multiple fields of study, including physics, chemistry, biology, sociology and economics, being broadly defined as the movement of a substance, physical quantity or set of entities from a locus of higher concentration (or cardinality) to a region of lower density. This motion is dictated by the direction of a specific gradient and is often nonlinear in nature, being colloquially associated with a ‘spreading out’ phenomenon. If unopposed, the diffusion process induces a gradual mixing of the diffused quantity in the space it acts upon, eventually resulting in a uniform distribution across the latter.

When used by itself, the term usually refers to particle diffusion, which designates the thermal motion of molecules exerting a stochastic force on their surrounding particles, ultimately propelling them into a Brownian motion (Adamczyk, 2013). The resulting flux of molecules is governed by the concentration gradient, the causal relationship between these variables being mathematically described via Fick’s laws (Fick, 1855). In the closely-related biological context, diffusion represents the net flow of particles through the semipermeable cell membranes, a vital form of passive transportation of both amino acids and nutrients across any living organism (Clark et al., 2018). In thermal engineering, on the other hand, the diffusion of heat is of central interest, constituting a type of heat transfer governed by the temperature gradient and described via the analogue of Fick’s laws, the Fourier’s law (Fourier, 1822). The latter gives rise to the heat equation, one of the most important differential equations in physics and engineering disciplines.

On the other hand, this thesis explores and examines diffusion processes from a public-health and sociological perspective, providing modelling solutions to important open problems in the field, such as estimating the effect of external

interventions for impeding them or optimizing the resource allocation for such interventions when they are scarce. These processes can take many different forms, but they typically appear as *epidemics* that propagate through various channels across some network of human interactions. In the conventional sense, epidemics refer to an increase, often abrupt, in the number of reported cases of a studied disease above the threshold value expected within a given area, commonly known as the endemic level (CDC, 2020). Throughout history, human health and society have faced serious challenges from epidemics, such as the notorious Plagues or the 1918 influenza pandemic, which have stimulated extensive research on their dynamics and prediction. A prominent example in our recent history is the global health crisis triggered by COVID-19, which has prompted the development and application of several novel methods to estimate its impacts and inform policy interventions (Tolles and Luong, 2020). In Section 1.3, we provide essential background on the most explored epidemic models for studying diffusion processes, including agent-based, equation-based, and data-driven models. What is more, we compare and contrast their key advantages and disadvantages, while also surveying some of their important applications and results.

When the diffusion process occurs over graphs, one should note that its dynamics are influenced not only by the specific features and mechanisms of the process itself, but also by the underlying *network structure*. While classic Erdős–Rényi graphs (Erdős and Rényi, 1959) can offer a preliminary insight into the diffusion behavior, they may fail to capture the properties of several social networks, such as the deviation from Poisson degree distributions and the presence of preferential attachment (Newman et al., 2002; Barabási, 2015). As such, more refined construction models are often required to account for these patterns of human interaction: growth models can build preferential attachment networks with power-law degree distributions (Barabási and Albert, 1999; Albert and Barabasi, 2002; Holme and Kim, 2002), rewiring approaches can construct small-world graphs featuring high clustering coefficients (Watts and Strogatz, 1998; Wang and Chen, 2003), while the configuration model can encompass any data-driven degree distributions (Bollobás, 1980; Newman et al., 2001). Since the network topology influences the epidemic outcome significantly, it is possible to exploit the signals stemming from this structural information to analyze or even control its evolution. Chapter 2 reviews the most popular graph generation approaches, surveys common network heuristics, and then introduces two main families of machine learning algorithms for extracting insights from a network’s structure: *graph kernels* (Kriege et al., 2020) and *graph neural networks* (Zhou et al., 2020). While focusing on the latter, our survey also describes several of their recent applications to modelling diffusion processes over interaction graphs that have either stimulated or directly impacted our work.

Building on these introductory chapters, our main chapters aim to address two crucial research questions:

1. *How feasible is contact tracing as an epidemic control strategy considering the challenges of interview processes in manual tracing and application adoption in digital solutions?*
2. *How can we learn policies that maximize the effect of testing, tracing and vaccination under budget constraints?*

Chapter 3 introduces our first contribution to the field, published as [Rusu et al. \(2021\)](#), which outlines a new individual mean-field setup that can analyze the interaction between two competing diffusion processes over arbitrary social graphs: the spread of the SARS-Cov-2 virus, characterized by a specialized compartmental formulation, and the propagation of contact tracing, aided by the strength of the testing strategy underpinning it. The introduced framework, adapted from the work of [Farrahi et al. \(2014\)](#), can address from a modelling perspective four of the open questions formulated by [Anglemyer et al. \(2020\)](#) in their Cochrane Review, shedding light on the efficacy of contact tracing systems against COVID-19: (1) the combined effects of digital and manual tracing can be studied via the proposed *triad network* method, which can realistically simulate the manual and digital tracing processes even when contacts are missed or individuals do not install the tracing application; (2) people with poor mobile network access can be factored in by the degree of overlap parameter, which controls the accuracy of the manual tracing graph; (3) individuals that have privacy concerns can be modelled via the application adoption rate; (4) the ethical and economical repercussions of false positives and false negatives of tracing can be scrutinized through the statistics our model readily captures, showcasing the existence of a trade-off between the number of people who are prevented from becoming infected and those who are wrongfully traced and isolated.

Conversely, Chapter 4 introduces an alternative approach to conducting pathogen testing, contact tracing and vaccination that is adaptable to the budgeting requirements and risk tolerances of regional policy makers, limiting the aforesaid false positives, while still breaking the high-risk transmission chains. The chapter, published as [Rusu et al. \(2022\)](#), introduces several agents that can rank individuals based on their role in their interaction network and the epidemic state over which this diffuses, demonstrating that testing, isolating or immunizing just the top-ranked individuals can attain adequate levels of containment without incurring the costs associated with the inefficiencies of conventional strategies. In addition, it compares the containment rates achieved by each of the policies we consider, demonstrating that a reinforcement learning actor, adapted from the graph neural architecture of [Meirom et al. \(2021\)](#) but efficiently trained via eligibility traces ([Sutton and Barto, 2018](#)), outperforms competitive network heuristics by up to 15%, while far surpassing the

standard random samplers by margins of 50% or more. The performance of the agents under scrutiny is evaluated on multiple network topologies and epidemic parameters, generated by both our individual-based mean-field model and a standard agent-based model, confirming the consistent superiority of the reinforcement learning policy.

All in all, the recent pandemic has shown the challenges of making complex public health decisions when limited evidence of their individual impact is available. In Appendix C, we present several simulation-based visualization techniques that can help policy makers evaluate the effects and impact of distinct interventions, whether they are implemented by human or autonomous agents. First, we use balancing contour plots to illustrate the trade-off between the infection reach and different combinations of public health measures, leveraging data from multiple grid-based simulations. Equipped with a cost-based estimate of ‘moving’ within this space, public authorities can directly employ these plots to decide where additional efforts are worth spending. Second, we demonstrate how a human-in-the-loop could visualize and assess the epidemic state and a control policy’s effectiveness when the latter is suggested by a deep learning agent, as described in Chapter 4. And third, we utilize several local interpretable models to generate figures that display how each feature value impacts the decision-making process of an autonomous agent. Our analysis reveals the importance of heuristic features at the beginning and towards the end of an outbreak, with the epidemic state information steering interventions in-between these phases. This chapter contains material that has been accepted for publication in the IEEE Pervasive Computing journal (Rusu et al., 2023).

It should be noted that, despite our focus on SARS-CoV-2 as a case study, the models developed in this thesis could be applied to other infectious diseases or diffusion processes with little to no adjustments to their internal mechanisms. In Chapter 5, we detail the wider impact of our methods, with potential applications spanning from future pathogen epidemics to analyzing information diffusion. Despite their differences, the latter is mechanistically equivalent to the former, as argued in Section 1.4, with individuals being ‘susceptible’ to getting ‘infected’ by a given piece of information, having a chance to also become ‘infectious’ themselves. As such, the term *epidemic* acquires a broader meaning in this sociological context, where it can denote any persistent and widespread phenomenon that can induce a form of social contagion. This formulation enables important problems, such as influence maximization or rumor detection in social media, to be solved using well-established epidemiological frameworks (Bettencourt et al., 2006; Jin et al., 2013; Kumar et al., 2020). Section 1.4.2 surveys some of the most popular mathematical models for information diffusion, while Section 5.2 further explains the mechanisms by which our epidemic simulation framework can be applied to this domain. In particular, we highlight the equivalence of testing and tracing programmes with monitoring and feedback processes in the informational domain. Moreover, we show how the problem

of controlling the spread of viruses can be formulated as an influence maximization task, enabling the use of the control agents introduced here to find optimal solutions.

1.2 Findings and contributions

The principal findings, contributions and developments outlined in this thesis are summarized in the following:

- A new model for simulating COVID-19 outbreaks based on *triad* topologies is introduced, enabling the study of realistic scenarios involving testing, manual and digital tracing: limited application uptakes, staffing inefficiencies, poor contact recollection of index cases etc. Using this, we show that digital tracing remains an effective pathogen control mechanism even at achievable uptakes (i.e. $< 50\%$) when combined with adequate levels of testing and manual tracing.
- A novel compartmental formulation is proposed, called SEIR-T, which separates the traced/isolated status from the infection states, thus ensuring the disease progression is independent of the tracing mechanism. This aspect gives versatility to our methodology, enabling it to be directly applicable to any existing compartmental descriptions of diffusion processes.
- The introduced simulation setup can be used to answer from a modelling perspective four of the open problems formulated by [Anglemyer et al. \(2020\)](#), including the necessity for models to consider mobile network and internet access issues, and the prevalence of tracing false positives. The latter are shown to be alarmingly numerous when tracing efforts are blindly scaled up for the purpose of achieving swift containments, motivating the employment of more targeted strategies.
- The control methodology of [Meirom et al. \(2021\)](#) is extended to accommodate the concurrent prioritization of testing, tracing, and vaccination in COVID-19 epidemics diffusing over interaction networks, in the presence of budgeting constraints. We call the resulting simulation-control framework *EpiCURB*.
- A new reinforcement learning agent, backed by graph neural networks, is demonstrated to attain up to 15% higher containment rates than competitive heuristics, and more than 50% compared to standard random samplers. The reasons for its success are hypothesized to be the superior timing of detection and the swift identification of nodes with high-risk transmission potential, factors which get examined in-depth.
- A novel and efficient adaptation of the Proximal Policy Optimization algorithm ([Schulman et al., 2017](#)), which enables online learning through eligibility traces,

is proposed, constituting the foundation of the reinforcement learning procedure that supports our targeted methodology for testing, tracing, and vaccination.

- A simple yet effective improvement to any type of epidemic control agents is presented: recollection of negative test results. This significantly increases the fraction of people that are prevented from infection.
- Several interventional strategies are assessed in their ability to control epidemics generated by either a multi-site mean-field or an agent-based model. The reinforcement learning policy is shown to work equally well in both setups.
- An effective approach to evaluate the impact of individual public health interventions when they are implemented in combination is proposed: generating contour plots that depict the efforts invested in each measure against the resulting containment level. This can aid policy makers in determining the optimal resource allocation for every intervention given a target infection reach, complementing prioritization strategies that depend on an input budget.
- Different visualization techniques for assessing the epidemic state and the deployed deep learning control agent are presented. Ranging from node embedding projections to dendrograms, these can inform policy makers of the pathogen's reach across the graph, the agent's overall performance, and the communities or regions most at risk. Starting from the latter, a novel and effective baseline for targeting mass vaccination campaigns is proposed.
- The GraphLIME method (Huang et al., 2022) is introduced as a dependable tool for explaining the decisions that are made by a graph neural network control agent, making them accessible for audit and public scrutiny. This technique allows us to establish the influence of network heuristics in the decision-making process, while revealing the periods when the epidemic state information becomes pivotal.

The following publications, awards and submissions are based on these contributions:

- Journal paper titled 'Modelling digital and manual contact tracing for COVID-19. Are low uptakes and missed contacts deal-breakers?', published in *PLOS ONE* (Rusu et al., 2021).
- Conference paper titled 'Flattening the curve through reinforcement learning driven test and trace policies', published in the *16th EAI International Conference on Pervasive Computing Technologies for Healthcare*, and recipient of the *Best Paper Award* (Rusu et al., 2022).
- Manuscript titled 'EpiCURB: Learning to Derive Epidemic Control Policies', accepted for publication in *IEEE Pervasive Computing* (Rusu et al., 2023).

1.3 Modelling the diffusion in pathogen epidemics

Pathogen epidemics are among the most prominently studied diffusion processes over interaction networks, as they exemplify how infections pervade among individuals over time and space, and how interventions can influence their patterns and consequences. The objective of epidemic modelling is to develop and implement mathematical or computational models that can simulate the transmission of infectious diseases in a population, and the impact of measures enacted to control or mitigate them. These can provide insights into the dynamics and features of epidemics, such as their transmission modes, growth rates, durations, and others. They can also assist in forecasting the future trajectory and impact of pathogen diffusion, as well as inform policy choices and public health actions.

There exist different types of epidemic models, depending on the level of detail, complexity, and realism they are built to incorporate. Some of the most prevalent and influential types are:

- **Equation-based models (EBMs):** These are mathematical models that use a set of differential equations to describe the average behavior or trends of the epidemic across different compartments, which typically separate individuals based on their health status. One of the most widely-used EBMs for modelling epidemics is the Susceptible-Infectious-Recovered (SIR) model, which [Kermack et al. \(1927\)](#) developed based on the earlier work of [Ross \(1916\)](#), to investigate the dynamics of highly-contagious diseases such as malaria and plague. The ordinary differential equations (ODEs) associated with this model can be consulted in Eq 1.1, where S is the number of susceptible individuals, I the number of infectious, and R the number of recovered/removed, β the transmission rate, and γ the recovery rate. EBMs can be adapted or modified in various ways to account for more realistic and/or pathogen-specific scenarios (e.g. Zika ([Aik et al., 2017](#)), Ebola ([Berge et al., 2017](#)), SARS-CoV-2 ([Zhao and Chen, 2020](#); [He et al., 2020](#); [Giordano et al., 2020](#))). As such, one can add more compartments (e.g. exposed but not yet infectious – E state in the SEIR formulation of Eq 1.2, asymptomatic, symptomatic, hospitalized), remove some of them (e.g. SI, SEI – no recovery possible), or cycle back to certain states (e.g. SIS, SIRS – nodes return to being susceptible after a while) to represent different stages or types of infection ([Zhao and Chen, 2020](#); [He et al., 2020](#); [Giordano et al., 2020](#); [Meirom et al., 2021](#)). One can also incorporate public health interventions into these models, such as vaccination, quarantine, isolation, and social distancing, that can significantly influence the dynamics of the simulated disease ([Dashtbali and Mirzaie, 2021](#); [Prabakaran et al., 2021](#)). Furthermore, one can relax the assumptions of constant parameters, allowing for time-varying transmission rates ([Kucharski et al., 2020](#); [Jing et al., 2021](#)). There exist three

main approaches to solving EBMs: analytical methods, which use mathematical techniques to find exact or approximate solutions in terms of functions or series (Kermack et al., 1927); numerical integration, which uses algorithms to approximate the solutions by discrete values (Okabe and Shudo, 2020); and stochastic simulation, which applies random processes to model the dynamics of the system as continuous-time Markov chains (Allen, 2017; Farrahi et al., 2015). The analytical methods usually have lower computational costs, but they may have limited applicability or feasibility for systems with higher complexity or nonlinearity (Kolk and Lerman, 1992). Numerical integration methods, on the other hand, enable the study of complex dynamics that cannot be described analytically, but their main limitation lies in the accumulation of approximation errors over time, impacting the accuracy of long-term predictions (Boyce et al., 2022). Lastly, a stochastic simulation of an EBM for epidemics requires sampling from the distribution of the inter-transition time of an individual to a different health state (Farrahi et al., 2015), thus reflecting the intrinsic randomness and variability of the underlying diffusion process. This hybrid approach is suitable for modelling scenarios with individual-level interactions, heterogeneity, or stochastic events, yet it also entails more intricacy and computational demand (expanded upon in the description of hybrid models).

$$\begin{aligned} \frac{dS}{dt} &= -\beta SI \\ \frac{dI}{dt} &= \beta SI - \gamma I \\ \frac{dR}{dt} &= \gamma I \end{aligned} \quad (1.1)$$

(1) SIR model

$$\begin{aligned} \frac{dS}{dt} &= -\beta SI \\ \frac{dE}{dt} &= \beta SI - \epsilon E \\ \frac{dI}{dt} &= \epsilon E - \gamma I \\ \frac{dR}{dt} &= \gamma I \end{aligned} \quad (1.2)$$

(2) SEIR model

- **Agent-based models (ABMs):** Also called individual-based models (Railsback and Grimm, 2011), these involve computational representations of autonomous agents, their behavior and interactions within a given dynamic environment, where each agent can have its own attributes and rules. ABMs can capture the dynamics of complex systems and phenomena, reflect the heterogeneity and adaptability of individual behavior, and model the feedback effects of the simulated environment. They can also incorporate different types of agents, such as humans, animals, vectors etc., depending on the nature of the disease and its hosts. ABMs are often used to study scenarios that involve spatial or temporal variations, behavioral changes, or stochastic events. To do so, they often account for either the interaction network structure of the population, ignoring other details about physical locations, or an explicit spatial representation of the population, where agents can move and interact. A

network-based ABM simulates the disease dynamics in discrete time intervals using a graph representation of individuals as nodes and their contacts as edges (Abueg et al., 2021; Meirom et al., 2021). At every time increment, a Bernoulli trial with predefined probabilities is performed for each link and each vertex of the network, the outcome of which determines the execution of interaction-level (e.g. transmission, contact tracing) and individual-level events (e.g. recovery, testing, isolation, vaccination). The probabilities for the former may be influenced by factors such as contact frequency and duration, often encoded as edge weights. In contrast, spatial models (Railsback and Grimm, 2011) use a spatial grid (Nardini et al., 2021) or a continuous space (Nadini et al., 2020) to represent the physical environment where individuals are located, with contact patterns governing the pathogen diffusion. Simulations of the latter can be done through both continuous- and discrete-time approaches. Compared to EBMs, ABMs tend to be relatively complex and resource-intensive to simulate, the involved cost often being justified by their degree of granularity and ability to monitor public health interventions at the individual level (Sukumar and Nutaro, 2012). Government-advising groups in the UK decided to employ this paradigm early on in the COVID-19 pandemic to estimate the effects of such interventions (Ferguson et al., 2020; Hinch et al., 2020b). A more recent study analyzed the combined effects of manual tracing with digital solutions, at various application uptakes, via a rich yet scalable ABM fitted to mobility data from several counties in Washington (Abueg et al., 2021).

- **Data-driven models:** These are models that rely on empirical data to fit or calibrate their parameters or functions, employing statistical or machine learning (ML) techniques to infer or forecast the epidemic trends. These models can utilize various data sources, such as case reports and hospital records (Murray, 2020; Fritz et al., 2022), mobility data (Kapoor et al., 2020; Lozano et al., 2021) or social media data (Ansell and Dalla Valle, 2023; Lucas et al., 2023), depending on their availability and quality. They can also employ different methodologies, such as parametric curve fitting (Lozano et al., 2021), regression analysis (Fritz et al., 2022), autoregressive time series analysis (Reinhart et al., 2021; Murray, 2020; Adiga et al., 2020), or deep learning (Lozano et al., 2021; Kapoor et al., 2020; Panagopoulos et al., 2021; Lucas et al., 2023), according to their objectives and assumptions.
- **Hybrid models:** The techniques reviewed above are often integrated in order to account for the complexity and heterogeneity that characterize the dynamics of an epidemic. Hybrid models can leverage the strengths and overcome the limitations of each individual method, providing more realistic or flexible representations of the diffusion process. Notable examples here include: enhancing compartmental models with agent-based features through the use of

contact networks, thus generating epidemic trajectories based on multi-site mean-field formulations, appearing as either individual-based mean-fields (IBMF), described via ODEs (Newman, 2010; Pastor-Satorras et al., 2015; Shah et al., 2020) or probabilistic equations of state transitions (Huerta and Tsimring, 2002; Farrahi et al., 2014, 2015), or pair approximation mean-fields (e.g. degree-based method of Pastor-Satorras et al. (2015), cluster-based approach of Huerta and Tsimring (2002); Newman (2010)); agent-based models employing compartmental formulations to track the state of the individuals (Abueg et al., 2021; Meirum et al., 2021; Hinch et al., 2020b); deriving time-varying transmission rates in EBMs (Di Domenico et al., 2020; Kucharski et al., 2020; Jing et al., 2021) or ABMs (Abueg et al., 2021) from real case reports or hospital records; inferring appropriate compartmental parameters using machine learning (Arik et al., 2020) or Bayesian techniques (Liu et al., 2020a; Flaxman et al., 2020); utilizing employment (Abueg et al., 2021) or demographic data (Riley and Ferguson, 2006) to generate or calibrate the interaction networks over which the epidemics spread; deriving meta-population models using networks obtained from mobility datasets that connect and permit the movement of agents between subpopulations, individually modelled by separate EBMs (Poletto et al., 2013; Rahmadani and Lee, 2020; Chang et al., 2021). Given a transmission rate β , a recovery rate γ , and the adjacency matrix A of the interaction network, the ODEs of an IBMF SIR model (Newman, 2010; Pastor-Satorras et al., 2015) track across time the average probabilities S_i , I_i and R_i for any vertex i to be in each of the characteristic epidemic states:

$$\frac{dS_i}{dt} = -\beta S_i \sum_j A_{i,j} I_j \quad \frac{dI_i}{dt} = \beta S_i \sum_j A_{i,j} I_j - \gamma I_i \quad \frac{dR_i}{dt} = \gamma I_i \quad (1.3)$$

Despite its high degree of granularity, the ODE formulation in Eq 1.3 may still be inadequate for modelling the complex dynamics that often underlies epidemics, as it relies on approximations and assumptions that can fail to capture the intricacy and uncertainty of the corresponding diffusion process. As discussed in Newman (2010), the contact network and the infection events may be subject to variability and stochasticity, which can significantly influence the dynamics and outcomes of the simulated epidemic. Moreover, numerical or analytical methods that directly solve the system may yield unacceptable errors or inaccuracies. These aspects require a more sophisticated modelling approach that can account for the correlations between the state probabilities of different nodes.

In Chapter 3, we introduce a novel IBMF hybrid model for the SARS-CoV-2 virus that is stochastically simulated, thus accounting for the intrinsic randomness and fluctuations of the viral transmission, while preserving the mathematical rigor of EBMs. Building on the previous work of Farrahi et al. (2014) and Huerta and Tsimring (2002), our experimental setup studies the antagonist interaction

between the spread of the SARS-Cov-2 virus and contact tracing. In this context, we present a compartmental formulation that captures the evolution of both processes across every node of the network, simulating them in continuous time using numerical methods. What is more, we analyze their behavior under different scenarios and parameter settings, deriving important public-health conclusions about the overall containment efficacy of “test and trace” strategies. In contrast, Chapter 4 carries out a comparative analysis of the performance of heuristic-based and learned epidemic control strategies on both the IBMF model introduced in this thesis, as well as an ABM setup, revealing the inherent modelling transferability that exists between these approaches.

1.4 Going beyond pathogen diffusion

Zooming out from the pathogen epidemic case study, we remark there are a considerable number of examples in the literature of exploiting epidemiological models to analyze other domains that involve complex systems with interacting agents, such as information diffusion, viral marketing, cybersecurity, or network analysis. For example, [Jin et al. \(2013\)](#) use the SEIZ compartmental model for information diffusion to study the propagation of fake news on Twitter, while [Mutlu et al. \(2021\)](#) employ it to predict how the information would cascade on the platform by taking into account the cognitive processing depth of users. In the context of viral marketing, [Gonçalves et al. \(2017\)](#) utilize the SIR model to analyze the mathematical properties of a viral campaign. In the cybersecurity domain, [Pastor-Satorras and Vespignani \(2001\)](#) employ an SIS model to study the spread of computer viruses, while [Prakash et al. \(2012\)](#) use the SI formulation to detect the source of malware in a network and how to trace back the infection path using the former’s topology. In network analysis, [Fanti and Viswanath \(2017\)](#) leverage a network-based diffusion model to infer the hidden identities of Bitcoin users from their message patterns in the blockchain, and how to deanonymize them using deep learning techniques. In addition to analyzing the state of the system and/or its agents, methods for node-level intervention have been suggested or indicated as feasible for most of these domains. For example, in information diffusion and viral marketing, it is often of interest to maximize the spread of positive influence or minimize the spread of negative influence by selecting the optimal set of nodes to initiate or block a diffusion process ([Kempe et al., 2003](#); [Meirom et al., 2021](#)). Analogous to real epidemics, one may want to minimize the risk of infection or maximize the robustness of a computer network by removing or immunizing the most vulnerable or influential nodes. In network analysis, one may want to maximize the anonymity or minimize the traceability of nodes by hiding or revealing their identities and relationships ([Fanti and Viswanath,](#)

2017). These subset selection problems can be formulated and solved using similar techniques to those we introduce in Chapter 4.

In the following subsections, we extensively explore the domain of information diffusion, which is both the most well-studied and the most amenable after pathogen epidemics to the simulation-control methods presented in this work. What is more, we provide a detailed overview of some of the principal modelling techniques for analyzing how information propagates in different settings.

1.4.1 Information diffusion problem

The problem of information diffusion appears in many different contexts as it represents a central mechanism for knowledge sharing, learning, and idea adoption. With the rise of the web and social media, people started to get exposed to an unprecedented amount of new information, thus opening up the potential for exploiting parts of this data flow for marketing purposes (Even-Dar and Shapira, 2007) or political gains (Anderson, 2017), shaping people’s perception about given products, ideologies or even historical events (Klein, 2020) in the process. This informational flow has also been shown to be a good predictor of human behavior, even in the absence of complete profile data (Bagrow et al., 2019). As such, predicting how a piece of information would spread through a social network, and potentially maximizing its influence on individuals, is of utmost importance for social sciences, civic engagement, business and electoral campaigns (Arnaboldi et al., 2016). At the same time, in today’s world, it becomes increasingly more crucial for the worldwide community to combat malicious acts in the processes of information diffusion, arising from the spread of fabricated news, violent extremist propaganda, and other dangerous content.

One of the most important and well-studied problems in information diffusion is influence maximization (Kempe et al., 2003; Wang et al., 2013; More and Lingam, 2019; Kimura et al., 2009). The objective in this domain is to identify the best set \mathcal{I} of nodes from a given graph that, when “activated”/“infected”, would maximize the extent of the information diffusion, often under constraints like limited budgets or time. For preventing malicious acts, it is often enough to detect and eliminate the malignant content, yet some situations require a more significant intervention to curb their negative effects, such as removing the perpetrator from the social network or finely understanding the nature of the spreading dynamics. To find the optimal set of influencers \mathcal{I} in the network, a standard hill-walking procedure is typically employed, maximizing at every iteration the expression $\iota(\mathcal{I} \cup \{v\})$ over all the nodes v , where $\iota(\cdot)$ measures the *influence degree* of the input set (Kempe et al., 2003; Kimura et al., 2009). This algorithm requires a reliable way to repeatedly evaluate ι , while also assuming that adding one node at a time to \mathcal{I} will converge to a good solution. A common yet inefficient approach to estimate ι is to run simulations of the underlying propagation

model, choosing the current elements of \mathcal{I} as infection seeds, and measuring the resulting information reach (Kempe et al., 2003). More resource-effective approaches have also been proposed, with notable examples including heuristic-based (Kimura et al., 2009) and node centrality-based methods (More and Lingam, 2019). In analogy with the problem of identifying optimal test subjects in pathogen outbreaks, these methods employ various criteria or techniques to rank the nodes and select the seed set, a task that can be substantially enhanced by deep learning models (refer to Section 2.9.2 and Chapter 4 for more details).

Going beyond social networks, information diffusion can be seen as a universal mechanism that transfers data among nodes in a graph, regardless of the nature of the vertices or the data. For instance, information diffusion can occur when a node aggregates information from its adjacent nodes to fill in missing values or update outdated entries. Another prominent example is given by the transmission of stimuli charges from one neuron to another in biological or artificial neural networks.

1.4.2 Modelling the diffusion of information

Among the most common ways of modelling information diffusion, we note three significant directions:

- **Random walks:** A common way to model information diffusion over networks is by employing Markov random walks (Neumann et al., 2016). Given an initial state $W_0 = v_i$, i.e. the source vertex, the one-step transition probabilities encoded in T , with $T_{j,k} = P(W_{t+1} = v_k \mid W_t = v_j)$, define a random walk on the network, starting at v_i and ending after t_f iterations: $W_i = \{W_t \mid t \in [t_0, t_f] \text{ and } W_0 = v_i\}$. When no prior information about the probabilities is known, it is customary to set T to the normalized graph Laplacian $\hat{L}^{rw} = D^{-1}A$, where D is the degree matrix (more on notations in Section 2.2). In this context, each node typically carries a d -dimensional representation $X_0 \in \mathbb{R}^{N \times d}$ of its prior information at time t_0 . For example, if the goal is to spread labels across the network, these vectors can be formed such that $y_i = \arg \max X_0(i, :)$ is the class of node v_i . The standard unconditioned diffusion process, given by $X_{t+1} = TX_t$, has the downside of converging to a constant steady-state, which is rarely informative. As a result, a predefined number of iterations is often used as a criterion for early stopping (Szummer and Jaakkola, 2002). Another major downside of the former method stems from the fact that connected vertices always inherit and retransmit their most recent update. One could restrict this behavior by making the transition matrix \hat{T} describe an absorbing random walk (see Eq 1.4), thus allowing for a set of diffused elements S to ‘stick’ to the nodes that receive them.

Such a procedure can safely be used for label propagation in a semi-supervised learning setting, where class changes after the initial labelling are to be avoided.

$$\hat{T}_{i,j} = \begin{cases} 0 & i \in S \text{ and } i \neq j \\ 1 & i \in S \text{ and } i = j \\ T_{i,j} & \text{otherwise} \end{cases} \quad (1.4)$$

Zhou et al. (2003) proposed a modelling technique that combines ideas from the approaches above, leading to a formulation that performs a weighted summation of the old and new representations. Let the transformation matrix be $T = \hat{L}^{sym} = D^{-\frac{1}{2}}AD^{-\frac{1}{2}}$, with α a tunable parameter, and Y a prior information matrix (e.g. $Y = X_0$), then the improved diffusion process becomes:

$$X_{t+1} = \alpha TX_t + (1 - \alpha)Y \quad (1.5)$$

Beyond their application to the domain of information diffusion, the ideas discussed here underlie many graph learning algorithms as well, such as diffusion kernels and graph neural networks (see Chapter 2 and Appendix A). For instance, the GCN model of Kipf and Welling (2017) draws inspiration from Zhou et al. (2003), Eq 1.5 resembling a linear GCN layer with a skip-connection.

- Epidemiological framework:** The random walk approaches described above are effective in tasks where the dynamics T of the diffusion process is static, and it is either known or entirely dependent on the network structure (e.g. deterministic label propagation). However, for problems with a high degree of stochasticity, such as analyzing influence maximization (Wang et al., 2013) or identifying rumors (Jin et al., 2013) in social networks, it is often preferable to deduce the underlying dynamics, and then represent it in a flexible, easily replicable format that allows for further simulations of increasingly more complex scenarios. Well-studied epidemiological frameworks are a convenient way to achieve all these goals. In particular, compartmental models have successfully been applied to a broad range of information diffusion problems. Despite their simplicity, they tend to provide accessible descriptions of the modelled dynamical systems based on mathematical equations, ensuring both reproducibility and rigor. Following the analogy with epidemic modelling, this domain has also adopted both non-graph and hybrid graph-based compartmental formulations (Li et al., 2019). The standard SIR model has also become a popular non-graph model for information diffusion (Li et al., 2019), but it is usually better suited for cases where the spread of information is transmitted exactly once by each individual, such as general news (Abdullah and Wu, 2011), rumors (Sudbury, 1985; Yan et al., 2022), or memes (Wang and Wood, 2011) that are both attractive and ephemeral. That being said, the model presents some limitations in describing




FIGURE 1.1: **The SEIZ information diffusion model.** Diagram of possible transitions is given on the left, while the characteristic ODE equations are presented on the right. The possible system states are: *S* (Susceptible), *E* (exposed), *I* (infectious), and *Z* (skeptical). Source: Jin et al. (2013).

the general case of information diffusion, as it assumes that individuals exposed to an idea will necessarily adopt and propagate it (Bettencourt et al., 2006). This assumption can lead to poor data fits in some cases (Rui et al., 2018; Jin et al., 2013), aspect that has motivated the emergence of many extensions that can also model the concept of indecision/skepticism, such as SPIR (Rui et al., 2018), SEIZ (Bettencourt et al., 2006; Jin et al., 2013), SHIR (Liu et al., 2016), and others.

The SEIZ compartmental formulation, illustrated in Fig 1.1, draws inspiration from the SEIR model, allowing for an ‘incubation’ period to exist before an affected node spreads any information further. Instead of recovering, nodes in SEIZ are allowed to be ‘skeptics’ to what they read or hear (state *Z*), effectively stopping the spread through their corresponding pathways. Entering in contact with either infectious (βI rate) or skeptics (bZ rate) leads to different outcome probabilities, with susceptible vertices having a chance to become exposed, directly infectious, or skeptics as a result. The basic SEIZ model assumes that individuals interact on a fully connected graph, and no hybrid variants that incorporate heterogeneous topologies have been proposed thus far. Jin et al. (2013) used the SEIZ model to study how true and false news spread on Twitter. They chose eight stories, four real and four fake, and estimated the SEIZ parameters for each one by fitting the predicted number of infectious users to the actual number of tweets sent at each time point, using a non-linear least squares algorithm. They found that the rate of believing and sharing a story, $R_{SI} = \frac{(1-p)\beta + (1-l)b}{\rho + \epsilon}$, is significantly higher for true news than for rumors, since the latter are generally less likely to be believed, but more likely to be rampantly shared by adopters.

In contrast, graph-based epidemiological models have received significantly less attention from the information diffusion literature, mostly due to the difficulty of recovering complete network structures (Li et al., 2019; Choudhury et al., 2010). Nevertheless, some studies have used graph-based formulations to study

various aspects of information diffusion. For instance, [Kumar and Sinha \(2021\)](#) combine a Bass approach ([Bass, 1969](#)) for modelling idea adoption with a spreading mechanism based on SIS to investigate how information propagates over social networks. Similarly, [Zhan et al. \(2019\)](#) use the SI model to investigate the impact of temporal networks on the efficiency of the diffusion process, while [Matsubara et al. \(2012\)](#) extend it to analyze “spikes”, which are the abrupt surges of activity or popularity that a piece of information can exhibit. Despite the prevalent use of social influence models for solving influence maximization tasks, some studies have also employed the SI model or its variants to simulate the underlying diffusion process ([Wang et al., 2013](#); [More and Lingam, 2019](#)).

- **Social influence models:** Another conventional method for understanding how information propagates through a network, predominantly used for social influence problems, is to simulate either the Independent Cascades (IC) or the Linear Threshold (LT) models ([Kempe et al., 2003](#)). The former involves a continuous mechanism, akin to the SI/SIR dynamics on networks, where each “active” node has a one-time opportunity of activating its “inactive” neighbors, determined by a fixed probability that does not depend on previous events. The strong link between the two modelling approaches is also acknowledged by [Kimura et al. \(2009\)](#), being an important motivator for the emergence of equivalent epidemiological formulations for addressing such tasks ([Wang et al., 2013](#); [More and Lingam, 2019](#)). In contrast, LT posits that each node has a threshold that indicates the proportion of its neighbors who must adopt a piece of information before it does so itself ([Granovetter, 1978](#); [Watts et al., 2007](#)). Similar to IC, the diffusion process is stochastic, with updates taking place in discrete time steps, and ending when no more adoptions are feasible. At each time increment, every node that has not adopted the information examines the fraction of its neighbors who have done so. If the fraction is equal to or higher than its personal threshold, the adoption is said to have occurred.

Chapter 2

Graphs: Generation and Modelling

2.1 Introduction

This chapter provides the essential background and literature review for generating and modelling network data, focusing on the techniques frequently utilized for studying diffusion processes. After introducing some general notation necessary for examining our work, we survey the most popular approaches for generating graph data, as well as two of the most widely-used methods for solving graph-based problems: network heuristics and machine learning algorithms. The generation approaches we discuss here enable us to analyze how the pathogen may diffuse across a broad range of interaction patterns, offering a realistic view of the heterogeneous spreading dynamics characteristic of respiratory diseases. The graph-based heuristics and machine learning methods, on the other hand, are pivotal for the control framework we propose in Chapter 4, as they are employed for ranking all network participants with the aim of recommending public health measures that effectively lower infection rates, while also adhering to strict budgeting constraints. Graph neural networks (GNNs) and their more expressive variants represent one of the key focuses of this thesis, as they enable us to effectually leverage the information contained within the interaction networks' complex structures and features. This aspect is essential for the success of our final solution to choosing the most beneficial interventions: a reinforcement learning (RL) agent, adapted from the work of [Meirom et al. \(2021\)](#), reviewed in Section 2.9.2, that exceeds all the competitors we benchmark in terms of containment levels.

2.2 Notations

A static graph G is a tuple $(V(G), E(G))$, where V and E are general-purpose functions mapping any such G to its sets of vertices and edges, respectively. The corresponding nodes and links may also have feature vectors $f_v(\cdot) \in \mathbb{R}^k$ associated with them, usually stacked as rows in a design matrix X (more common for nodes), and/or labels $l(\cdot)$, where $l: V(G) \cup E(G) \mapsto \Sigma$ is a map to some common alphabet Σ . We assume without loss of generality that an edge's weight $w_G(v, u) = A_{v,u}$ is the first (and sometimes the only) element of its feature vector, where A denotes the adjacency matrix of graph G . Although link-weighted versions exist for many network-related quantities and algorithms, it is common for unweighted variants to be used instead, where every weight is effectively assumed to be 1. Unless explicitly stated otherwise, we employ the unweighted version of these routines in this thesis.

Aside from the adjacency matrix A , all graphs can also be represented via the degree matrix D (diagonal matrix containing the degrees of each vertex), the graph Laplacian $L = D - A$, or a variant of the normalized graph Laplacian: $L^{sym} = D^{-\frac{1}{2}}LD^{-\frac{1}{2}}$ or $L^{rw} = D^{-1}L$. For directed graphs, we can redefine the above by replacing D with the in-degree D_{in} or out-degree D_{out} matrices. If a specific routine only works on directed networks, an undirected graph can be made directed by doubling its set of edges such that the resulting in-degree and out-degree of each vertex stays equal:

$D_{in} = D_{out} = D$. When self-loops are artificially added to all nodes, the prior quantities are denoted with ' \sim ' (e.g. $\tilde{A} = A + \lambda I$, for $\lambda \geq 1$). Some algorithms also use a modified version of the aforementioned Laplacians: $\hat{L} = I - L$. We say that two networks G_1 and G_2 are isomorphic, denoted with $G_1 \simeq G_2$, if a bijection $\psi: V(G_1) \mapsto V(G_2)$ exists, such that $(v, u) \in E(G_1)$ if and only if $(\psi(v), \psi(u)) \in E(G_2)$. Determining network isomorphism is neither NP-complete, nor is there any known polynomial-time algorithm for solving it (Schöning, 1988). A classic heuristic for obtaining approximate solutions for this problem, which unfortunately provides no guarantees, is the Weisfeiler-Lehman (WL) test (Weisfeiler and Lehman, 1968). This involves iteratively refining the labels of all nodes by aggregating information from their neighbors, followed by a comparison of the final representations of the networks to be compared. As the number of nodes k that get updated at once in each iteration grows, the k -WL test becomes increasingly more powerful. Aside from testing isomorphism, this heuristic plays a central role in deriving several graph kernels and in characterizing the expressivity of graph neural networks. As a matter of fact, a firm link between 1-WL and standard GNN models has recently been established (Xu et al., 2019). This link has resulted in the development of architectures that attain the expressivity level of 1-WL, and techniques that overcome some of its limitations, both of which are exploited by our modelling solution in Chapter 4. The 1-WL label update rule for iteration i can be mathematically written as follows:

$$l^i(v) = \text{relabel}((l^{i-1}(v), \text{sort}(\{l^{i-1}(u) \mid u \in N_G(v)\}))) \quad (2.1)$$

Each node v has a particular neighborhood in the network it is part of, denoted with $N_G(v) = \{u \in V(G) \mid (v, u) \in E(G)\}$. This neighborhood determines the degree of v in graph G , $\text{deg}_G(v) = D_{v,v} = |N_G(v)|$. The degree distribution of graph G , given by $P(\text{deg}_G(v) = x) = \frac{1}{n} \sum_{v=1}^n \mathbb{I}(\text{deg}_G(v) = x)$, where \mathbb{I} is the indicator function, is a useful measure for characterizing its heterogeneity and the diversity of its connectivity patterns. We denote the set of all paths from node v to any node u as $P_G(v, u)$, with $\sigma_G(v, u) = |P_G(v, u)|$, and define the shortest path distance between them as $d_G(v, u) = \min_{p \in P_G(v, u)} \sum_{(a,b) \in p} w_G(a, b)$. The latter is used by several network heuristic methods (see Section 2.4).

If there are no predefined vertex labels, a few iterations of the WL update rule can be run, using default values in-place of the missing ones (e.g. $\text{deg}_G(v)$). If a network of interest changes over time, we can model each individual time snapshot as a static graph G_t . As such, the quantities defined above can be repurposed for dynamic networks by annotating them with the corresponding timestamp as a subscript (or superscript when further indexing is needed).

Two approaches that we touch upon in this thesis, diffusion kernels (Kondor and Lafferty, 2002) and graph neural diffusion (Chamberlain et al., 2021), are directly associated with the heat diffusion equation in physics and its homogeneous variant (i.e. when the diffusivity is a constant scalar – $g(u, t) = c$). As such, we remind the reader of these expressions in Eqs 2.2 and 2.3, where $x(\cdot, t)$ represents the distribution of some property $u \in U$ at time t , while ∇ , div and Δ are the gradient, divergence, and Laplace operators, respectively.

$$\frac{\partial x(u, t)}{\partial t} = \text{div}[g(u, t) \nabla x(u, t)] \quad (2.2)$$

$$\frac{\partial x(u, t)}{\partial t} = \text{div}[c \nabla x(u, t)] = c \Delta x(u, t) \quad (2.3)$$

2.3 Network generation

Network generation methods are techniques that can create graphs with certain properties and characteristics, such as size, density, degree distribution, clustering coefficient, community structure, and many others. These methods can be classified into different categories based on their underlying principles and assumptions: (1) random graph models, where the edges are randomly assigned between nodes and no further constraint is put on the topology; (2) growth models, where the nodes and edges are added sequentially according to some factors like the desired degree distribution or fitness of the graph; (3) rewiring models, where the connections of an

existent network are altered according to specific rules; (4) geometric models, where the spatial position of the nodes guides the creation of edges; and (5) learning models, where GANs or evolutionary algorithms optimize the set of links according to real data. We present a brief survey of some of the most prevalent classical graph generation methods, which we employ extensively throughout our thesis:

- **Erdős-Rényi model:** Developed by Paul Erdős and Alfred Rényi in 1959 (Erdős and Rényi, 1959), this is probably the most prevalent method for generating random networks. There exist two distinct formulations of the Erdős-Rényi (ER) model: $G(N, M)$ and $G(N, p)$. The former creates a network with N nodes and M randomly chosen edges, while the latter creates a network with N vertices and connects each pair of nodes with probability p , thus having an expected number of edges equal to $\frac{pN(N-1)}{2}$. The average degree K of the graphs can be virtually fixed by using the first formulation, with $M = \frac{NK}{2}$. The ER model represents a useful tool for analyzing various problems in combinatorics and probability (Bollobás, 2001; Janson et al., 2011), computer science (Mitzenmacher and Upfal, 2005), physics (Albert and Barabasi, 2002), biology, or social sciences (Newman and Park, 2003). It is worth mentioning that, although the epidemiological literature has widely adopted it, this graph model can be unsuitable for capturing the interaction patterns of certain real social networks (Newman et al., 2002). However, the inherent ability to accommodate the characteristics of randomly mixed populations (Keeling and Eames, 2005) makes the ER graphs adequate vehicles for studying pathogen outbreaks in public places, such as stores or mass transit conveyances (Abueg et al., 2021). Random mixing models, in turn, have been shown to offer acceptable estimates of the total epidemic size when the transmission probability is high or the infectious period is relatively small (Smieszek et al., 2009), as has been the case for several COVID-19 breakouts. As such, we consider this graph generation method an acceptable baseline for analyzing the spread of SARS-CoV-2, employing it extensively in Chapter 3 to study the effects of contact tracing.
- **Configuration model:** First introduced by Béla Bollobás in 1980 (Bollobás, 1980), and later extended in Newman et al. (2001) and Newman (2003), this approach builds random graphs from an arbitrary degree sequence, expected as a parameter. In Chapter 4, we sample multiple such sequences from a degree distribution fitted to real contact tracing data in order to obtain realistic estimates of our control method's effectiveness.
- **Barabási-Albert model:** First described by Albert-László Barabási and Réka Albert in 1999 (Barabási and Albert, 1999), this represents a popular growth model that utilizes a preferential attachment mechanism to generate scale-free (SF) graphs, which feature power-law degree distributions. The method receives

N and an integer m_1 as parameters, and constructs networks by sequentially attaching new nodes, each with m_1 edges preferentially attached to the high degree vertices, which eventually become *hubs*. The average degree can be virtually fixed by using a modified variant, called dual Barabási-Albert (Moshiri, 2018), which expects N , K , and 2 integers m_1 and m_2 as parameters. For each new node added to the graph, there will be m_1 incident edges with probability p , or m_2 edges with probability $1 - p$, that are preferentially attached to existing nodes, where p gets calculated according to:

$$p = \frac{NK - 2m_2(N - \max(m_1, m_2))}{2(N - \max(m_1, m_2))(m_1 - m_2)}.$$

It is a known fact that the SARS-CoV-2 virus is an overdispersed infectious agent (Endo et al., 2020; Adam et al., 2020), and like many other pathogens with a high epidemic potential (Lloyd-Smith et al., 2005), its diffusion is largely driven by ‘superspreading’ events (Lewis, 2021). As such, SF graphs tend to offer a sounder representation of the viral transmission pathways than other network classes since superspreaders can be adequately modelled as the aforesaid hubs (Zenk et al., 2020; Kojaku et al., 2021). Consequently, this generation approach offers an appropriate benchmark for training control agents to handle more heterogeneous diffusion dynamics, such as that of COVID-19 (see Chapter 4).

- **Watts-Strogatz model:** Introduced by Duncan J. Watts and Steven Strogatz in 1998 (Watts and Strogatz, 1998), this is a widely-used rewiring model that produces graphs with several small-world (SW) properties, including short average path lengths and large clustering coefficients, which are highly representative of real interaction networks. The Watts-Strogatz method generates a ring of size N , then creates clusters by joining each node to its k_{nn} nearest neighbors, after which shortcuts are made by rewiring edges with probability p_r . Clusters have been shown to be an important transmission catalyst of the SARS-CoV-2 virus (Liu et al., 2020b), and thus modelling them may offer more realistic representations of the spreading mechanism.
- **Holme-Kim model:** Proposed by Holme and Kim (2002), this approach expects N , one integer m_{HK} , and a probability p_Δ as parameters, growing the graph by adding m_{HK} incident edges for each new node, with a chance p_Δ to construct a triangle from every new edge. This process creates networks that exhibit both SF and SW properties, making the Holme-Kim (HK) method an ideal instrument for capturing the spreading dynamics of respiratory diseases, such as COVID-19, since the generated graphs would typically feature both superspreaders and realistically-high clustering coefficients. This generation method is the focus of a detailed exploration in Chapter 3, where it is used to assess the impact of contact tracing strategies.

- **Stochastic-block model:** Developed by [Holland et al. \(1983\)](#), this is a rewiring model that partitions the nodes into disjoint blocks of sizes s , and then assign edges between vertices according to a parameter matrix of probabilities P that controls the intra- and inter-block connectivities. This method can be used to create flexible clusters, which have internal connections determined by the diagonal elements of P and external links governed by the off-diagonal elements of P , enabling a more realistic modelling of small communities that aids our demonstration of the advantages of more targeted public health interventions (for more details, refer to Section 4.4.4).

2.4 Network heuristics

Network problems often require finding satisfactory solutions in a timely and feasible manner. Network heuristics are methods that enable the discovery of such solutions by making use of sensible rules, derived from experience or intuition. They do not ensure optimality, but they frequently achieve sufficiency in a wide range of difficult problems. Within the broad spectrum of applications, we note finding appropriate strategies for testing ([Meirom et al., 2015](#)) and vaccinating against a pathogen ([Preciado et al., 2014](#)), routing data ([Tounsi, 2021](#)), or maximizing social influence ([Murata and Koga, 2018](#)), making a computer network more efficient or secure ([Sun et al., 2022](#)), and analyzing a graph's structure or function ([Gupta et al., 2016](#)). One can divide network heuristics into three categories, based on where they operate: node-level, edge-level and graph-level.

Heuristics acting at the node level focus on the properties or characteristics of individual vertices and how they affect the network structure or function. For example, node-level heuristics can measure the centrality of vertices, which reflects their importance or influence in the graph based on different criteria. They can also help us understand the roles and behaviors these nodes exhibit in the network. For instance, one possible application of these measures is to use them for ranking the vertices in a contact tracing network, and selecting the most 'influential' set for testing or vaccination prioritization. We show in Chapter 4 that this strategy can result in significantly higher containment rates than standard random samplers.

Commonly-used node-level heuristics include:

- **Degree centrality ([Freeman, 1978](#)):** The number (or fraction) of edges connected to the parameter node v , i.e. $c_D(G, v) = \text{deg}_G(v)$. Nodes with higher degree centrality are more influential or connected in the network.
- **Betweenness centrality ([Freeman, 1978](#)):** The number (or fraction) of shortest paths between any pair of vertices that pass through the node v , i.e.

$c_B(G, v) = \sum_{s \neq v \neq t} \frac{\sigma_G(s, t|v)}{\sigma_G(s, t)}$. Nodes with higher betweenness centrality are more central or bridging in the network.

- **Closeness centrality (Freeman, 1978):** The inverse of the average distance from a node v to all other vertices in the network, i.e. $c_C(G, v) = \frac{N-1}{\sum_{u=1}^{N-1} d_G(v, u)}$. Nodes with higher closeness centrality are more accessible or reachable in the network.
- **Katz centrality (Katz, 1953):** Computes the centrality of node v based on the centralities of its neighbors, i.e. $c_K(G, v) = \alpha_K \sum_u A_{v,u} c_K(G, u) + \beta_K$, where α_K and β_K are parameters. Nodes with higher Katz centrality are more connected or active in the network.
- **Eigenvector centrality (Bonacich, 1972):** A special case of Katz centrality where $\alpha_K = \lambda_{max}^{-1}$ and $\beta_K = 0$, with λ_{max} being the largest eigenvalue of the adjacency matrix A . Nodes with higher eigenvector centrality are more influential or prestigious in the network.
- **PageRank (Brin and Page, 1998):** A variant of the eigenvector centrality that accounts for the direction and quality of edges, combined with a random jump assumption. Can be obtained from the Katz centrality by setting $\beta_K = \frac{1-\alpha_K}{N}$ and replacing A with AD_{out}^{-1} , with α_K usually fixed at 0.85. Nodes with higher PageRank are more authoritative or popular in the network.
- **Clustering coefficient (Watts and Strogatz, 1998):** The fraction of possible triangles through the parameter node v that exist, i.e. $cc(G, v) = \frac{2T_G(v)}{deg_G(v)(deg_G(v)-1)}$, where $T_G(v)$ is the number of triangles through vertex v . Nodes with higher clustering coefficient are more embedded or clustered in the network.

In contrast, edge-level heuristics make decisions based on the attributes of the links in a network, often forming paths through the graph. For instance, a heuristic operating at the edge level might be applied to identify which paths are most vulnerable or valuable for an attacker to exploit (Chae et al., 2022). Classic approaches for finding minimum spanning trees, such as Prim's (Prim, 1957) and Kruskal's algorithms (Kruskal, 1956), or shortest paths $d_G(v, u)$, e.g. Dijkstra's algorithm (Dijkstra, 1959), can also be considered edge-level heuristics. Finally, graph-level heuristics consider the global structure or behavior of the network as a whole and how it affects its performance or resilience. They are frequently employed to measure the centralization of networks, which indicates the degree of concentration or inequality of ties among nodes (Padgett and Ansell, 1993; Simmons et al., 2018). These heuristics can also help optimize network design or configuration, based on some predefined criteria that reflect the desired properties or objectives. For example, the small-world networks we utilize here must by definition feature high graph-level clustering coefficients, which describe their average local cohesion, and low values for the average shortest-path length heuristic, which measures its global efficiency (Watts and Strogatz, 1998).

2.5 Machine learning with network data

Many of the challenging problems in the growing field of data science involve the study of relationships between structured objects. In bioinformatics, data concerning molecular networks, which serve as the backbone of molecular activity, is commonly used to predict a protein's function and interaction patterns, or even infer some of the functionalities that the complex cellular machinery performs (Sharan and Ideker, 2006). Manipulating this structure at more granular levels has also enabled scientists to find new drugs and treatments for common diseases (Bongini et al., 2021; Wu et al., 2021). Optimization problems in mathematics and physics can often be represented via networks, thus becoming solvable via nonspecific frameworks (Toenshoff et al., 2020; Cappart et al., 2021). Aggregating travel data for traffic prediction and estimating times of arrival naturally renders large datasets of connected structured objects, which ought to be leveraged using efficient graph models (Lange and Perez, 2020). In the context of social interaction networks, numerous such structures have been processed using machine learning in order to understand their patterns and evolutionary nature (Kumar et al., 2006), improve recommender systems (Fan et al., 2019), or predict likely links to emerge in the process (Dong et al., 2013; Yuan et al., 2019), among many other use cases.

Similar to the network heuristics discussed earlier, we can formulate different ML tasks involving network data based on the desired layer of abstraction: the node level (e.g. node classification), the edge level (e.g. link prediction) or the graph level (e.g. graph isomorphism). These tasks can be effectively solved by learning algorithms that can exploit both the network structure and the labels associated with the afferent nodes and edges. Graph kernels and graph neural networks are two of the most popular examples of such approaches. These techniques are discussed in greater detail in the subsequent sections, where their theoretical foundations and some of their notable variants are presented.

2.6 Kernel methods and graph kernels

Kernel methods are classic machine learning algorithms that learn from encoding the *similarities* between all pairs of datapoints, rather than using feature representations of the data directly. More formally, let $k: D \times D \mapsto \mathcal{R}$ be a continuous symmetric function which measures the similarity between all pairs of datapoints $(x_i, x_j) \in D \times D$. Then k is said to be a valid *kernel* on D if it can also act as the inner product of some arbitrary Hilbert space \mathcal{H} (Gretton, 2019), i.e. for a specific feature map $\phi: D \mapsto \mathcal{H}$, $k(x_i, x_j) = \langle \phi(x_i), \phi(x_j) \rangle_{\mathcal{H}}$. \mathcal{H} can in theory be infinitely dimensional, as is the case with RBF kernels of the form $k_{RBF}(x_i, x_j) = \exp\left(-\frac{\|x_i - x_j\|^2}{2\sigma^2}\right)$, but by

utilizing a mapping from the cross product of the original domain to the inner product of \mathcal{H} , one can leverage the expressivity of this higher-dimensional space without explicitly operating on its elements (i.e. the *kernel trick*) (Mohri et al., 2018). The downside of such an implicit approach is that all similarities need to be precomputed and stored in a *Gram matrix* K , with $K_{i,j} = k(x_i, x_j)$, operation which can be expensive for large datasets. By construction, k has to be positive semi-definite (p.s.d.), fulfilling Mercer’s condition, and due to Moore-Aronszajn theorem (Aronszajn, 1950), the corresponding \mathcal{H} is always unique, acting as a reproducing kernel Hilbert space (RKHS) of functions. The latter is essential for solving optimization problems via the kernel trick, yet learning with indefinite kernels is occasionally possible (Luss and d’Aspremont, 2008; Gu and Guo, 2012; Johansson and Dubhashi, 2015).

Currently, there exists a large spectrum of efficient kernelized methods suitable for most common learning objectives: supervised classification and regression (e.g. SVMs, Gaussian processes (Rasmussen, 2004)), statistical dependence testing or distance measurement (e.g. HSIC (Gretton et al., 2005)), unsupervised dimensionality reduction (e.g. kernel Principal Component Analysis – PCA (Schölkopf et al., 1997)), clustering (e.g. kernel k-means (Tzortzis and Likas, 2008)), density estimation (e.g. KDE (Silverman, 1986)), and many others. HSIC (or Hilbert-Schmidt Independence Criterion) is a kernel-based method to quantify the statistical dependence of two random variables. It can be used for feature selection and model fitting, as well as for explaining predictions, as we describe in Section 2.8 and empirically demonstrate in Appendix C.6. HSIC is based on the idea of embedding the joint and marginal distributions of the input random variables into an RKHS, and then computing the distance between the embeddings using the Hilbert-Schmidt norm (Gretton et al., 2005). Given two variables X and Y of size N , their afferent Gram matrices K_X, K_Y , and the centering matrix $H = I_N - \frac{1}{N}\mathbb{1}_N\mathbb{1}_N^T$, where I_N is the N -dimension identity matrix and $\mathbb{1}_N$ is the N -dimension vector of ones, the empirical approximation of HSIC for (X, Y) can be written as Eq 2.4, while its normalized variant takes the form of Eq 2.5.

$$\text{HSIC}(X, Y) = (N - 1)^{-2} \text{tr}(K_X H K_Y H) \quad (2.4)$$

$$\text{NHSIC}(X, Y) = \text{tr}(K_X H K_Y H) \quad (2.5)$$

The field of graph kernels has been rapidly evolving since the early 2000s, with many variants being developed and applied to different domains (Kriege et al., 2020). Like several network heuristics, these kernels can operate at the graph, the subcomponent, or the node level. Higher-order structures generally require the aggregation of lower-level (or *base*) kernels to form the final representations. Although this thesis does not directly employ them, we present for completeness several such aggregation mechanisms in Appendix A.1. Among the common examples of base p.s.d. kernels, we find the vertex-level label and feature kernels $k_{vl}(x, y) = k(l(x), l(y))$ and $k_{vf}(x, y) = k(f_v(x), f_v(y))$, where k can be a Kronecker delta, linear or RBF kernel, and

the pair-level shortest path kernel $k_{sp}((u, v), (w, z)) = \max(0, c - |d(u, v) - d(w, z)|)$ (Borgwardt and Kriegel, 2005), where c is set via cross-validation and $k(\cdot, \cdot) = 0$ if either path does not exist. More generally, however, one can construct custom p.s.d. Gram matrices from any symmetrical matrix M via the exponentiation: $K_\beta = e^{\beta M}$, with β a free parameter (Kondor and Lafferty, 2002). By differentiating the above w.r.t. β , a differential equation $\frac{d}{d\beta} K_\beta = M K_\beta$ with initial condition $K_\beta(0) = I_N$ arises. This indicates a gradual process that can transform K_β into a structure-aware kernel matrix, provided that the chosen generator M carries such information within. By setting $M = -L$, we can obtain a node-level *heat diffusion kernel*, whose name alludes to the similarity between the resulting differential form and the heat equation in Eq 2.3. The computation of the afferent Gram matrix requires the diagonalization of M as $H = Q\Lambda Q^{-1}$, followed by the exponentiation of the diagonal matrix of eigenvalues:

$$K_\beta = e^{\beta H} = Q e^{\beta \Lambda} Q^{-1} \quad (2.6)$$

As part of our initial evaluation of graph-based methods, partially reproduced in Appendix E.3, we train several SVMs with various configurations of the heat diffusion kernel. These experiments show that graph kernels, although competitive in some cases, tend to lack expressiveness in node-level tasks compared to recent GNN models. As such, the rest of this thesis will focus on the latter family, using the former only as a source of comparison, while outlining possibilities for further investigation.

2.7 Graph neural networks

Graph neural networks are state-of-the-art deep learning algorithms for graphs that can generate node embeddings based on the knowledge contained within the neighborhood of each vertex. The earliest studies in the space of GNNs defined the model as two functions: a transition function f_θ , and an output function o_θ (Gori et al., 2005; Scarselli et al., 2009). The first expresses the dependence between a node v and its neighborhood, while the latter controls how the final outputs are produced. Both of these are learned by neural networks using backpropagation through time, and can be expressed in vectorized form as follows:

$$X(t+1) = F_\theta(X(t), Y) \quad \hat{Y}(t+1) = O_\theta(X(t+1), Y),$$

where X encompasses the node representations, Y are some labels, and \hat{Y} are the predicted outputs. The two functions form the basis of what later came to be known as the *message-passing* paradigm (Gilmer et al., 2017), which quickly became dominant in the field due to its effectiveness and computational efficiency (Bronstein, 2022). By stacking multiple GNN layers, multiple rounds of message passing can be performed,

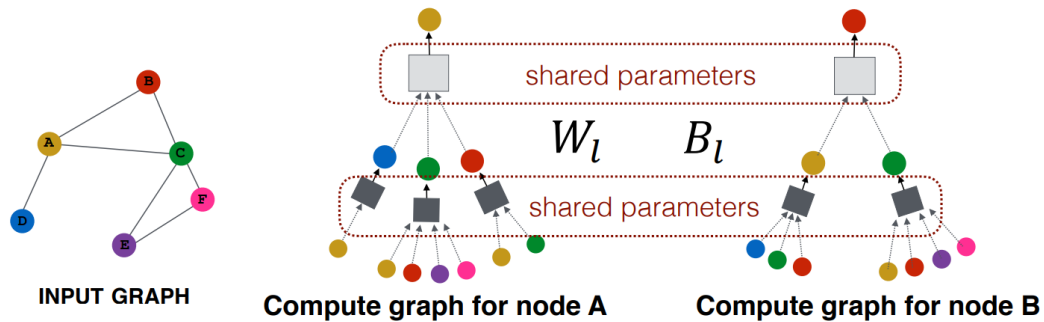


FIGURE 2.1: **Graph neural networks intuition.** Every node in the input graph determines a separate compute graph. The squares represent neural networks with shared weights W_l and biases towards the current representation B_l per each layer l . In this context, layers refer to hops within the network, i.e. for each vertex v , layer 0 are its node features x_v , layer 1 aggregates information from its immediate neighborhood (features from adjacent nodes and their common edges) according to W_1 and B_1 etc. Source: Leskovec (2021).

resembling WL updates. The node embeddings become increasingly more refined as the aggregated representations that reach them contain information stretching to more hops away (see Fig 2.1 for details). Adding nonlinearities, such as the standard ReLU function (Nair and Hinton, 2010), between layers is a common practice, and these typically improve the expressivity of the embeddings.

That being said, with increasing depths, conventional GNN architectures frequently suffer from feature *oversmoothing* (Hoang and Maehara, 2019; Oono and Suzuki, 2019) and *bottlenecks* (Alon and Yahav, 2022). As a result, unlike CNNs, GNNs often demand fewer layers for elevated performances. The first of these common issues occurs when the information gathered from numerous GNN layers makes node embeddings too smooth to be distinguished from each other, problem tightly connected to the equivalence between classic GNN architectures and low-pass filtering (Hoang and Maehara, 2019). The second issue is a natural consequence of having to fit the knowledge acquired from long-range hops into a fixed-dimensional vector, operation analogous to using CNNs with oversized receptive fields. This bottleneck effect can be partially mitigated by replacing the last GNN layer with a fully-adjacent (FA) layer that connects all nodes, thus easing information flow and preventing long-range signals from getting lost (Oono and Suzuki, 2019). Recently, models that learn more diverse information-flow pathways that are not constrained by the graph structure have also been proposed. These techniques facilitate building deep GNN architectures that significantly lessen oversmoothing and bottlenecks (e.g. GRAND by Chamberlain et al. (2021), described in more detail in Appendix A.2, or Neural Sheaf Diffusion by Bodnar et al. (2022)).

After the introduction of Graph Convolutional Networks (GCN) (Kipf and Welling, 2017), this model and its several variants have quickly become the norm in graph-based learning. The original GCN model is based on earlier work on spectral

graph convolutional neural networks by Bruna et al. (2014), but it circumvents the inefficient eigendecomposition of the graph Laplacian by utilizing a linear layer-wise model, as shown in Eq 2.7, where the parameters at layer l are Θ_l (this quantity encompasses both W_l and B_l from the illustration in Fig 2.1).

$$X^{(l+1)} = \hat{L}^{sym} X^{(l)} \Theta_l \quad (2.7)$$

Hamilton et al. (2018) bring forward a generalization of the formulation above that allows representations to be aggregated in more than a single way (see Eq 2.8), thus increasing the overall flexibility of the model. In the GraphSAGE equation below, $x_v^{(l)}$ is the representation of node v at layer l , \parallel is the concatenation operator, W_l and B_l are the weights of layer l , while AGG is a generalized aggregator function, which can be an *AveragePool*, a *SumPool*, a *MaxPool*, or even an arbitrary operator (e.g. Multi-Layer Perceptron – MLP, Long-Short Term Memory network – LSTM). The authors of GraphSAGE also propose a layer-wise neighborhood subsampling technique that makes learning on large networks more feasible.

$$x_v^{(l+1)} = (W_l \cdot AGG(\{x_u \mid u \in N_G(v)\}) \parallel B_l x_v^{(l)}) \quad (2.8)$$

The training performance on large-scale graphs can also benefit from simpler GNN models, which bypass the message passing routine by propagating the initial X through powers of graph operators (e.g. \hat{L}^{sym}), before utilizing an MLP to get the final outputs. This type of model is sometimes known as Graph-Augmented MLP (GA-MLP), and has been shown to be efficient at distinguishing most pairs of non-isomorphic graphs, albeit failing to be as expressive as classic message-passing GNNs (Chen et al., 2020). One of the most common examples of GA-MLPs is the SGC model described by Eq 2.9, which completely eliminates all nonlinearities (i.e. $X^{(L)}$ is the final output) in favor of using a single linear transformation of the node features X and a power-law formulation of the graph Laplacian (Wu et al., 2019). Similar in nature is the Graph Filter Neural Network (gfNN), which works by applying an MLP on top of $X^{(L)}$ (Hoang and Maehara, 2019). SIGN (Scalable Inception Graph Neural Networks) by Frasca et al. (2020) takes a slightly different approach by precomputing all the transforms $\hat{L}^i X$ up to a given power k (i.e. $i \leq k$), and then training a more efficient MLP by inputting these transformations. That being said, the standard GNNs tend to have a better expressivity, being the natural choice for smaller networks.

$$X^{(L)} = (\tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}})^k X \Theta \quad (2.9)$$

The general aggregation mechanism of GraphSAGE, which is common for most GNN architectures, equally weighs all the neighboring information unless an arbitrary operator is used. Motivated by the accomplishments of attention mechanisms in

natural language processing (Brown et al., 2020; Devlin et al., 2019; Vaswani et al., 2017) and computer vision (Bao et al., 2021; Khan et al., 2022), multiple variants of GNNs that use learned attention weights to compute weighted averages of the neighborhood representations have been suggested (Veličković et al., 2018; Dwivedi and Bresson, 2021; Brody et al., 2022). The Graph Attention Network (GAT) proposed by Veličković et al. (2018) leverages a variant of the multi-headed Bahdanau attention (Bahdanau et al., 2016) to obtain superior node embeddings. Its layer update rule is given by Eq 2.13, where α are node-level attention coefficients and W is a learned parameter matrix. In order to compute the coefficient $\alpha_{v,u}$ of node v to its neighbor u , a softmax over $N_G(v)$ is performed over the scores $e_{v,u}$ (refer to Eq 2.12). These scores can be calculated via either the original formulation (Eq 2.10) or the more expressive GATv2 variant (Eq 2.11) (Brody et al., 2022), where \parallel is the concatenation operator, $x_v^{(l)}$ and $x_u^{(l)}$ are current node representations, while a is a learned vector.

$$e_{v,u}^{GAT} = \text{LeakyReLU}(a^T [Wx_v^{(l)} \parallel Wx_u^{(l)}]) \quad (2.10)$$

$$e_{v,u}^{GATv2} = a^T \text{LeakyReLU}([W \parallel W][x_v^{(l)} \parallel x_u^{(l)}]) \quad (2.11)$$

$$\alpha_{v,u} = \text{softmax}_{u \in N_G(v)}(e_{v,u}) \quad (2.12)$$

$$x_v^{(l+1)} = \sigma(\sum_{u \in N_G(v)} \alpha_{v,u} Wx_u^{(l)}) \quad (2.13)$$

Aside from assessing their performance empirically, the expressive power of GNNs is often quantified using the WL test. The similarity of GNNs and the 1-WL update rule has initially been noted by Xu et al. (2019). However, the authors have demonstrated that the standard *AveragePool* and *MaxPool* neighborhood aggregators make models that are strictly less powerful than 1-WL. To overcome this limitation, they propose a new architecture, the Graph Isomorphism Network (GIN), that can be as powerful as the 1-WL procedure. A GIN layer takes the form of Eq 2.14.

$$X^{(l+1)} = \text{MLP}(\tilde{A}X^{(l)}) \quad (2.14)$$

Developing GNNs that are more expressive than 1-WL has been the focus of major efforts since the introduction of GIN, with many approaches now existing: breaking node symmetries by augmenting the features with random values (Sato et al., 2021), modifying the message passing rule (Beaini et al., 2021), or changing the input graph structure (e.g. k -GNN by Morris et al. (2020), Invariant Graph Networks (IGN) by Maron et al. (2019)). Though computationally intensive, the latter have been shown to exactly match the expressive power of k -WL. Surprisingly, however, their performance on benchmark datasets has consistently trailed behind more classic GNN architectures (Dwivedi et al., 2020; Bronstein, 2020). In contrast, the other approaches do not significantly affect runtimes and frequently enhance performances.

The preliminary analysis of GNN formulations presented in Appendix E.3 indicates a slight advantage for GAT and GCN on the standard Planetoid datasets: Cora, Citeseer and PubMed (Yang et al., 2016). However, since the performance differences among these tend to be small and inconsistent, the architectural experiments are repeated for the node ranking problem of Chapter 4. The latter reveal that a combination of the improved GATv2 and GIN renders superior results overall. The final training procedure we use is further enhanced by employing several of the techniques and good practices we detail above for mitigating oversmoothing and bottlenecks, such as low layer counts, final FA layers, and random node features. For more details on our epidemic control model, please refer to Section 4.3.5.

2.8 Explaining graph neural networks

A widely-used approach for explaining predictions in deep learning involves perturbing the inputs and fitting local explainable models to each data point and its corresponding perturbations. LIME (Ribeiro et al., 2016) and SHAP (Lundberg and Lee, 2017) are two popular examples of this methodology. Although the above are directly applicable to GNNs, they do not possess the capability to leverage the structural information of graph data or capture nonlinear relationships between the inputs and the outputs. To solve these limitations, GraphLIME has been proposed (Huang et al., 2022). GraphLIME replaces the local perturbations matrix with stacked node features selected from a node’s neighborhood, fitting multiple nonlinear interpretable models using HSIC (see Section 2.6). In Appendix C, we utilize GraphLIME to explain the testing decision taken by an RL epidemic control agent, making them accessible for auditing processes.

The GraphLIME criterion can be written as Eq 2.15, where $\beta_1, \dots, \beta_d \geq 0$ are the feature importances, d is the number of features, ρ is the Lasso coefficient, $\bar{\mathbb{K}}$ and $\bar{\mathbb{Y}}$ are the normalized doubly-centered Gram matrices describing the kernel space created from the feature inputs and the model’s outputs, respectively, while $\|\cdot\|_F$ designates the Frobenius norm. The equivalent form of the criterion, shown in Eq 2.16, where f_k is the k -th feature vector and y is the output vector, clearly isolates the (normalized) empirical HSIC values for each feature-feature and feature-output pair, highlighting the minimum redundancy maximum relevancy nature (Peng et al., 2005) of the GraphLIME solution. That is, $\forall m, n \in [1, d]$, $\text{NHSIC}(f_m, y) = \text{tr}(\bar{\mathbb{K}}^{(m)} \bar{\mathbb{Y}})$ is to be maximized in order to better capture the variation in the output space, while $\text{NHSIC}(f_m, f_n) = \text{tr}(\bar{\mathbb{K}}^{(m)} \bar{\mathbb{K}}^{(n)})$ gets minimized to reduce the informational redundancies between each feature pair.

$$\min_{\beta \in \mathbb{R}^d} \frac{1}{2} \left\| \bar{Y} - \sum_{m=1}^d \beta_m \mathbb{K}^{(m)} \right\|_F^2 + \rho \|\beta\|_1 \quad (2.15)$$

$$\min_{\beta \in \mathbb{R}^d} \frac{1}{2} \sum_{m,n=1}^d \beta_m \beta_n \text{NHSIC}(f_m, f_n) - \sum_{m=1}^d \beta_m \text{NHSIC}(f_m, y) + \frac{1}{2} + \rho \|\beta\|_1 \quad (2.16)$$

2.9 Applications of graph-based ML to diffusion processes

2.9.1 Identifying patient zero

One of the important open problems in network science is to identify the source(s) of a diffusion process, be it a viral outbreak, a rumor, or a cyberattack (Shah et al., 2020; Waniek et al., 2022). This problem has profound implications for various domains, including but not limited to public health, social media, and cybersecurity. By finding the origin of the diffusion, sometimes called ‘patient zero’, one can estimate the basic reproduction number R_0 of an epidemic, isolate and quarantine the most influential spreaders, or detect and counteract the malicious actors who disseminate false information. Most methods proposed for this task can trace back the diffusion from a single node, while others can handle multiple sources as well (Shelke and Attar, 2019). Typically, all these approaches rely on preexisting knowledge about the interaction graph’s structure. A notable exception, however, is the contact tracing procedure proposed by Waniek et al. (2022), which applies the model we introduce in Chapter 3 and a depth-based contact search that simulates tracing interviews to locate (or approximate the region of) ‘patient zero’.

When the full interaction network is known, GNNs can be employed to pinpoint the infection origin. For example, Shah et al. (2020) train several GCN models using epidemic simulation data derived from the IBMF model described in Eq 1.3 to identify the region of ‘patient zero’. Using one-hot representations of the states of each node at time t as feature vectors, the proposed GCNs can perform well in predicting the source node or the source neighborhood of the diffusion. By constraining the analysis of the dynamics to the early stages of the spread (i.e. $t \rightarrow 0$, $S_i \rightarrow 1$, $I_i \approx 1 - S_i$ and $R_i \approx 0$), the authors also define a time horizon after which the average infection probability is statistically no different than 1: $t_{max} \approx \frac{\log N}{\gamma(R_0 - 1)}$, where $R_0 = \frac{\beta\lambda_1}{\gamma}$ is the pathogen’s basic reproduction number. After t_{max} is passed, detecting ‘patient zero’ should in theory become considerably harder, irrespective of the method utilized. That being said, the authors fail to emphasize that this result does not necessarily hold beyond the early stages of an outbreak for all network types. For instance, in a line graph, t_{max} would not be enough to ensure (almost) all the vertices get infected. Nonetheless, the success of their GCN approach has significantly motivated our application of GNNs for epidemic control in Chapter 4.

2.9.2 Controlling the diffusion dynamics with GNNs and RL

As we highlight in Section 3.5, testing, tracing or vaccinating individuals at random may not be the most suitable public health strategy when only a limited number of resources are effectively available. Instead, it would be desirable to rank and target candidates based on their potential to generate large infection cascades. This problem is similar in nature to influence maximization, in that one attempts to maximize the gains of a generalized campaign by influencing a bounded set of vertices. The incomplete observability of the interaction graph environment represents the main challenge of this task, as it is often impossible to have access to all the node states when the periodic decisions have to be made. Given this partial view, an optimal policy needs to balance the infection/influence potential of each individual with the probability they become affected.

Meirom et al. (2021) introduce a novel actor-critic RL model to address the problem of ranking individuals for testing during an epidemic or maximizing influence in a social network. Their control framework for the first use case is diagrammed in Fig 2.2, while a more thorough outline is provided in the subsequent paragraphs. We extend this work in Chapter 4 and Appendix C to enable the inclusion of contact tracing and vaccination prioritization, proposing different visualization techniques for assessing the controlled diffusion process state, and improving runtimes by training the underlying deep learning module online. What is more, we apply the updated method on more realistic COVID-19 scenarios, demonstrating transferability between the IBMF model presented in Chapter 3 and a pure ABM setup. For more details regarding our epidemic control model, please refer to Section 4.3.

In this setting, we consider an social graph whose structure changes over time, while its set of vertices remains fixed: $G_t = (\mathcal{V}, E(G_t))$, where $\mathcal{V} = V(G_0) = \dots = V(G_T)$. The observable environment is organized as a temporal multi-graph with timestamped edges constructed by aggregating all the *known* $G_{t'}$ up to time t – denoted in this thesis with $G_{t' \leq t}$ – and the *already-observed* vertex states. The authors employ a single-layer GNN to track the diffusion dynamics in G_t (i.e. network D updates representations in the neighborhood of positively-tested nodes), and another k -layer GNN to account for the overall structure and long-range dependencies of $G_{t' \leq t}$ (i.e. network I updates representations in the entire connected component across time of a positively-tested vertex). This model formulation can readily utilize any fine-grained transmission probabilities $T_{v,u}(t)$ available by casting them as aggregation weights in D , $A_{v,u}^t = T_{v,u}(t)$, and edge features in I . Conversely, if $T_{v,u}(t)$ cannot be reliably estimated, the authors suggest dropping network D from the architecture. In addition to these probabilities, I uses the time delay of each edge in $G_{t' \leq t}$ as a separate feature.

At every time step t , both the *actor* and the *critic* are aware of: $G_{t' \leq t}$, a percentage of the node states at $t = 0$, as well as all the past test results. That is, for each $0 < t' < t$, the

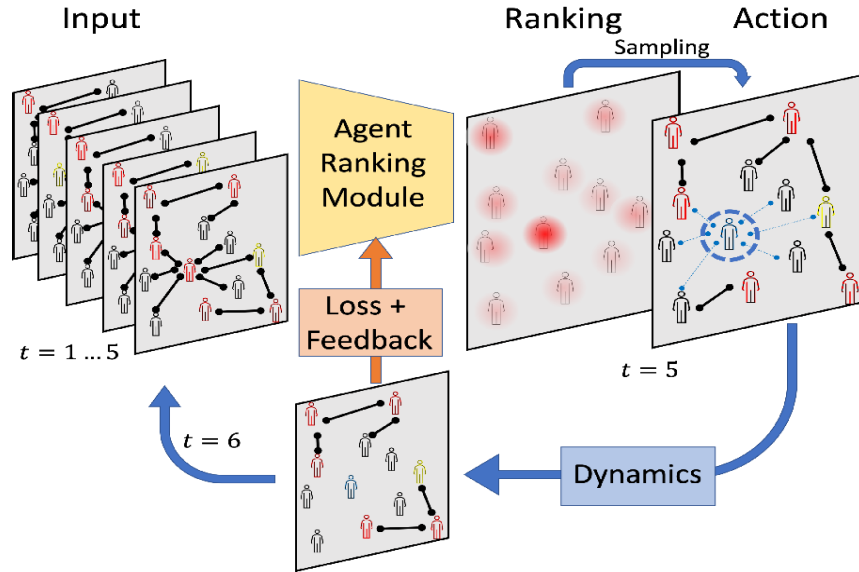


FIGURE 2.2: **Visualization of the RL+GNN approach.** In this toy example, the ranking module receives as input a sequence G_1, \dots, G_5 , constructs a temporal multi-graph from this, and feeds it to the actor ranking module. The latter outputs a probability distribution over actions, which ultimately gets sampled by the actor. The blue node is chosen in this scenario and then gets isolated, the resulting new state being aggregated to the original input to obtain the next timestamp $t = 6$. At the same time, the ranking module's parameters are updated according to a PPO loss involving the reward of this new state (based on the RL criterion above) and the critic's feedback. Source: [Meirom et al. \(2021\)](#).

actor already chose a set of vertices for testing $a(t')$, thus knowing at time t whether any $v \in a(t')$ has previously been infected. The RL agent is optimized using a Proximal Policy Optimization loss ([Schulman et al., 2017](#)), with actions sampled by the actor based on the vertex scores s_t output by the *ranking module*. The latter is effectively an MLP over the node features, the outputs of the 2 GNNs $d_t \in D(\cdot)$, $i_t \in I(\cdot)$, as well as the final hidden representation of the previous timestamp h_{t-1} . Every selected action at time t reveals the infection status of a given node v , while also detaching its edges in case of a positive test result. To encourage exploration, the score-to-probability function for sampling a_v from the action space A (i.e. actor choosing v for testing) is given by Eq 2.18, where s_t is obtained from Eq 2.17, while the episodic parameter ϵ_{ep} can be tweaked to smooth out the sampling distribution.

$$\begin{aligned} h_t &= \text{MLP}_{\theta_C}(h_{t-1}, f_v(t), d_t, i_t) \\ s_t &= \text{MLP}_{\theta_F}(h_t, h_{t-1}, f_v(t)) \end{aligned} \quad (2.17)$$

$$P(A = a_v) = \frac{x_{t,v}}{\sum_u x_{t,u}}, \text{ with } x_{t,v} = s_{t,v} - \min_u s_{t,u} + \epsilon_{ep} \quad (2.18)$$

The RL optimization criterion the authors utilize takes the form $\min \sum_t \gamma_r^t (I(t) + E(t))$, where the $I(t) + E(t)$ term represents the total number of infected up until time t in a SEI model (i.e. no recovery permitted), while γ_r is a

discount factor. Note that we denote the discount factor with a subscript to avoid confusion with recovery rate γ . The work employs a critic that has an almost identical architecture to the ranking model, but with a separate set of parameters and a single scalar output: the initial input of MLP_{θ_F} gets through an element-wise *MaxPool* operation on rows before being fed through its layers, thus outputting a system state value rather than a vector. In contrast, we suggest in Chapter 4 that the actor and the critic share the embedding layers, implementation detail that reduces training time considerably by design while still being effective.

Compared to more conventional supervised learning (SL) or network heuristics, the RL testing model of Meïrom et al. (2021) is shown to attain a remarkable performance w.r.t. the epidemic outcome. That being said, the authors neglect two important mechanisms of public health policy in their analysis: namely, the possibility of quarantining individuals based on contact tracing and conferring immunity through vaccination. Since their agents can isolate nodes or mark them as immune only after a positive test result has been recorded, their framework is not sufficiently lifelike. As previous studies have demonstrated (Abueg et al., 2021; Ferretti et al., 2020), and we confirm in Chapter 3, the effects of viral epidemics are significantly reduced when contact tracing is employed. At the same time, vaccination has been proven to be effective in preventing infections (Andrews et al., 2022; Bruxvoort et al., 2021; Lopez Bernal et al., 2021). Our approach in Chapter 4 aims to optimize the efforts for each of these interventions, limiting the extent of viral epidemics in a more realistic way, while striving for resource efficiency.

2.9.3 Further related work in graph-based ML for diffusion processes

Besides the above-mentioned works, we also acknowledge several related studies that have applied GNN architectures in distinct contexts involving diffusion processes: infection forecasting (Kapoor et al., 2020; Panagopoulos et al., 2021; Wang et al., 2020b), detection and screening (Yu et al., 2021), full population state estimation (Tomy et al., 2022), and vaccination policy optimization (Jhun, 2021). Among these applications, infection forecasting with GNNs has attracted a lot of attention after several companies have shared population mobility datasets to support research and policy making during the COVID-19 pandemic. This data is usually anonymized and aggregated at the level of geographical regions, and thus the modelled graphs in this context represent networks of communities or regions, rather than human interactions. For more information about this use case, please consult Appendix A.3.1.

Despite the potential of graph kernels for studying diffusion processes, this area has received little attention in comparison to GNNs. A rare example of such an effort, Oettershagen et al. (2021) develops several methods to transform temporal networks

into suitable representations for applying graph kernels to model and classify diffusion processes. These methods are discussed in more detail in Appendix A.3.2.

Chapter 3

Modelling Contact Tracing

3.1 Problem overview

The epidemic started in Wuhan, China by the SARS-CoV-2 virus has uncontrollably spread through communities all around the world, rapidly becoming a major global threat which, as of August 2023, was responsible for more than 769 million infected people and 6.9 million deaths (WHO, 2023). Prompted by the scale of this disease, cross-disciplinary teams started working against the clock to develop reliable pathogen spreading models that could be used to assess the effectiveness of different public health interventions. Since imposing a general lockdown proved economically unbearable for most countries, the attention significantly shifted to less restrictive yet partially successful measures, such as educating the public to socially distance, deploying large-scale testing, and quarantining contacts through various tracing mechanisms (Dighe et al., 2020). The latter proved rather challenging for the traditional interview-based approaches, mostly attributable to staffing problems and poor recollection on the part of the interviewees (Garry et al., 2021). As a result, digital alternatives were quickly sought after by several governments. These were successfully deployed in many states, most of them relying on either a Bluetooth solution, such as the Exposure Notification system (GAEN) (Google and Apple, 2020), or a geolocation-based software, similar to the Integrated Disease Surveillance Programme in India (Garg et al., 2020). That being said, the efficiency of these strategies remains largely dependent on the application adoption rates and the behavioral patterns of their userbase (i.e. self-isolation compliance, respecting the usage guidance, keeping the tracing device turned on etc.). Although some have suggested an application uptake of at least 50% would be needed at the population level to contain the epidemic (Ferretti, 2020), others showed via simulations that 60% would be enough to stop the spread without requiring further interventions (Hinch et al., 2020a). Nonetheless, the adoption levels generally quoted in the literature as 'sufficient' remain mostly unattainable due to privacy concerns and internet access

limitations. The picture gets even more intricate when the aforementioned behavioral issues are widespread in the active users' communities or if inadequate testing regimes and manual tracing procedures are employed.

Motivated by the limited evidence we have on the success of contact tracing in the face of such challenges, this chapter proposes four key advancements:

1. We build on the seminal work of [Farrahi et al. \(2014\)](#) and propose a new *multi-site* (or *individual-based*) *mean-field* model to analyze the joint impact of the aforesaid limitations on the efficacy of contact tracing, looking at varied scenarios involving various degrees of manual tracing network overlap Γ and digital tracing adoption r through the use of *triad* topologies. Our code is made publicly available to facilitate further experimentation (refer to Appendix B.1).
2. We propose separating the 'traced' status from the infection states, thus allowing for a node to get isolated at all times, unless it has reached a final state, i.e. recovered or dead, while also ensuring self-isolation can end due to non-compliance or term expiration, all without impacting an individual's normal disease progression. This modelling choice also makes our approach directly compatible with most of the other compartmental models.
3. We conduct a thorough analysis of numerous parameter settings and network topologies, including synthetic graphs generated using the methods we describe in Section 2.3, such as Erdős-Rényi ([Erdős and Rényi, 1959](#)), scale-free ([Holme and Kim, 2002](#); [Barabási, 2015](#)), and small-world networks ([Watts and Strogatz, 1998](#); [Wang and Chen, 2003](#)). Furthermore, we simulate similar outbreaks over a real social graph constructed from the Social Evolution dataset, which captures the proximity and location of several students who participated in a location tracking experiment ([Madan et al., 2012](#)). By comparing the simulated contact tracing outcomes in these two settings, we offer valuable insights into how various modelling parameters can affect the success of contact tracing.
4. We show that the effectiveness of contact tracing is influenced by not only the accuracy of the tracing network, but also by several other factors, such as the testing rates, tracing reliability, staffing and delays, public-health communiqués, isolation conformity, among others, and that finding an optimal configuration of these factors given a country's epidemiological situation is crucial for a swift viral containment. Additionally, we demonstrate that significant reductions in the peak of infections and the total number of deaths can be achieved even when lower uptakes are registered ($r < 0.4$) or the interviewing process misses several contacts ($\Gamma \leq 0.5$), given small tracing delays and the appropriate levels of testing and self-isolation compliance. Finally, we observe that the combined effects of manual and digital tracing can drive the effective reproduction number R_e below 1 even when neither is very efficient.

3.2 Related work

As discussed in Section 1.3, many compartmental formulations have been proposed for the SARS-CoV-2 virus – e.g. SIDARTHE (Giordano et al., 2020), SUQC (Zhao and Chen, 2020), SEIR with added quarantined and hospitalized compartments (He et al., 2020) etc. This chapter introduces a variation of the compartmental model proposed by the French National Institute of Health and Medical Research to study the impact of lockdown exit strategies on the spread of COVID-19 (Di Domenico et al., 2020).

That being said, individual-level public health interventions have typically been investigated using ABMs (Ferguson et al., 2020; Hinch et al., 2020b). A prominent example, Abueg et al. (2021) investigate the combined effects of manual tracing with digital solutions, at various application uptakes, via a rich yet scalable ABM fitted to mobility data from different counties in Washington. Their work provides one of the most robust modelling evidences for the efficacy of digital tracing, even when the latter operates at lower levels of application adoption.

At the intersection of these two paradigms lies the category of individual-based mean-field models, which combine the mathematical rigour and superior generalizability of EBMs with the ability to leverage locality information regarding every individual. Similar to certain ABMs, the infection spreads over a predefined network that can either be random (Rozhnova and Nunes, 2009) or inferred from real data (Farrahi et al., 2014), yet unlike ABMs, the dynamics are fully characterized by state transition equations. By stochastically simulating a similar model, Huerta and Tsimring explore the effects of contact tracing in a generic epidemic scenario (Huerta and Tsimring, 2002; Tsimring and Huerta, 2003). Farrahi et al. (2014) take their idea a step further by restricting the tracing propagation to a subset of the infection network, thus accounting for the inherently noisy nature of this process. Despite their powerful modelling capabilities, both approaches are limited by the underlying compartmental formulation they consider (i.e. SIRT), which makes several unrealistic assumptions that do not generalize to real viral diseases: inter alia, the recovery is conditioned on tracing, susceptibles can never be wrongfully isolated, and a traced person remains noninfectious for the full duration of the epidemic. Our modelling approach fixes these issues by separating the traced/isolated status from the infection state, thus allowing the ‘active’ nodes (i.e. not hospitalized, recovered or dead) to become traced or exit self-isolation after a certain amount of time without changing their afferent disease progression. Concurrently, this modification enables one to simulate the effects of contact tracing independently of the compartmental model utilized.

Branching process models, which have gained popularity in the last few years (Jacob, 2010; Lashari and Trapman, 2018), have also been applied to study the effects of manual contact identification used in conjunction with digital tracing solutions at various uptakes on the spread of COVID-19 (Plank et al., 2020). The corresponding

simulations show that an effective manual tracing needs to be coupled with an application uptake of at least 75% to achieve containment, although smaller adoption rates can decrease the reproduction number R_e if combined with other public health interventions. Our results are in accordance with the latter observation, but they also show that, given the right testing and tracing regimes (including an adequate self-isolation compliance), lower and achievable adoption levels are often enough to significantly reduce the viral diffusion, subject to the network's connectivity patterns.

3.3 Methodology

3.3.1 Compartmental model outline

Given the increasing evidence that basic SIR frameworks are unable to capture the complexity and heterogeneity of SARS-CoV-2 epidemics (Moein et al., 2021), we introduce a novel compartmental formulation that accounts for several of its particular features. We show a representation of this model in Fig 3.1, where the state transitions are labelled with their corresponding time-dependent probabilities. The model attains a final configuration when all non-susceptible nodes become either recovered (R) or dead (D). A description of the epidemic parameters utilized, together with the values we consider for each of them, can be consulted in Table 3.1.

3.3.2 Network propagation mechanism

We use a predefined network for the infection spread, and a subnetwork for each type of contact tracing: one for the manual and another for the digital process. This method enables us to simulate either one tracing strategy alone, called *dual* topology – Fig 3.2, or both concomitantly, called *triad* topology – Fig 3.3. Connected vertices in the true infection graph are to be considered 'close contacts', as previously defined by authoritative institutions like the CDC (CDC, 2020).

The tracing graphs represent subset views of the true contacts network, where missing edges correspond to application misuse in the digital setting or contacts not recalled in the manual interviewing process, while isolated vertices account for individuals that never run a government's digital solution or are effectively unreachable. Be that as it may, people can at times overestimate the number or the duration of their social interactions (Mastrandrea et al., 2015), and thus it is possible that tracers are occasionally pursuing erroneous links. Even though our model can simulate 'false' contacts, similarly to Farrahi et al. (2014), we consider their occurrence quite rare during a global pandemic (and thus negligible), since the public health staff tends to be better trained, while the general public pays more attention to their interactions.

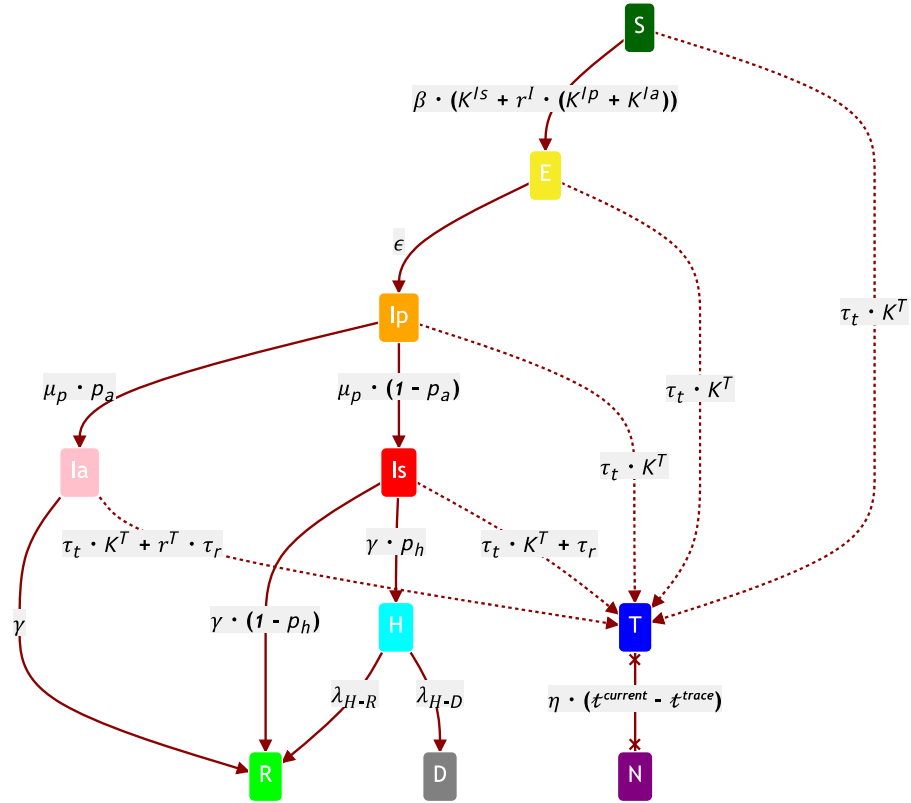


FIGURE 3.1: **The SEIR-T Compartmental Model for COVID-19.** Each node in the interaction network has 2 allocated variables: an infection state and a tracing status. The infection states from top to bottom are: *S* - susceptible; *E* - exposed but not infectious; *I_p* - infectious, presymptomatic; *I_a* - infectious, asymptomatic; *I_s* - infectious, symptomatic; *H* - hospitalized; *R* - recovered/removed; *D* - dead. At any point in time, a node’s tracing status can either be *T* (traced and isolated) or *N* (not traced/isolated or non-compliant). Each state transition has a certain time-dependent probability $p_{s_1 \rightarrow s_2}$; the edge labels here represent both $\frac{p_{s_1 \rightarrow s_2}}{\Delta t}$, and the λ rate of the corresponding exponential to sample from in the continuous-time simulations.




FIGURE 3.2: **Final state of an epidemic simulation over a dual topology.** Infection spreads with respect to the neighborhoods of the first network (here a SF graph); the second network represents a digital tracing view at uptake $r = 0.5$.

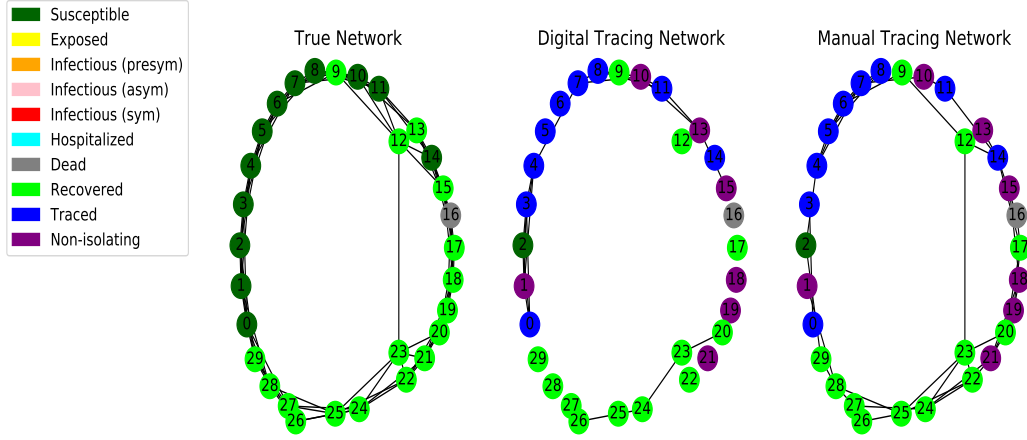


FIGURE 3.3: **Final state of an epidemic simulation over a triad topology.** Infection spreads with respect to the neighborhoods of the first network (here a SW graph); the second network represents a digital tracing view at uptake $r = 0.5$, while the third a manual tracing view with overlap $\Gamma = 0.5$.

TABLE 3.1: **Compartmental model parameters used in our simulations.**

Parameter	Value(s) used	Description
β	0.0791	Transmission rate corresponding to $R_0 = 3.18$. Obtained via maximum likelihood estimation performed on hospitalization data by Di Domenico et al. (2020).
K^X	\mathbb{R}	Function mapping nodes to the total number/weight of connections to neighboring nodes in state $X \in \{I_p, I_a, I_s, T\}$ for a given network.
r^I	0.5	Relative infectiousness of I_p and I_a compared to I_s . This is still disputed: ≈ 0.5 according to Grassly et al. (2020) and Di Domenico et al. (2020), but weak evidence as per McEvoy et al. (2021).
ϵ^{-1}	3.7	Latency period, measured in days. Source: Di Domenico et al. (2020).
p_a	0.2/0.5	Probability of being asymptomatic. Disputed value: 0.2 used by Di Domenico et al. (2020) and Mizumoto et al. (2020), but 0.5 according to Keller et al. (2020) or Oran and Topol (2020).
μ_p^{-1}	1.5	Presymptomatic period, measured in days. Source: Ferretti et al. (2020).
p_h	0.1/0.2	Probability of being hospitalized for <i>adults</i> and <i>seniors</i> (assumed almost 0 for children). Equivalent to p_{ss} in Di Domenico et al. (2020).
γ^{-1}	2.3	Infectious period considering the mean generation time 6.6 days. Source: Di Domenico et al. (2020).
λ_{H-R}	0.083/0.033	Daily rate of recovery for <i>adults</i> and <i>seniors</i> (assumed almost 0 for children). Source: Etalab (2020).
λ_{H-D}	.0031/.0155	Daily rate of deaths for <i>adults</i> and <i>seniors</i> (assumed almost 0 for children). Source: Etalab (2020).
τ_t	[0-0.5]	Contact tracing rate. Encompasses multiple related phenomena: the tracing latency/efficiency due to staffing/server reliability, depending on the type of tracing; the likelihood of remaining isolated given the number of traced neighbors. Ranges from no tracing (0) to every 2 days on average (0.5).
τ_r	(0-0.5]	Testing / Random tracing rate. Ranges from almost no testing (0.001) to every 2 days on average (0.5).
r^T	0.8	Relative probability for I_a to be tested positive (against I_s). Assume testing E and I_p rarely happens or results in false negatives most of the time.
η	0/0.001	Non-compliance / Self-isolation exit rate. Scaled by the time elapsed since the beginning of isolation, $t^{\text{current}} - t^{\text{traced}}$.

We control the subsetting of the infection graph via two interlinked parameters: the *degree of overlap* $\Gamma = \frac{K-Z_{rem}}{K}$ and the *uptake rate* $r = \frac{N-N_{utn}}{N}$, variables which ultimately determine the values of N_{utn} and N_{ute} (refer to Eq 3.1). To be more explicit, the input Γ and the infection network's mean degree K are utilized to calculate Z_{rem} , the average number of edges per node to get hidden from a tracing view. The latter effectively corresponds to marking as untraceable $N_{ute} = \frac{NZ_{rem}}{2}$ of the edges in the interaction graph. Similarly, the selected r and the total number of nodes N are used to establish how many vertices are to be made completely untraceable in a particular tracing subnetwork: $N_{utn} = N(1 - r)$. This chapter showcases simulations in which the first of these two parameters describes the accuracy of manual tracing, whereas the second quantifies the adoption of a digital solution. That being said, our model supports exploring more complex scenarios, where both the overlap and the uptake can be varied for a single tracing view. A full description of the network-related variables involved in our modelling procedure can be consulted in Table 3.2.

$$N_{utn} = N(1 - r) \qquad N_{ute} = \frac{NZ_{rem}}{2} = \frac{NK(1 - \Gamma)}{2} \qquad (3.1)$$

Throughout our experiments, we assume a 'traced' individual (i.e. having tracing status T) automatically enters self-isolation, thus making it impossible to infect others or get infected unless it reaches the 'non-isolating' status N . This can happen either as recommended by authorities, after 14 days, or wrongfully, with a probability of η scaled by the time elapsed since isolating. In addition, we presume that a node's probability to get infected proportionately increases with the amount of infectious neighbors it has in the contacts network, while the likelihood of being traced and compliant with self-isolation recommendations is directly proportional to the number of adjacent T nodes it has within each of the tracing subnetworks.

TABLE 3.2: Network-generation parameters.

Parameter	Value(s) used	Description
N	1N	Population size of the infection network.
K	10/20	Average degree of the infection network. This is only configurable in the ER experiments.
m_{HK}	10	Random edges to add for each new node in the Holme-Kim networks (Holme and Kim, 2002).
p_{Δ}	0.2	Probability of completing a triangle after adding a random edge.
Γ	[.1, 1]	Degree of overlap between infection network and a tracing subgraph. Used to calculate Z_{rem} , which in turn reveals N_{ute} (number of untraceable links).
r	[.1, 1]	Uptake rate (between infection network and a tracing subgraph). Used to calculate N_{utn} (number of untraceable nodes).

3.3.3 Simulation overview

The baseline simulations in this chapter are performed on Erdős–Rényi random graphs with varying population sizes and average degrees. Despite the limitations of this class of networks, as discussed in Section 2.3, [Tsimring and Huerta \(2003\)](#) argue that the epidemic dynamics of their model remains “qualitatively similar” between ER realizations and the empirically-based class of SW graphs, since both have well-defined epidemic thresholds. Given that we implement tracing in an analogous manner, we can also anticipate similar levels of tracing efficiency for both graph generation approaches. At the same time, we consider our modelling choice of sampling nodes for contact tracing according to their count of isolated neighbors to be similar in essence to the frequency-based tracing procedure proposed by [Kojaku et al. \(2021\)](#), and thus we expect superspreaders within SF graphs to be preferentially targeted by our control framework, avoiding potential imbalances the corresponding hubs may cause. It is, therefore, sensible to assume the simulation mechanism we employ remains suitable for assessing the effects of tracing over a broader range of network models, beyond ER. We provide further evidence for this in a series of additional experiments on more realistic SF and SW networks, presented in Sections 3.4.3 and 3.4.5, offering a more comprehensive analysis of our model.

In contrast to the above, Section 3.4.6 investigates the effects of digital and manual tracing in a viral outbreak simulated over a real social network, representing a tightly-connected community of 74 students and graduates from MIT who agreed to have their location and interactions monitored via WLAN and Bluetooth scans over an entire academic year (detailed exploratory analyses of the dataset can be examined in [Madan et al. \(2012\)](#) and [Farrahi et al. \(2014\)](#)). In our simulations, this dynamic network changes daily over a period of 31 weeks, its links being *weighted* by the aggregated number of Bluetooth proximities recorded between their corresponding corner points on each particular day. In the static settings presented thus far, K^X represents a function mapping nodes to the total number of neighbors in state $X \in \{I_p, I_a, I_s, T\}$ (see Fig 3.1). To account for the dynamic weights present in this dataset, however, all K^X terms get replaced by a time-dependent function K_t^X , given by Eq 3.2, where $K_{norm} = 10$ is a normalization factor that ensures the average function value remains above 1, $w_t^X(n)$ is the sum of edge weights incoming from those neighbors of node n which are in state X at time t , while $\langle W \rangle$ is the overall static average weight. The latter represents an average over days of the average total weight per node, calculated using Eq 3.3, where $D = 216$ is the number of days within the considered 31-week period, $N = 74$ is the number of nodes for which we have contact data, and $w_t(n)$ is the *total* weighted degree of node n at time t , irrespective of state.

$$K_t^X(n) = \frac{K_{norm} \cdot w_t^X(n)}{\langle W \rangle} \quad (3.2)$$

$$\langle W \rangle = \frac{\sum_t^D \sum_n^N w_t(n)}{D \cdot N} \quad (3.3)$$

In this work, the time intervals between two state changes of the same type (i.e. matching source and target compartments) are assumed to form an exponential distribution, with the λ rate equal to the corresponding transition label displayed in Fig 3.1. Choosing this distribution for timing the infection propagation, in particular, keeps our approach in line with many previous epidemiological works relying on compartmental formulations (Kröger and Schlickeiser, 2020; Farrahi et al., 2014; Ma, 2020), while also being in accordance with the findings of different cohort studies involving wearable tracking devices that have reported roughly-exponential decays in their participants' histogram of interactions (Stehlé et al., 2011; Farrahi et al., 2015). We acknowledge similar cohort studies have found heavier-tailed distributions based on power laws to be more compatible with the time intervals between successive interactions, citing the bursty nature of social dynamics as the determining factor (Cattuto et al., 2010; Starnini et al., 2013). However, the corresponding data fit has often proved imperfect, while extensive comparisons against exponentials have rarely been performed. In the epidemiological setting, more realistic and flexible distributions for the infection waiting times have been suggested, such as Erlang, a special case of the Gamma distribution (Krylova and Earn, 2013; Liu et al., 2020a), or Weibull, a generalization of the latter (Streftaris and Gibson, 2012; Lipsitch et al., 2003; Van Mieghem and Liu, 2019). These can effectively model the non-Markovian nature of some epidemics, but they typically require more parameters to be estimated, making the analysis of the diffusion process more difficult. That being said, the more convenient exponentials have been shown to provide a particularly good fit to epidemiological data when the pathogen's mean generation time is correctly fixed (Krylova and Earn, 2013) or the mean infection duration is smaller (Vergu et al., 2010). Both of these conditions should hold for the SARS-CoV-2 outbreaks we investigate.

For efficiency, we simulate the COVID-19 epidemics using Gillespie's algorithm (Gillespie, 1977), which has been shown to be stochastically exact to and faster than the Monte Carlo method (MC) for both static and dynamic network-based diffusion processes (Vestergaard and Génois, 2015). Compared to a continuous-time MC simulation, which entails sampling the next transition for all the possible state changes, discarding all but the most 'recent' event (Farrahi et al., 2015), Gillespie's procedure directly draws the time elapsed until the next transition and identifies the state change most likely to have taken place within that period. A detailed pseudocode for our event sampling procedure is provided in Algorithm 1.

Algorithm 1 Pseudocode for event sampling. Minimum-time event is sampled directly. Rates get recalculated only for last updated node and its neighbors.

```

1: global variables
2:    $t$                                 ▷ Current global time of the simulation
3:    $n_u$                                 ▷ Node ID of the last transition update
4:    $N_I$                                 ▷ The infection contacts network
5:    $N_T$                                 ▷ List of tracing networks (IDs 0 and 1)
6:    $I_f$                                 ▷ Getter of possible infection propensity functions
7:    $T_f$                                 ▷ Getter of possible tracing propensity functions
8:    $\lambda$                              ▷ Dict mapping nodes to possible transitions (rate,  $S_{to}$ )
9: end global variables
10: procedure SAMPLE_NEXT_EVENT
11:   ▷ Collect nodes that need rates updating: last and neighbors
12:    $update\_nodes \leftarrow N_I.neighbors(n_u) \cup \{n_u\}$ 
13:   ▷ Update  $\lambda$  for  $update\_nodes$ 
14:   for each  $n \in update\_nodes$  do
15:      $\lambda.pop(n)$                     ▷ Invalidate node in the dict of rates
16:      $S_I \leftarrow N_I.state(n)$         ▷ Get current infection state
17:      $S_T \leftarrow N_T.state(n)$         ▷ Get current tracing state
18:     ▷ Update rates based on node state and the true network
19:     for each  $(f_r, S_{to}) \in I_f(S_I)$  do                                ▷  $f_r$  is a rate function
20:       ▷ rate based on neighborhood of  $n$  in  $N_I$  and a scalar
21:        $rate \leftarrow f_r(N_I, n)$ 
22:       ▷ Add (rate, state) to the  $\lambda$  dictionary of rates
23:        $\lambda(n) \leftarrow \lambda(n) \cup \{(rate, S_{to})\}$ 
24:     end for
25:     ▷ Update rates based on tracing state and 2 tracing nets
26:     for each  $(f_r, S_{to}) \in T_f(S_T)$  do
27:       ▷  $rate_i$  depends on tracing neighborhood of  $N_T(i)$ 
28:        $rate_0, rate_1 \leftarrow f_r(N_T(0), n), f_r(N_T(1), n)$ 
29:        $\lambda(n) \leftarrow \lambda(n) \cup \{(rate_0, S_{to}), (rate_1, S_{to})\}$ 
30:     end for
31:   end for
32:   ▷ Convert  $\lambda$  into 2 lists related by the map: rate < - > (node,  $S_{to}$ )
33:    $rs, nts \leftarrow convert(\lambda)$ 
34:   ▷ Gillespie sampling of the minimum exponential time
35:    $\lambda_{min} \leftarrow \sum_{rate \in rs} (rate)$ 
36:    $t_{min} \sim Exp(\lambda_{min})$ 
37:   ▷ Next time point is the current  $t$  + the minimum sampled
38:    $t_u \leftarrow t + t_{min}$ 
39:   ▷ Categorical sampling of the actual transition
40:    $id_u \sim P(X = k) = \frac{rs(k)}{\lambda_{min}}$                                 ▷ Base rate over sum of rates
41:    $n_u, S_u = nts(id_u)$                                                 ▷ Change last-updated node
42:    $S_{past} = N_I.state(n_u)$ 
43:   ▷ Create Event  $e$  dictionary which shall be used to update the network states,
   neighborhood counts and the epidemic stats
44:    $e = \{("id", n_u), ("from", S_{past}), ("to", S_u), ("time", t_u)\}$ 
45:   return  $e$ 
46: end procedure

```

3.3.4 Metrics under consideration

Aside from scrutinizing the number of individuals in each compartment over time (please also refer to Appendix B for more such evaluations), we assess the efficacy of different contact tracing strategies ($C_{\theta;\tau_t}$, under different τ_t) by looking at their achieved *peak suppression* P_{sup} throughout all our simulations, thus comparing them against the corresponding no-tracing scenario ($C_{\theta;\tau_t=0}$) in which all parameters θ but τ_t are left unchanged. Mathematically, this can be expressed through Eq 3.4, where I_{max} is a function mapping parameter configurations C_θ to the average peak of infections recorded across multiple runs.

$$P_{sup} = I_{max}(C_{\theta;\tau_t=0}) - I_{max}(C_{\theta;\tau_t}) \quad (3.4)$$

Since the inception of the COVID-19 pandemic, the majority of the literature on epidemiological modelling and public-health messages alike have analyzed different nonpharmaceutical interventions in relation with their impact on R_e , the effective reproduction number (Kajitani and Hatayama, 2021; Di Domenico et al., 2020). For the more realistic scenarios we explore in Sections 3.4.5 and 3.4.6, we also estimate the R_e value after $t = 7$ days since t_0 . To do so, we input the recorded exponential growth rate λ to Eq 3.5, thus following the Wallinga and Lipsitch methodology (Wallinga and Lipsitch, 2007). The generation time distribution for our simulated SARS-CoV-2 epidemics is assumed to be $Gamma(\alpha = 1.87, \beta = 0.28)$ (Cereda et al., 2020), its moment-generating function being denoted with $M(\cdot)$. To calculate λ from the incidence rate $c(t)$ recorded within time window $[t_0, t_0 + t]$, we use Eq 3.6 together with the initial number of infected $c(t_0)$.

$$R = \frac{1}{M(-\lambda)} = \frac{1}{(1 - \frac{\lambda}{\beta})^{-\alpha}} = (1 + \frac{\lambda}{\beta})^\alpha \quad (3.5)$$

$$c(t) = c(t_0)e^{\lambda t} \implies \lambda = \frac{\log c(t) - \log c(t_0)}{t} \quad (3.6)$$

3.4 Results and discussion

3.4.1 Variation induced by population size

Our initial simulations using ER graphs have suggested the degree of variability across runs scales with the number of nodes. To confirm this, we design an experiment in which we vary the population size, $N \in \{200, 500, 1000, 2000, 5000, 10000, 20000\}$, while keeping the other parameters fixed at: average degree $K = 10$, *dual network* tracing with uptake $r = 0.5$ (overlap Γ implicitly derived), asymptomatic probability $p_a = 0.2$, contact tracing rate $\tau_r = 0.1$, and testing rate $\tau_t = 0.1$, with one infectious

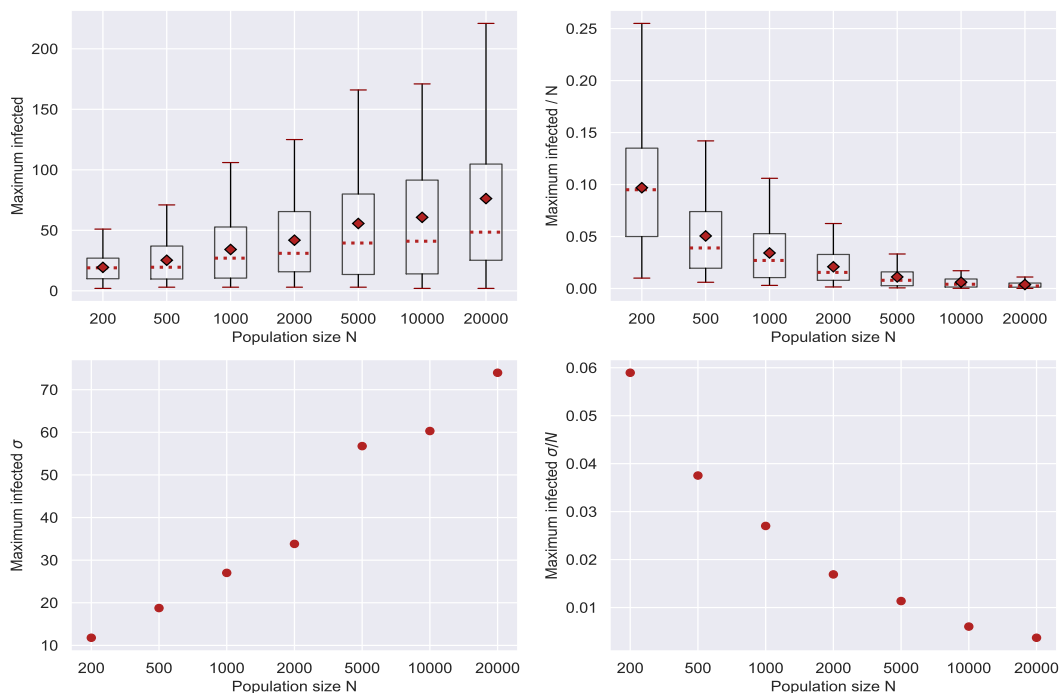


FIGURE 3.4: **Uncertainty of simulation results w.r.t. the infection peak.** Averages from 80-100 runs involving different population sizes, $K = 10$, $\tau_t = \tau_r = 0.1$, $p_a = 0.2$. On top, boxplots with quartiles represented via whiskers, medians via orange lines, and averages via green triangles; standard deviations given below. The left-hand side shows absolute values, whereas on the right all variation levels are scaled down by N .

individual set for time t_0 . We note that, for $\tau_r = \tau_t > \beta$, contact tracing is expected to engulf the infection percolation in the limit. However, by choosing an uptake value considerably smaller than 1, we ensure our variance analysis remains significant since many of the randomly-generated tracing views end up producing a much slower discontinuation of the otherwise quickly-contained infection cascades. This results in a high probability for finite outbreaks to occur during the early stages of the simulations (i.e. above the epidemic threshold for enough time).

The statistics in Fig 3.4 represent averages over several simulations conducted with each of the 10 different network initializations picked by a random sampler, filtering out those iterations which registered less than three overall infected (for a total of 80-100 simulations overall per each value of N). This data confirms the absolute variance in peaks of infection increases as the network expands, aspect which can be explained by the growing difference between early-stopped and full-blown outbreaks. As the size increases, however, the uncertainty in estimating the relative percentage of these maximal points decreases accordingly (following $\approx \frac{1}{\sqrt{N}}$). Consequently, a choice of $N = 1000$ yields a reasonable standard deviation of nearly 3% of N , while $N = 10000$ produces a smaller variability of $< 1\%$ across runs. As such, $N \in [1000, 10000]$ represents a statistically precise range for our model, given a randomly-mixed population, while also being computationally more efficient than

larger values. To ensure representativeness, we use both extremes in our experiments, with the remark that we adjust for the corresponding difference in variances by simulating 7 different networks with 15 random seeds each for $N = 10000$, but 50 networks with 15 seeds for $N = 1000$. We note this precision result may not hold for structured populations, but we adopt these setups in our SF/SW experiments as well for consistency.

3.4.2 Initial exploration of tracing overlap in larger populations

Going forward, we want to assess the effect of varying a tracing network's accuracy (i.e. overlap) in an outbreak involving a large community of $N = 10000$ individuals. To achieve this, we use the following parameter configuration: an average degree $K = 10$, *dual network* tracing with overlap $\Gamma \in \{0.1, 0.2, \dots, 1\}$ (uptake r is implicitly derived), asymptomatic probability $p_a = 0.2$, a tracing rate $\tau_t \in \{0.01, 0.04, 0.07, 0.1\}$, a testing rate $\tau_r \in \{.001, 0.01, 0.04, 0.07, 0.1\}$, and a non-compliance rate $\eta = 0$ (assuming everybody self-isolates until they are no longer infectious), with a single I_p node sampled at time t_0 . The epidemic outcomes get averaged over 105 runs, as aforesaid.

Fig 3.5 shows that a sub-optimal test rate, such as $\tau_r = 0.001$, leads to inconclusive results, where the variance induced by the stochasticity of the process shadows any benefit obtained through contact tracing. With better testing, clearer patterns start to emerge: The higher the contact tracing rate, the better the peak suppression is and the faster it gets approached (according to Fig 3.6). As τ_r becomes even more effective, smaller tracing network overlaps are needed to swiftly reduce that maximum point.

Looking at the tracing rate, a moderate value of $\tau_t \in \{0.04, 0.07\}$ achieves a delay in the peak for smaller Γ , but this can occasionally lead to a prolonged epidemic, especially for overlaps in the 'noise' region, such as $\Gamma = 0.11$, since the initially-uninfected areas may get incorrectly traced, so the epidemic has the chance to gain momentum once those individuals exit self-isolation. In contrast, noticeable reductions with no such side effect can be observed for $\Gamma \geq 0.5$. On the other hand, a small value of $\tau_t = 0.01$ seems unable to produce any positive outcome. In real life, the latter scenario would occur if the tracing programme was very slow, missing too many contacts as a consequence, or if the digital contacts application failed to promptly notify many of its active users. Another noteworthy occurrence in Fig 3.6 is the bimodality of some curves. This effect has been previously noted for larger tracing rates and overlaps (Farrahi et al., 2014), being a rare artefact of peak reductions that happen too rapidly and cannot be further sustained by a $\tau_t < \beta$.

Aside from outlining the effects of different testing strategies and tracing network overlaps, this experiment also hints at which parts of τ_r 's and τ_t 's parameter spaces are more relevant for exploration. To aid our search, we plot heatmaps of the achieved

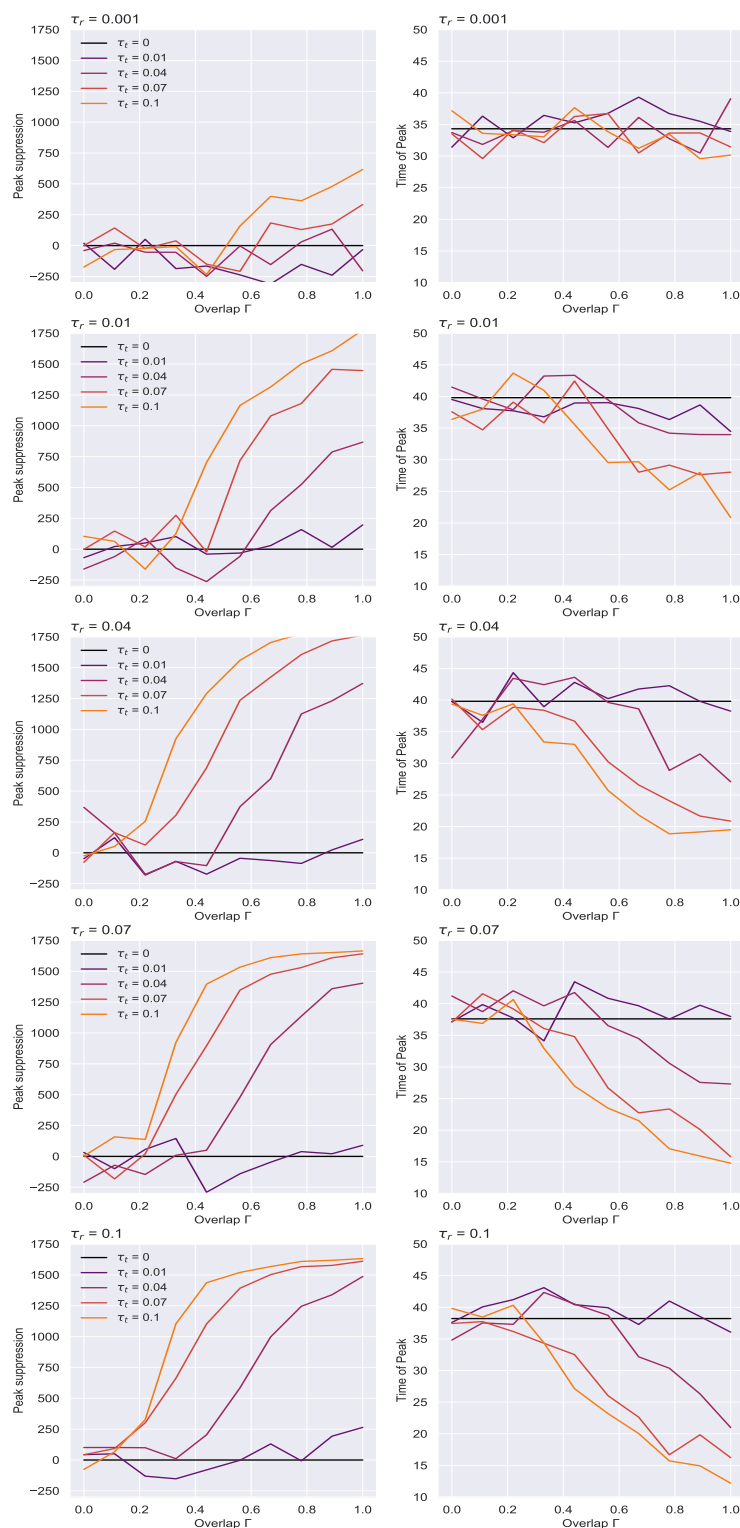


FIGURE 3.5: ER network – Peak suppression (left) and the time of peak (right) at various tracing network overlaps. Values are averaged over 105 runs, representing results for $N = 10000$, $K = 10$, $p_a = 0.2$. The suppression is calculated by subtracting the average maximal infected point given by each parameter configuration from the average point obtained with no tracing ($\tau_t = 0$, colored in black). Excepting $\tau_r = 0.001$ and $\tau_t = 0.01$, which produce inconclusive and noisy results, the effectiveness of an epidemic containment strategy expectedly scales with the testing and the tracing rates.

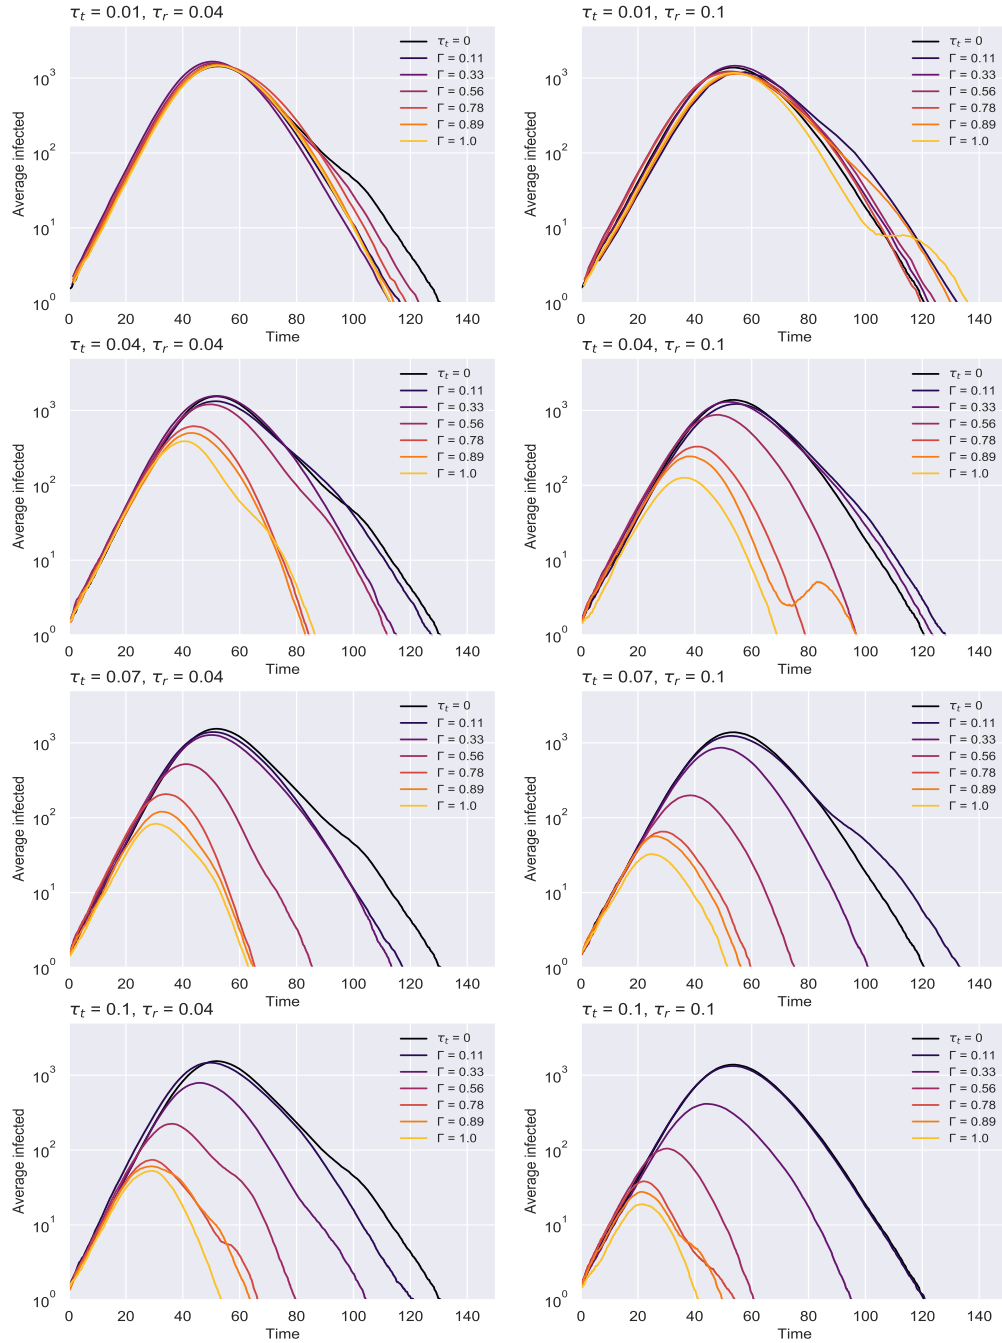


FIGURE 3.6: ER network – Epidemic evolution over time given a less efficient (left) and a more effective (right) testing regimes. Results averaged over 105 simulations, obtained for $N = 10000$, $K = 10$, $p_a = 0.2$. As the contact tracing rate increases, the accuracy of the network given by Γ becomes more important for ‘flattening’ the curves. The case with no contact tracing ($\tau_t = 0$) is colored in black.

peak suppression for different levels of overlap (see Fig 3.7), and observe, as a result, that significant outcomes (i.e. distinguishable from simulation noise for $\Gamma \approx 0.5$ and beyond) are obtained when $\tau_t \geq 0.1$ and $\tau_r \geq 0.04$, while values ≥ 0.1 for both parameters should fall within the ‘adequate’ region of a large spectrum of Γ values.

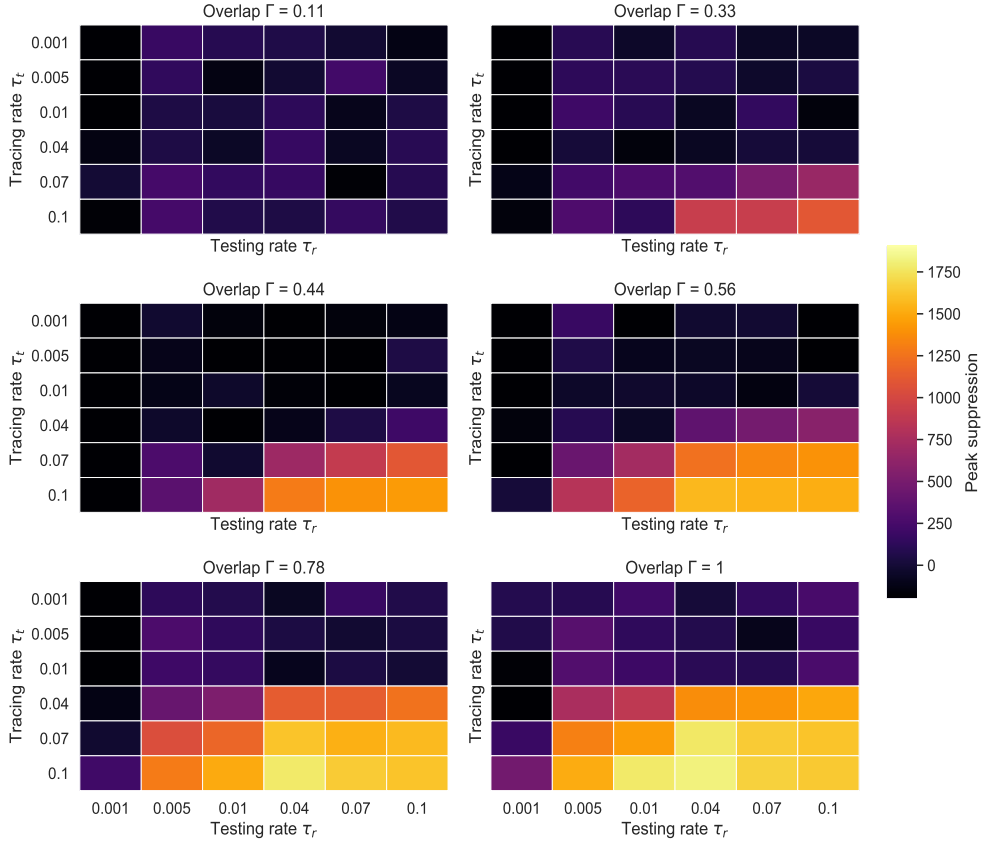


FIGURE 3.7: ER network – Heatmaps of achieved peak suppression for various testing and tracing rates. $N = 10000$, $K = 10$, $p_a = 0.2$. Averaged over 105 runs.

3.4.3 Hospitalizations and deaths when improving tracing overlaps

While keeping $\tau_r = 0.1$ and $\tau_t \in \{0.1, 0.2\}$, we explore how varying the tracing overlap Γ influences the number of hospitalizations and deaths in our model. Using multiple Barabási-Albert networks, with $N = 10000$ and average degree $K = 10$, an asymptomatic probability of $p_a = 0.2$, and a non-compliance rate $\eta = 0.001$ (with automatic isolation exit after 14 days), we follow the evolution of current hospitalizations and total deceased across time for three age-based groups. To avoid runs in which the epidemic gets quickly contained by chance, the simulation starts with 10% of the nodes in the I_p state – $c(t_0) = 10\%$ of N .

As such, Fig 3.8 displays on each row the epidemic outcome of a different groups: one consisting only of adults, another with seniors only, and a mixed-age group representative for the United Kingdom. To model the latter, we calculate p_h , λ_{H-R} , and λ_{H-D} by averaging the values corresponding to each age group (i.e. adults, seniors and children), using recent demographic estimates as weights (Roskams, 2022). As expected, the ratio between different tracing policies remains fairly consistent across all age groups. Additionally, an inverse relationship can be observed between the degree of overlap and both hospitalization and death counts; as the former increases,

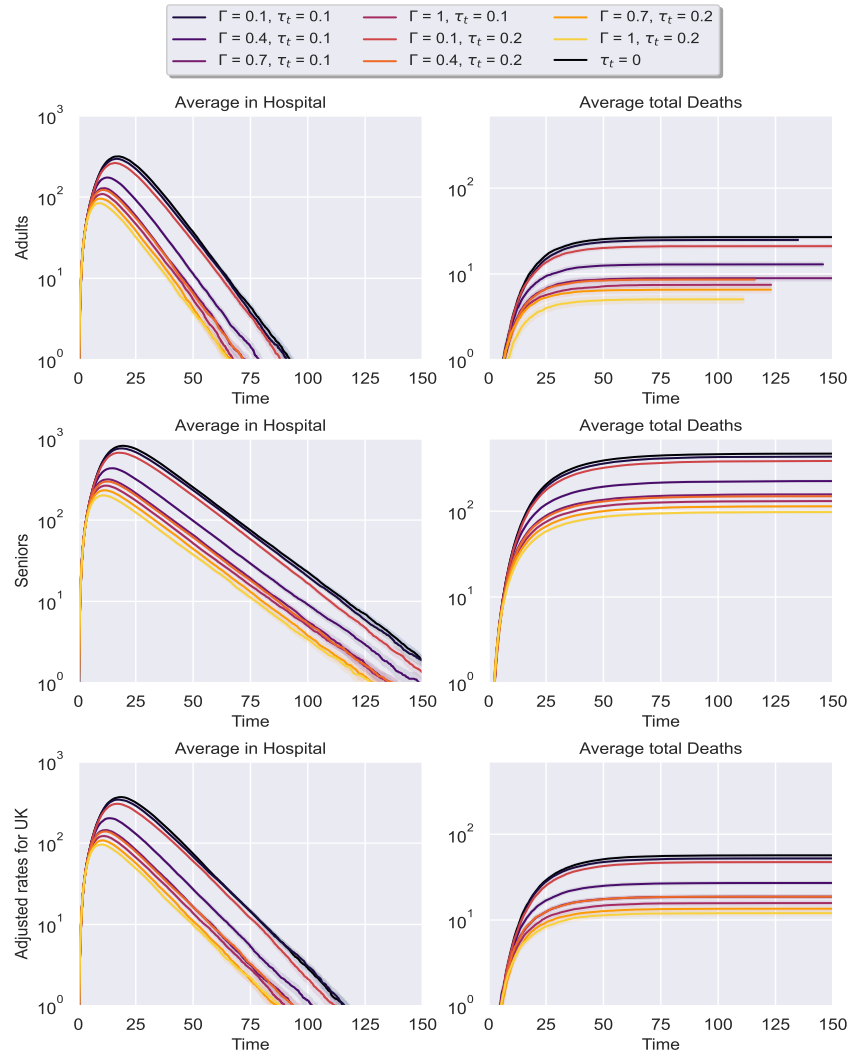


FIGURE 3.8: **SF network – Hospitalizations and total deaths over time.** Setting $\tau_r = 0.1$, $N = 10000$ and $K = 10$. Each row represents a different community: adults only, seniors only, and mixed (adjusted rates based on UK demographics). The 95% confidence intervals resulted from 105 runs are displayed around each line. The case with no contact tracing ($\tau_t = 0$) is colored in black.

the latter two decrease. Compared to no tracing, an efficient programme running at a rate $\tau_t = 0.2$ and benefiting from an infection network coverage $\Gamma \geq 0.7$ is shown to reduce the hospitalization peak by more than half an order of magnitude (on a log-10 plot), resulting in many saved lives. The close proximity (or even overlap for $\Gamma = 0.7$, $\tau_t = 0.1$ and $\Gamma = 0.4$, $\tau_t = 0.2$) of some curves pertaining to contrasting parameter configurations, where the differences between τ_t and Γ are of opposite signs, suggests that public health personnel can often compensate for the lower network accuracies by increasing their tracing rate (and vice-versa).

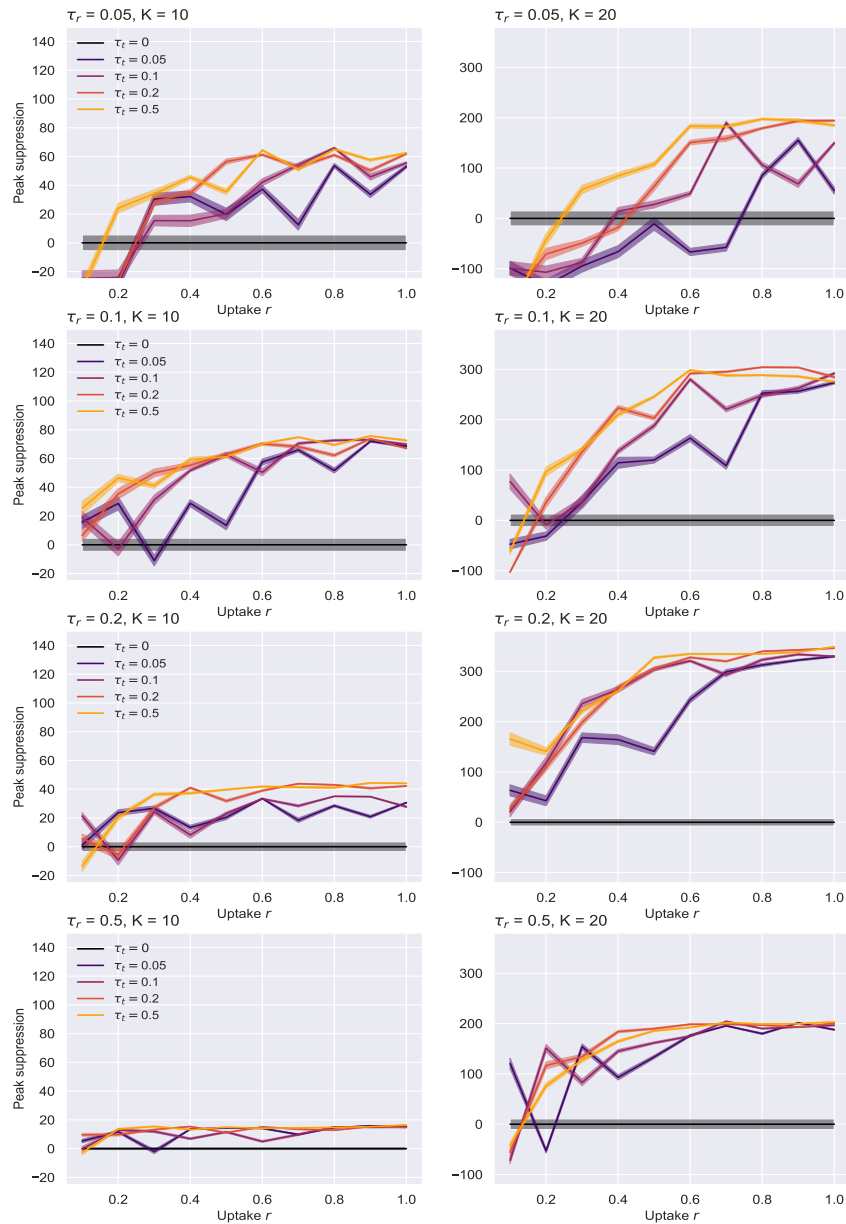


FIGURE 3.9: ER network – Uptake rate r against peak suppression. Suppression is difference in peak to no tracing, i.e. $\tau_t = 0$. $N = 1000$, $p_a = 0.2$, $\eta = .001$. $K = 10$ given on the left, $K = 20$ on the right. The case with no tracing ($\tau_t = 0$) is colored in black. Lines are plotted with the 95% confidence intervals from 750 runs.

3.4.4 Effects of average degree and app uptake

Further, we analyze the impact of the application uptake in scenarios with different average degrees ($K \in \{10, 20\}$), and more appropriate testing and tracing strategies – i.e. $\tau_t, \tau_r \in \{0.05, 0.1, 0.2, 0.5\}$. For this trial, we set $N = 1000$, the asymptomatic probability to $p_a = 0.2$, and the non-compliance rate to $\eta = 0.001$ (with automatic isolation exit after 14 days), selecting a single I_p node as the infection seed. The results are averaged over 750 simulations to reduce the variance induced by the smaller N .

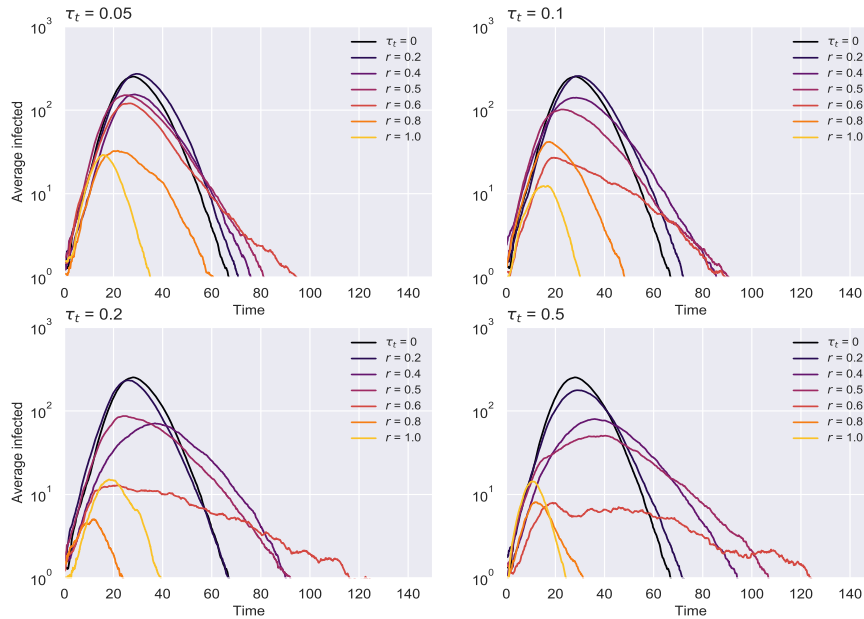


FIGURE 3.10: ER network – Epidemic evolution over time for $\tau_r = 0.1$, $N = 1000$ and $K = 20$. Results averaged over 750 runs. No tracing ($\tau_t = 0$) is given in black.

Fig 3.9 shows the peak suppression achieved by each strategy given a specific adoption level. For $\tau_t = 0.05$, uptakes $r \leq 0.5$ generally render results within the noise region. Improving the contact tracing rate, however, leads to a noticeable decrease of this maximal point, even at smaller adoption levels. This is particularly true in the larger average degree case. Interestingly, deploying a wider-scale testing programme alone ($\tau_r = 0.5$) seems to lead to a considerable spread reduction, which makes contact tracing less beneficial at achievable uptakes (even entirely profitless in the scenario where $K = 10$).

Our findings indicate that a testing rate of $\tau_r = 0.1$ remains suitable in conjunction with different contact-isolation regimes not only for the previous experiment with $N = 10000$, but also for these smaller scale scenarios. Further examining the effect of such a testing regime on the evolution of the spread (see Fig 3.10), it transpires that the epidemic curves significantly ‘flatten’ for uptakes $r \geq 0.4$, the effect becoming more apparent as the contact tracing rate increases.

Interestingly, our epidemic data suggests that higher uptakes do not always guarantee better epidemic outcomes (e.g. $r = 1$ ends up with a higher peak than $r \in \{0.8, 0.9\}$ for $\tau_r \geq 0.2$). This effect arises from the early isolation of a large number of susceptibles (scenario akin to a partial lockdown), which ends up exerting a significant and unpredictable influence on the transmission chain when these individuals eventually resume their social interactions.

3.4.5 Combining digital tracing with an imperfect manual process

In this section, we study a more realistic scenario in which digital solutions complement an inherently imperfect interview-based tracing system. To that end, a *triad network* topology is employed, with digital tracing happening at a rate of $\frac{1}{\tau_t}$ days on average, over one subgraph given by the uptake $r \in \{0.1, 0.25, 0.4, 0.55, 0.7, 0.85, 1\}$, while the manual process gets carried at a slower pace of $2 + \frac{1}{\tau_t}$ days on average, over a third network view whose edges have been randomly removed according to the degree of overlap $\Gamma \in \{0.1, 0.25, 0.4, 0.55, 0.7, 0.85, 1\}$. For the purpose of this experiment, we make use of a more representative graph structure for the SARS-CoV-2 transmission based on the Holme-Kim model (Holme and Kim, 2002), which features both a SF degree distribution and a SW clustering coefficient. The network parameters chosen here are: $N = 1000$, $m_{HK} = 10$ (number of random edges to add for each new node; this replaces K in Eq 3.1 for calculating N_{ute}) and $p_{\Delta} = 0.2$ (probability of making a triangle after adding a random edge). The simulation starts with 10% of the nodes in the I_p state. The other parameters remain unchanged from the previous section, including the number of total runs.

The most notable trend in Figs 3.11 and 3.12 is the monotonicity of each curve with respect to r , while the differences in the epidemic outcome between different τ_t values are considerably more apparent than in the previous experiment. This is a direct consequence of the increased number of infected people selected for time t_0 , which prevents simulations from averaging over too many early-stopped runs. Considering the scale that the COVID-19 pandemic has reached and the unavoidable presence of a delay between the infection onset and the debut of tracing, scenarios such as this one are more likely to occur, and therefore of a greater interest (Hinch et al., 2020b; Hellewell et al., 2020).

Fig 3.11 shows the degree of peak suppression achieved by utilizing digital and manual tracing solutions when compared to a scenario in which no contact tracing is performed. These results suggest that, as the efficacy of the interview-based process increases (i.e. less contacts get missed), lower and achievable application adoption rates (20-50%) are sufficient to effectively reduce the maximal point of the epidemic. When the tracers are eventually able to ‘see’ the full network of contacts ($\Gamma = 1$), varying r no longer impacts the spread significantly, as should be expected. In contrast, a very good testing regime ($\tau_r \geq 0.2$) can partially compensate for an inefficient manual tracing system ($\Gamma = 0.1$) within the aforementioned uptake range.

Our estimate of $R = 3.20$ for minimal interventions (i.e. $\tau_r = 0.05$ and no tracing) during this scenario’s first week falls within the confidence interval of the basic reproduction number $R_0 \in [3.09, 3.24]$ derived in Di Domenico et al. (2020) by applying the next-generation approach (Diekmann et al., 1990) on a model fairly similar to ours. Fig 3.12 demonstrates that with good testing regimes ($\tau_r \geq 0.1$) and a

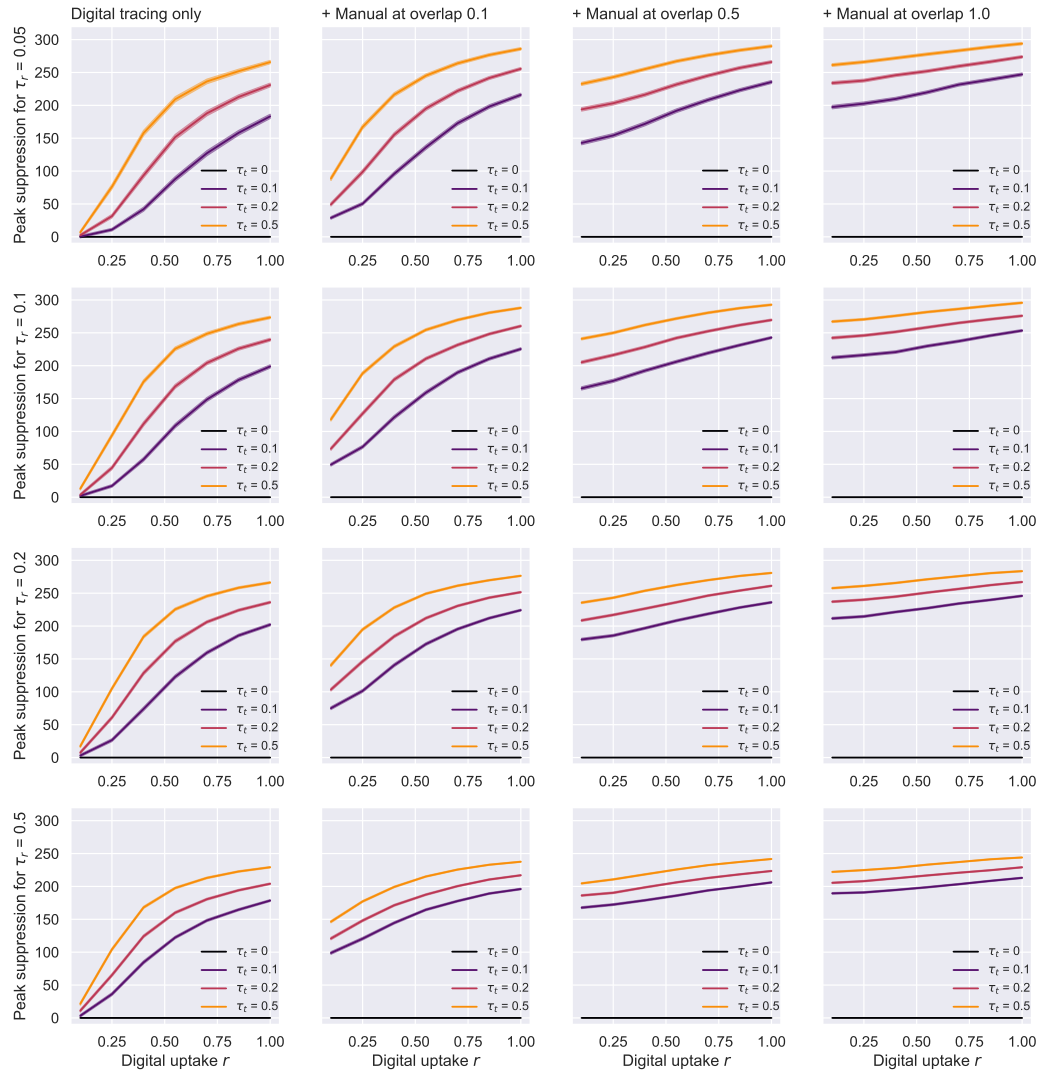


FIGURE 3.11: **HK network – Uptake rate r against peak suppression.** The results here correspond to networks generated with parameters $N = 1000$, $m_{HK} = 10$, $p_{\Delta} = 0.2$, $p_a = 0.2$, $\eta = .001$. On the left, we present a scenario in which only digital tracing is conducted, whereas the next 3 columns represent simulations with a combination of digital tracing on a second network, and manual tracing over a third network with various overlaps: 0.1, 0.55, 1. The 95% confidence intervals are also displayed. The case with no tracing ($\tau_t = 0$) is colored in black.

reasonable manual tracing in place ($\Gamma \geq 0.5$), achievable uptake levels are enough to limit R_e to a value close to 1. In contrast, digital tracing alone fails to significantly reduce the spread unless both the testing and the adoption rates are very high. Similarly to what could be observed in Fig 3.11, uptakes play a minor role in the infection proliferation if tracers are able to track the whole contact network eventually, yet this scenario is rather unlikely in real life. Interestingly, most of the trends outlined in the peak suppression charts are faithfully mirrored by the evolution of R_e in the first week of the simulation. This confirms that efficient contact tracing in the early stages of an outbreak is essential for containing a virus like SARS-CoV-2 (Shah et al., 2020).

Even though peak suppression remains a good metric for assessing the benefits of

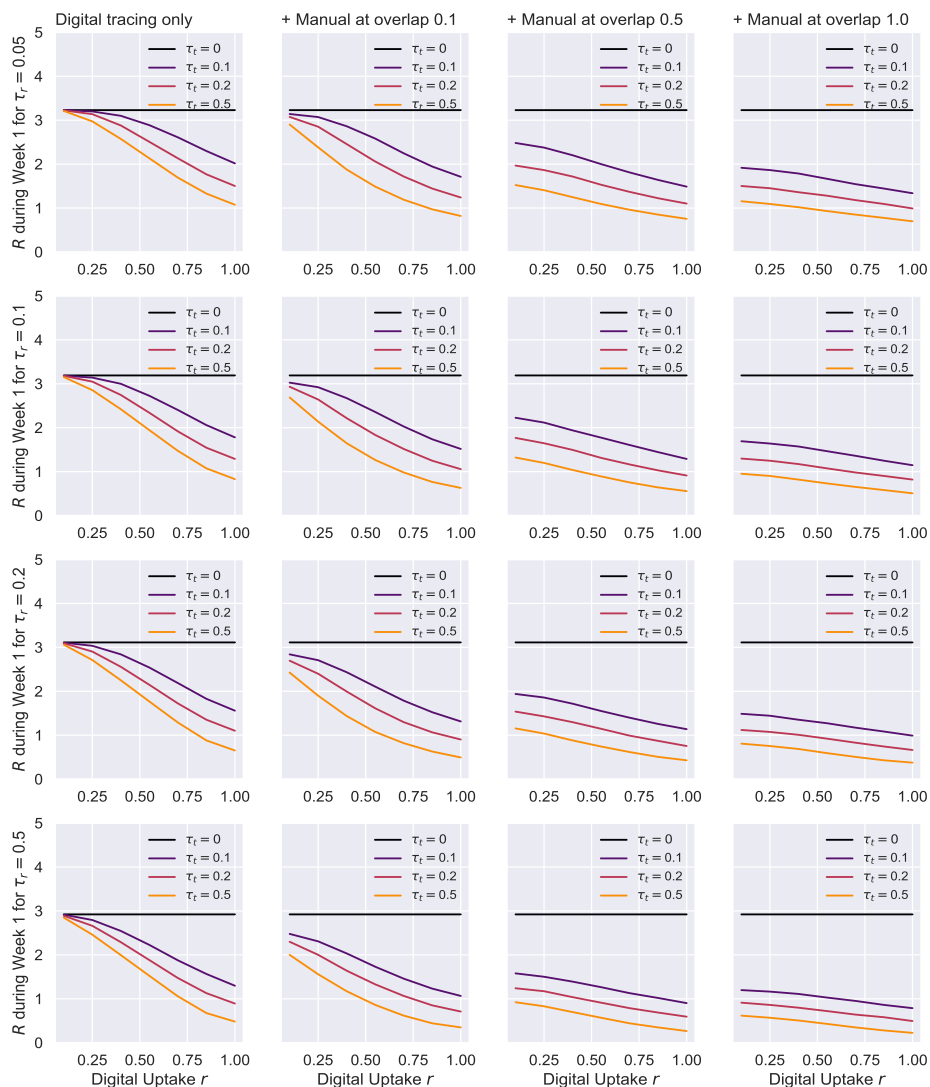


FIGURE 3.12: **HK network – Uptake rate r against the effective reproduction number.** Results here correspond to networks generated with $N = 1000$, $m_{HK} = 10$, $p_{\Delta} = 0.2$, $p_a = 0.2$, $\eta = .001$. On the left, we present scenarios in which only digital tracing is conducted, whereas the next columns represent simulations with a combination of digital tracing on a second network, and manual tracing over a third network with various overlaps: 0.1, 0.55, 1. The case with no tracing ($\tau_t = 0$) is colored in black.

public interventions, policy makers are more often interested in what combinations of these measures can quickly bring R_e to acceptable levels. With this in mind, we plot the contour lines of the R_e values produced by multiple configurations of interview-based network overlaps, testing and digital tracing adoption rates (see Fig 3.13, but also Fig B.5). In line with previous work (Mancastropa et al., 2021), most configurations reveal that scaling up the process of manual tracing has larger benefits than improving the application usage rates. Given the *estimated* uptake rates in 2020 of GPAW (2020), Finland and Iceland had around 40%, Ireland had 35%, the UK had 30%, and Germany and Norway had 27% of their populations participating in digital tracing. In this case, a high contact tracing rate ($\tau_t = 0.5$) combined with an effective

testing regime ($\tau_r \geq 0.2$) can lower R_e below 1 even if tracers fail to identify up to half the contacts ($\Gamma \geq 0.5$). Should adoption improve to 50%, the aforesaid effect would be obtained with a testing rate half as good. In contrast, a moderate tracing rate only becomes effective when a large-scale testing programme gets deployed ($\tau_r \geq 0.5$) or bigger uptakes are achieved ($r > 50\%$). We note the quoted uptakes are based on total download counts, which may not always reflect the true number of active users.

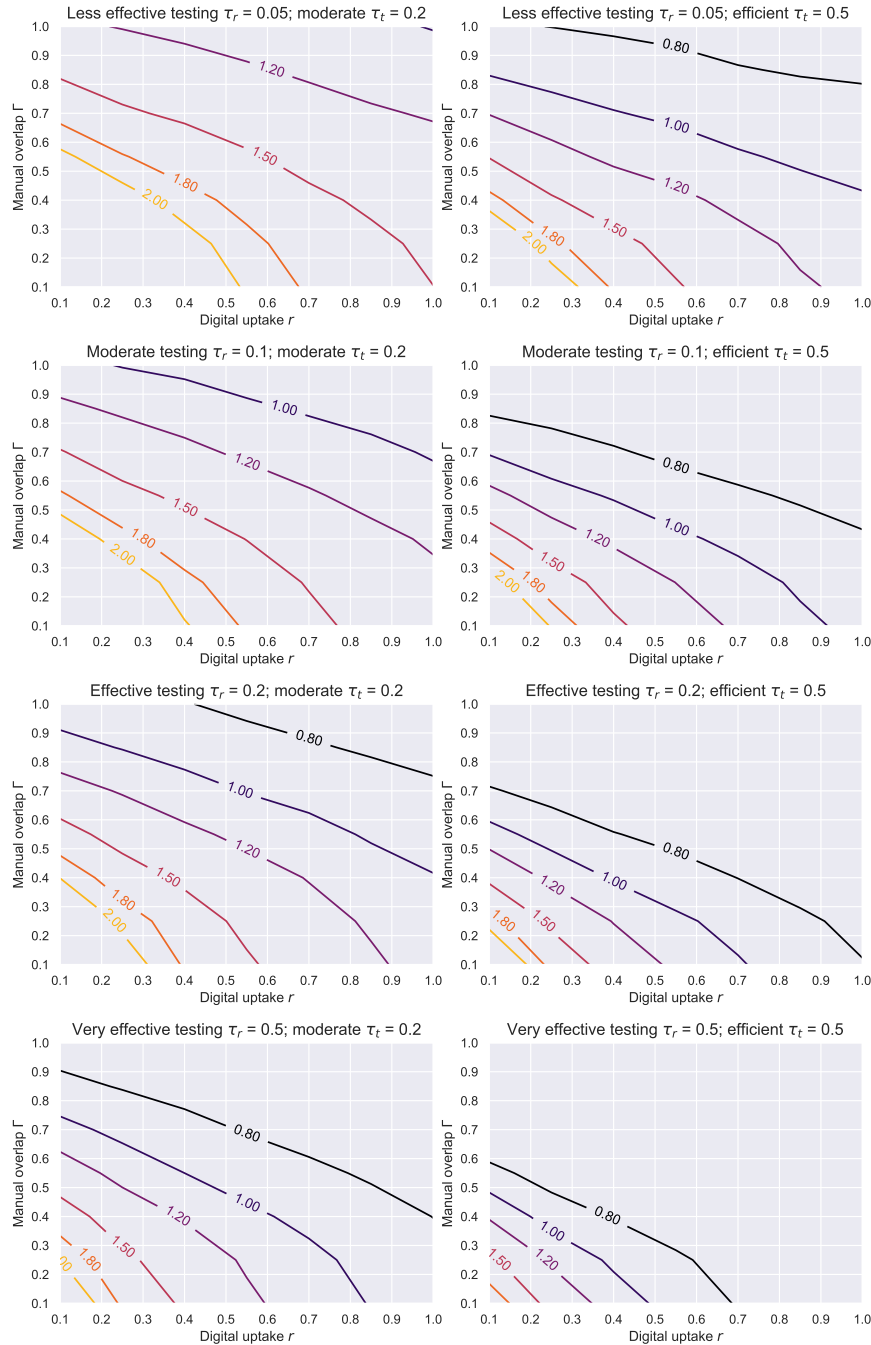


FIGURE 3.13: **HK network – Contour plots of R_e based on the level of manual tracing overlap Γ and digital tracing uptake r .** The results here are for a networks generated with parameters $N = 1000$, $m_{HK} = 10$, $p_{\Delta} = 0.2$, $p_a = 0.2$, $\eta = .001$. Each line represents a different testing level τ_r , while the columns correspond to a moderate (left) and efficient (right) level of tracing effort and isolation compliance given by τ_t .

3.4.6 Contact tracing efficiency in a real social network

Lastly, we evaluate the ability of digital tracing to curb an epidemic simulated over a real social network, in the presence or the absence of manual contact tracing. In this scenario, both the population size N and the average degree K are data-driven, with the latter also changing dynamically ($N = 74$, $K_{t_0} = 5.62$ at time t_0). Given that the network represents a tightly-knit community (static average degree $K_{static} > 60$), we investigate a broader range of testing and tracing rates: $\tau_r \in \{0.1, 0.2, 0.5, 1, 1.5\}$, $\tau_t \in \{0.1, 0.2, 0.5, 1, 1.5, 2\}$. The uptakes r , the overlaps Γ , and the initial incidence $c(t_0)$ are left unchanged from the last passage, while the relative delay between digital and manual tracing is kept at 2 days on average. The probability of becoming an asymptomatic case following exposure is fixed at $p_a = 0.2$ for the purpose of our initial discussion, but a comparison to the case in which $p_a = 0.5$ can be consulted at the end of this section.

The trends in Figs 3.14 and 3.15 follow the same pattern as those in Figs 3.11 and 3.12, respectively. This is because there is a positive and consistent relationship between the testing and tracing rates, and the achieved containment. At the same time, with an adequate overlap $\gamma \geq 0.5$, low and achievable uptakes consistently attain significant peak reductions. These configurations drive the R_e estimate of the first week below 1.5 even when the testing rate is smaller than 0.5. In contrast, with higher uptakes, the degree of overlap becomes less influential for the epidemic outcome. Interestingly, the benefits of increasing the tracing effort τ_t beyond the value of 1 remain minimal across all parameter configurations.

Fig 3.16 shows the 2D contours of the estimated R_e value during the first 7 days of the simulation for the entire parameter range, except for τ_t which is kept ≤ 1 based on the previous observation. When comparing to Fig 3.13, we can see that a significantly faster testing strategy is needed in this case to contain the epidemic and force R_e below 1. This is a consequence of dealing with outbreaks in such a densely-connected community, where the virus spreads too quickly to afford testing at lower rates than 0.5, if the objective is to maintain R_e subunitary. Similarly, an efficient $\tau_t \geq 0.5$ is necessary for achievable uptakes to attain the aforesaid goal. Therefore, when public health resources are scarce, it is advisable to limit the movement or impose lockdown measures within such hubs.

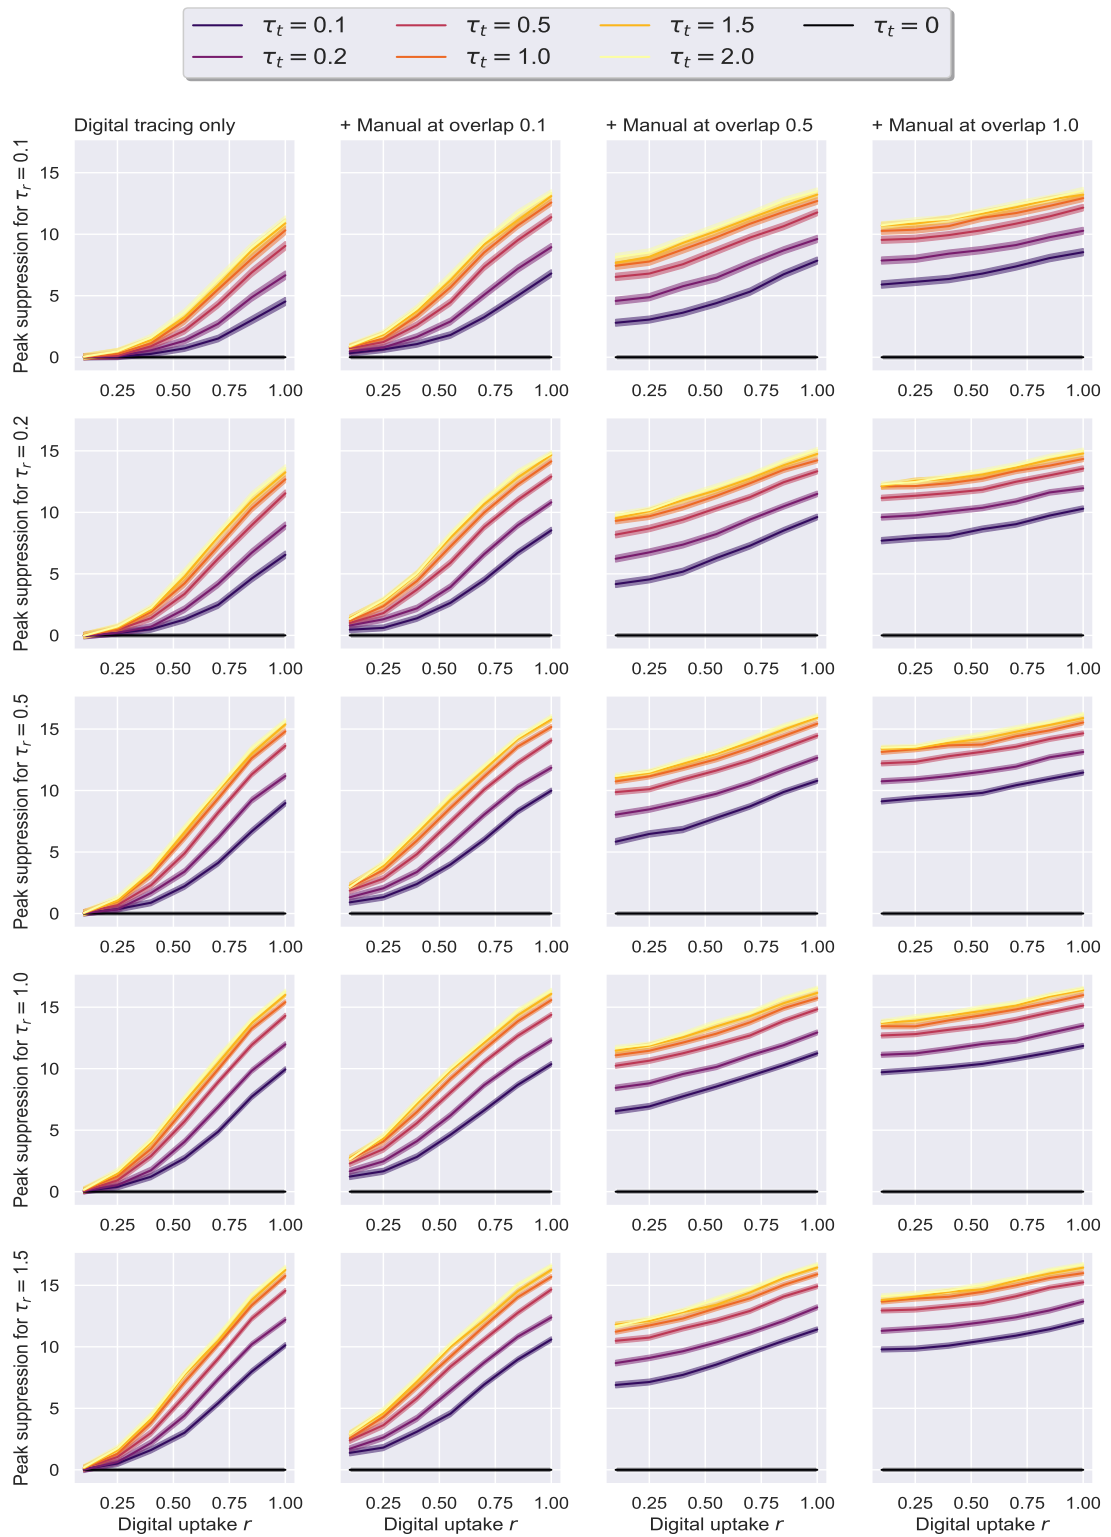


FIGURE 3.14: **Social Evolution – Uptake rate r against peak suppression.** The results here correspond to the real Social Evolution network, dynamic over the studied period of 31 weeks, $p_a = 0.2$, $\eta = .001$. On the left, we have a scenario in which only digital tracing is conducted, whereas the next 3 columns represent simulations with a combination of digital tracing on a second network, and manual tracing over a third network with various overlaps: 0.1, 0.55, 1. The 95% confidence intervals are displayed. No tracing case is colored in black.

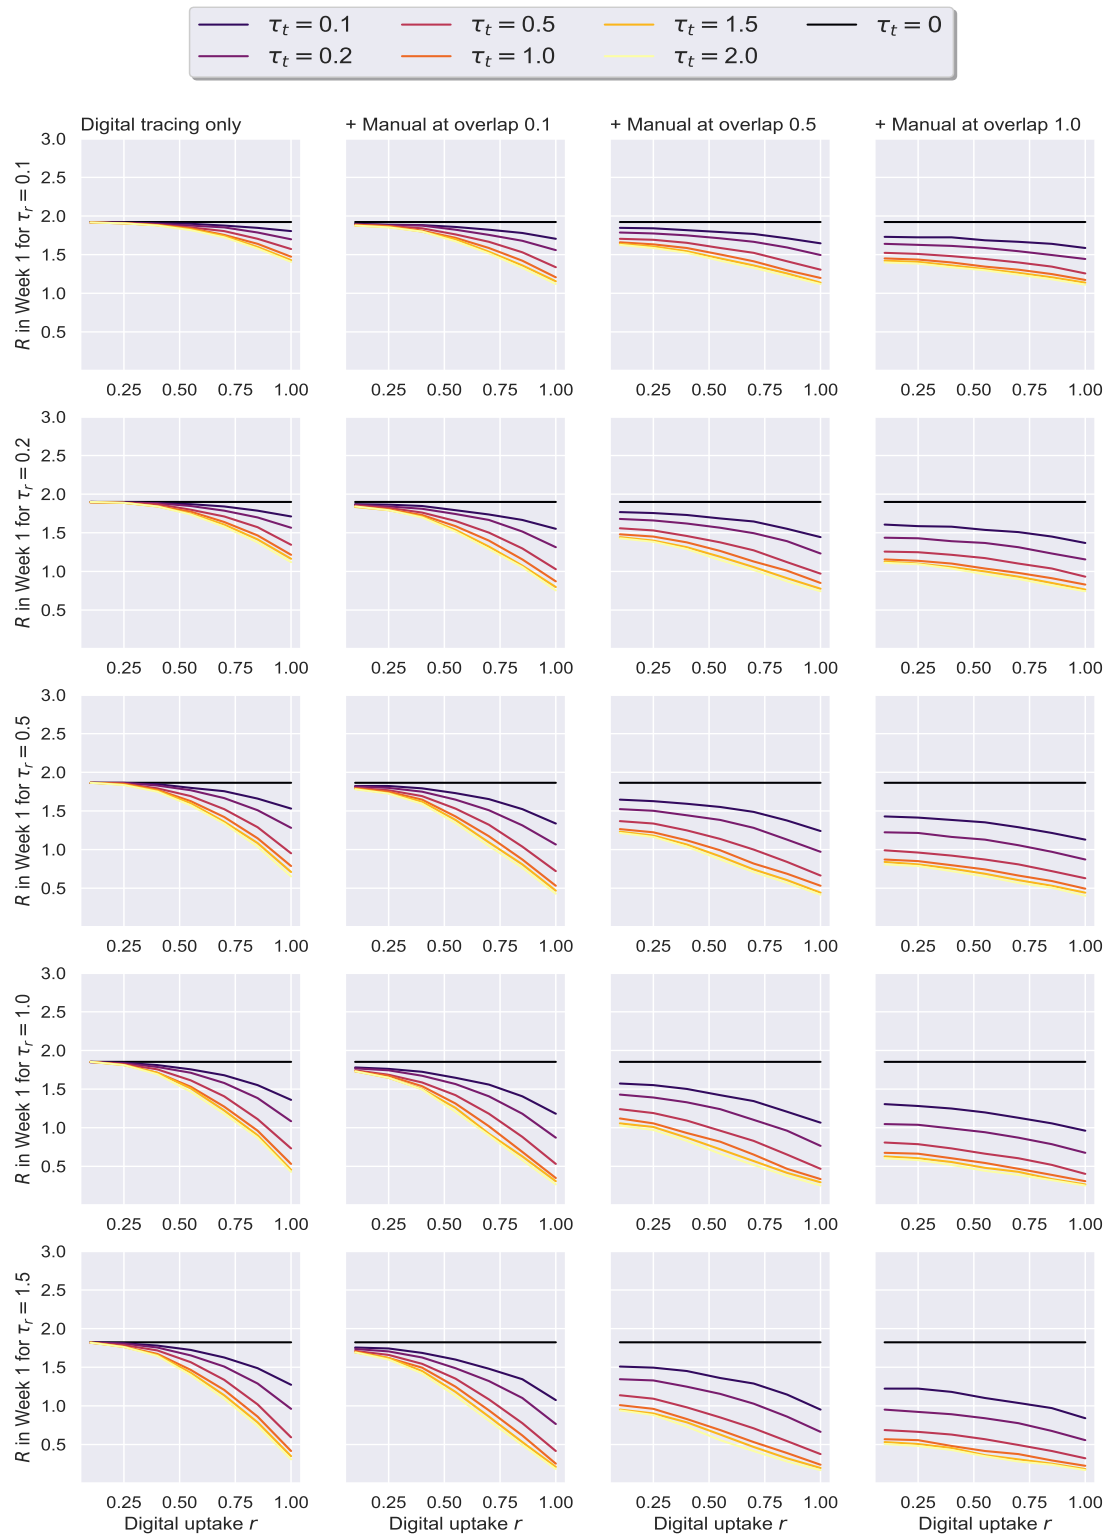


FIGURE 3.15: **Social Evolution – Uptake rate r against the effective reproduction number.** The results here correspond to the data-driven Social Evolution network, $p_a = 0.2$, $\eta = .001$. On the left, we have a scenario in which only digital tracing is conducted, whereas the next columns represent simulations with a combination of digital tracing on a second network, and manual tracing over a third network with various overlaps: 0.1, 0.55, 1. No tracing case is colored in black.

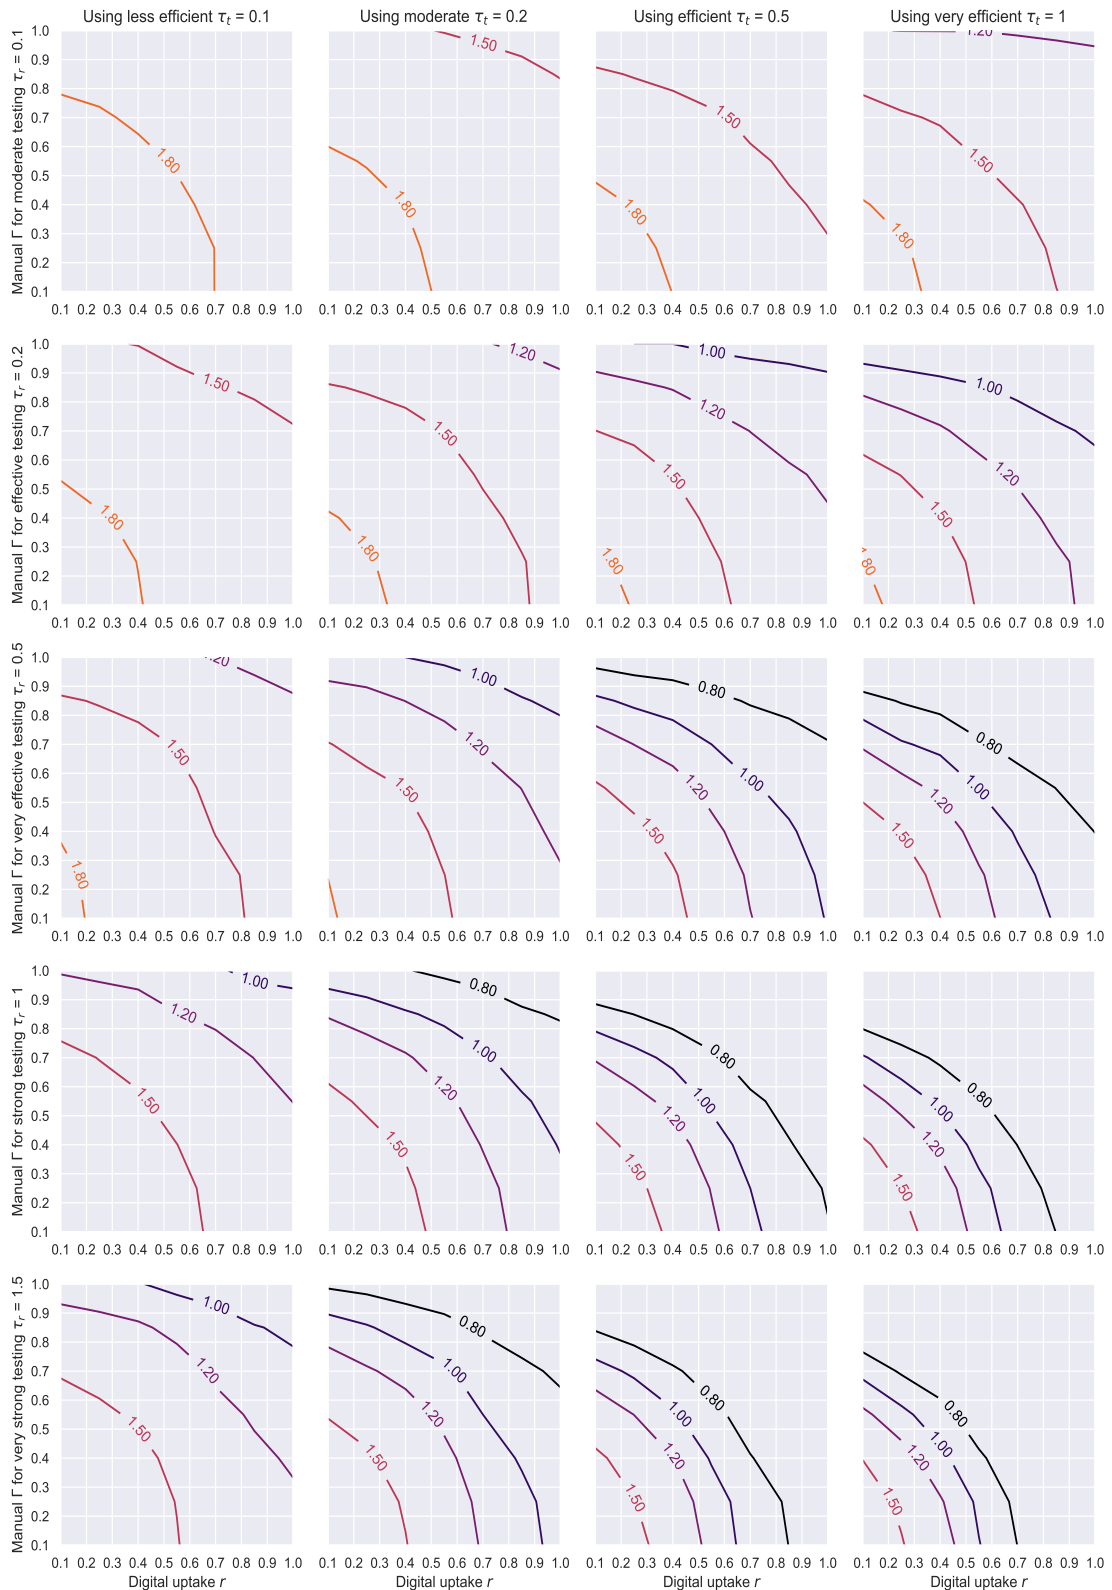


FIGURE 3.16: **Social Evolution – Contour plots of R_e based on the level of manual tracing overlap Γ and digital tracing uptake r .** The results here correspond to the real Social Evolution network, dynamic over the studied period of 31 weeks, $p_a = 0.2$, $\eta = .001$. Each line represents a different testing rate τ_r , while the columns showcase a less efficient (left), a moderate (center-left), an efficient (center-right) and a very efficient (right) level of tracing effort and isolation compliance τ_t .

Finally, we investigate whether the efficacy of tracing appreciably changes when different p_a values are considered. An asymptomatic node is assumed to be less infectious – $r^I = 0.5$, but also less likely to get tested positive – $r^T = 0.8$, so the epidemic dynamics should significantly differ when varying this probability. Remarkably, however, we observe the benefits of contact tracing do not significantly fluctuate across the two studied values in the majority of the scenarios under scrutiny (see Figs 3.17 and 3.18). As shown in Fig 3.18, the most apparent differences in R_e are registered when less accurate tracing networks ($\Gamma, r < 0.5$) and less effective testing rates ($\tau_r < 0.5$) are employed. That is to say a suboptimal ‘test and trace’ policy leads to more people getting infected when $p_a = 0.2$, yet the afferent higher rate of average infectiousness can be offset by the smaller likelihood of nodes testing positive in the $p_a = 0.5$ setting, ultimately leading to minimal differences for more adequate policies.

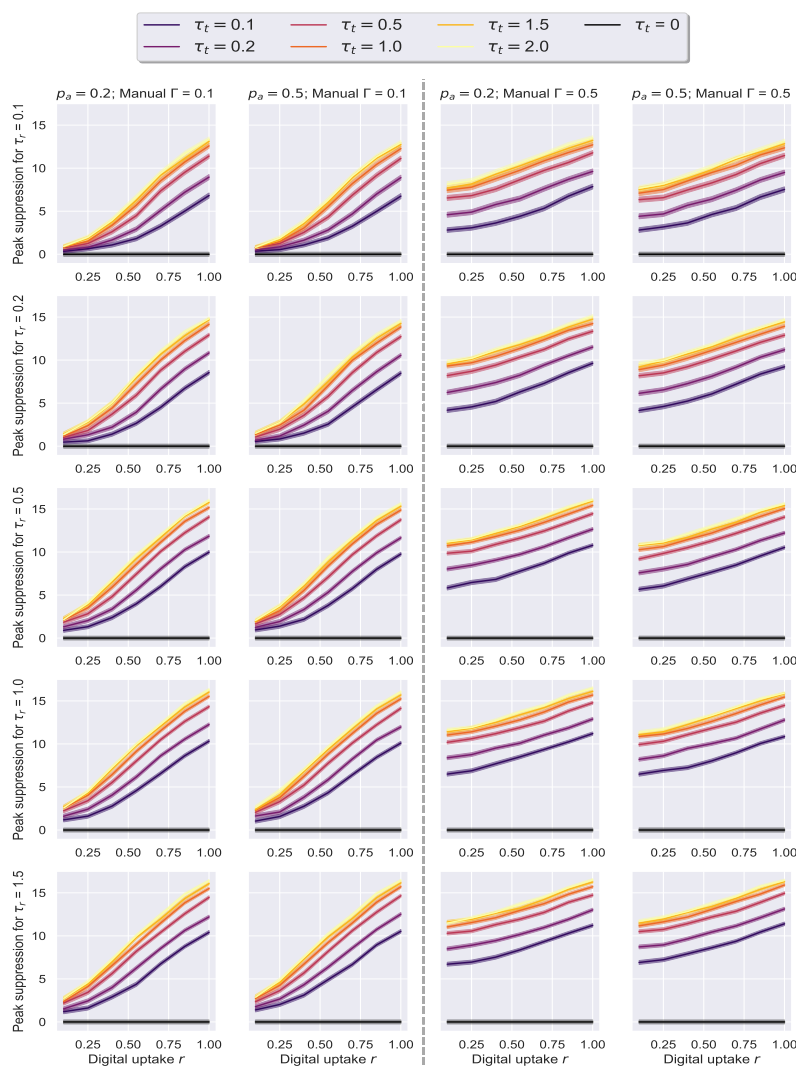


FIGURE 3.17: **Social Evolution – Uptake rate r against peak suppression, for different p_a values.** The results here correspond to the real Social Evolution network, dynamic over the studied period of 31 weeks, $\eta = .001$, and either $p_a = 0.2$ (on the left of each pair) or $p_a = 0.5$ (on the right of each pair). The left-quadrant pairs represent a triad network scenario with manual overlap $\Gamma = 0.1$, while the right quadrant shows $\Gamma = 0.5$. 95% confidence intervals are displayed. No tracing case is colored in black.

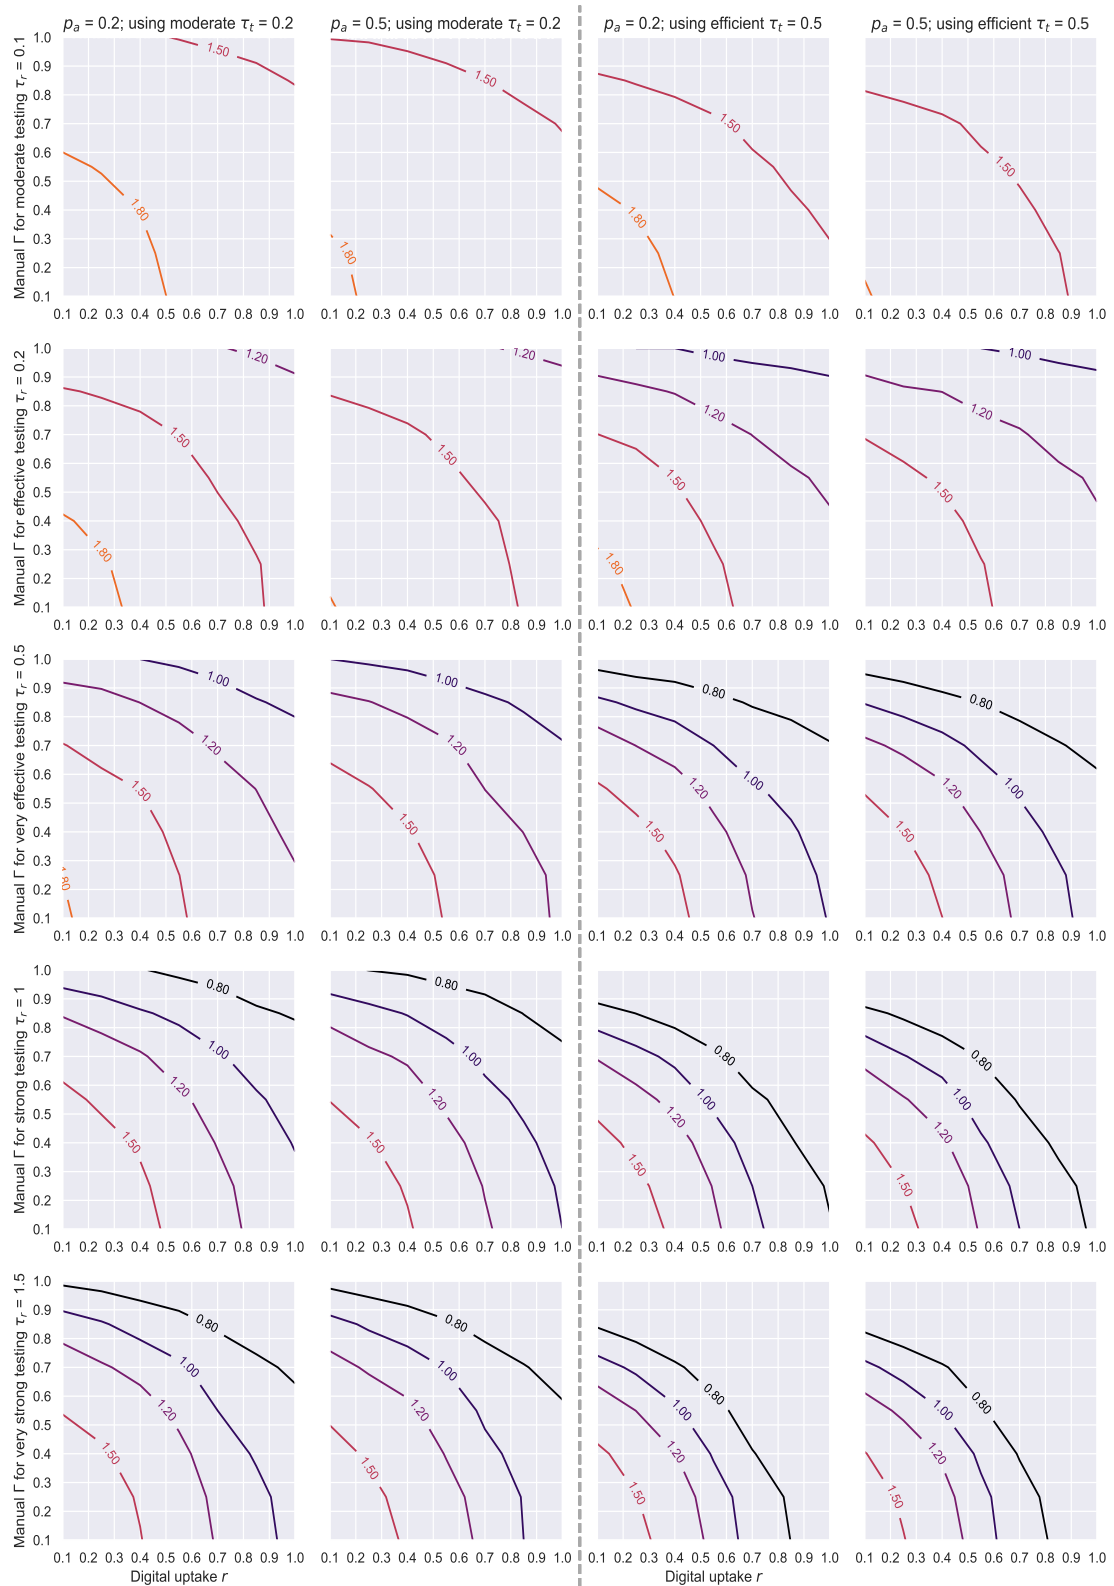


FIGURE 3.18: Social Evolution – Contour plots of R_e based on the level of manual tracing overlap Γ and digital tracing uptake r , for different p_a values. The results here correspond to the data-driven Social Evolution network when $\eta = .001$, with each pair of charts describing $p_a = 0.2$ on the left and $p_a = 0.5$ on the right. Each line represents a different testing level τ_r , while the columns showcase combinations of one p_a value together with a level of tracing effort given by τ_t .

3.5 Case study conclusion

Drawing on the seminal work of [Farrahi et al. \(2014\)](#), we introduce a new methodology that can assess how “test and trace” strategies might impact the transmission dynamics of complex viral epidemics, such as those caused by COVID-19. Following a comprehensive analysis of our model’s parameters, the procedures described here can be utilised to predict how the SARS-CoV-2 virus would spread through those communities where some indication of the interview-based network overlap and/or the digital tracing uptake exists. To facilitate such endeavors, we make our entire codebase open-source.

The approach we propose can address from a modelling perspective four of the open questions formulated by [Anglemyer et al. \(2020\)](#) in their Cochrane Review: the combined effects of digital and manual tracing can be studied via the triad network topology, populations with poor access to the internet may be factored in by the degree of overlap Γ , individuals that have privacy concerns or accessibility issues can be represented in the system via the application adoption rate r , while the repercussions of balancing false positives and false negatives of tracing can be assessed through the statistics our simulations readily capture (for examples, please consult Fig B.2 and Fig B.1). Consequently, the model we put forward is versatile enough to answer many important research questions, beyond what we demonstrate here.

The simulations we conduct show that digital tracing remains a viable solution for reducing the peak of an outbreak, as well as the effective reproduction number R_e , even when its adoption levels are lower. At the same time, a less efficient interview-based process, which misses up to half the contacts, can still contain the spread if coupled with achievable application uptakes (i.e. 30-40%) and appropriate testing regimes. This is because the overlap Γ , our measure of network recovery in manual tracing, has more impact when increased than the uptake of the digital method, which makes the aforementioned range for the adoption rate adequate. This finding corroborates the argument of [Mancastroppa et al. \(2021\)](#) that manual tracing is generally more effective than digital methods, even though the former presumes substantial delays. For highly-connected communities, the scale of the testing strategy becomes appreciably more important for the effectiveness of ‘test and trace’ programmes. In this case, the peak reduction is strongly correlated with the speed and the impact of the tracing system, and the corresponding public-health communiques. The faster and more efficient the latter are, the more likely the involved communities are to comply with the self-isolation recommendations issued by authorities, as more of their contacts get traced and isolated (facets essentially captured by the τ_t rate).

We would like to emphasize that the parameter ranges under scrutiny in this chapter are by no means exhaustive. Therefore, we leave for future exploration studying the effects of extensively varying the average degrees of the random networks, the

non-compliance rates, the initial infections, or the variant-specific asymptomatic and hospitalization probabilities. Looking at such diverse scenarios would allow one to better estimate the shortfalls of contact tracing when different variants of concern are circulating, discover factors that may have introduced significant inefficiencies into the strategies adopted by many countries (e.g. high non-compliance levels (Lewis, 2020)), while also ensuring the variability induced by early-stopped simulations is curbed.

We appreciate that our modelling approach, which utilizes random testing and frequency-based contact tracing, may not be the most optimal or have the highest level of realism for the implementation of “test and trace” programmes. A more efficient strategy would be to direct mass-testing campaigns to the graph hubs or dynamically adapt tracing efforts in the highly-impacted regions of the diffusive space, thus improving the epidemic outcome. To finely control the network dynamics in such an informed manner, we can employ a combination of graph neural networks and gradient-based reinforcement learning techniques. Please refer to the next chapter for more details on this alternative methodology.

Future work could leverage mobility datasets recorded during the COVID-19 pandemic to infer a broader range of network structures, or derive time-dependent estimates of the transmission rate, as previously described in Liu et al. (2020a). Additionally, all the viral-specific parameters could be tailored to the epidemiological situation of a country of interest by fitting them to governmental data reporting on the number of COVID-19 hospitalizations or deaths recorded there. An example of fitting a simpler SEIR compartmental formulation to the death counts registered in various US states is given in Appendix D.2.

Chapter 4

EpiCURB: Epidemic Control Using RL for Budget allocation

4.1 Introduction

The recent outbreak of the SARS-CoV-2 virus has significantly transformed the way we strategize and respond to the spread of infectious diseases. Extreme control measures like full-scale lockdowns were shown to have adverse effects on the worldwide economy and human well-being (Joffe, 2021; Knox et al., 2022), with opposition against such disruptive interventions sharply growing across 2021 among the adults living in the United Kingdom.¹ In response, less rigid health policies, such as encouraging the public to practice social distancing, deploying comprehensive testing schemes, and conducting large-scale contact tracing (Dighe et al., 2020), were enacted to reduce the burden of infection rates. Despite the advent of highly effective vaccines (Andrews et al., 2022; Bruxvoort et al., 2021; Lopez Bernal et al., 2021), governments around the world continued to provide extensive financial support for testing and tracing systems for several more months, fueled by evidence of their previous success (Fetzer and Graeber, 2021; Matrajt and Leung, 2020). Prompted by the emergence of milder SARS-CoV-2 variants (Sigal, 2022), concerns about the limitations (Mercer and Salit, 2021) or the societal impact (Clair et al., 2021; Martinez-Garcia et al., 2022) of these programmes, and growing evidence of public non-compliance (Davis et al., 2021), many administrations were later forced to adjust their resource allocation strategy for these programmes. In early 2022, the United Kingdom implemented the “living with COVID” strategy, which aimed to reduce costs and avoid major disruptions, such as the ‘pingdemic’ crisis that resulted from

¹Surveys by YouGov: <https://tinyurl.com/yougov-early>, <https://tinyurl.com/yougov-attitudes>.

the widespread self-isolation of people through the NHS COVID-19 application (Pandit et al., 2022; Rimmer, 2021). Unfortunately, the indiscriminate reduction of public health efforts did not yield the desired outcome in the initial phases, as evidenced by the significant resurgence in cases that ensued,² a trend that rapidly spread throughout Europe (Henley, 2022). At the same time, with the gradual waning of vaccine protection over time (Ferdinands, 2022; Leung and Wu, 2022), and an apparent decrease in demand for further vaccine doses among healthy adults (Smith, 2022; Su et al., 2022), similar surges remain a possibility in the future.

In this chapter, we introduce *EpiCURB*, a novel approach for enacting “test, trace and vaccinate” programmes that allows them to adjust to the budget and risk criteria of a region, while minimizing the transmission chains generated by the underlying viral diffusion. While we analyze several approaches for targeting the aforementioned processes in a COVID-19 setting under strict budget constraints, we demonstrate that a reinforcement learning agent can discover efficient and generalizable policies that outperform the other baselines in terms of the achieved containment rates. We evaluate our findings on various scenarios of epidemic, budget and contact network configurations, demonstrating the adaptability of our proposed method. Furthermore, we show that even the static non-learning agents substantially surpass conventional untargeted strategies, especially when enhanced with negative test recollection.

The chapter introduces three main novel contributions to this field, as follows:

1. We devise new strategies for targeting public health measures in pragmatic contexts where minimizing the economic and societal costs remains a priority, concomitantly restricting the testing, tracing and vaccinating efforts to higher-risk individuals. To that end, we present a range of novel and existing policies that are implemented or learned by specialized agents using centrality, neighborhood or epidemic state factors, and contrast them with conventional methods, such as sampling based on randomness, acquaintance or contact frequency. Our RL agent, backed by a graph neural architecture inspired from the development of Meiom et al. (2021), is shown to outperform competitive heuristics in all three tasks by up to 15% in the containment rate, while far surpassing the standard random samplers by margins of 50% or more, despite being trained using a simple test prioritization setup with partially observable information. By clustering the resulting node embedding space, a novel community-based strategy for vaccination is also shown to be highly effective.
2. We assess our approach on a wide range of epidemic scenarios simulated by two COVID-specific transmission dynamics: the individual-based multi-site

²Infection survey by ONS: <https://tinyurl.com/ons-covid19>.

mean-field model introduced in Chapter 3 and a standard agent-based model, featuring diffusion parameters adapted from Di Domenico et al. (2020) to an SEIR compartmental formulation. The performance of the introduced agents is evidenced across these diverse settings, confirming not only their transferability between the aforesaid modelling families, but also the merit of the RL policy.

3. In addition to analyzing the numerical results and epidemic curves resulted from simulating our control policies over multiple interaction networks, we also appraise the decision taken by the corresponding learning-based agents, studying the mechanism by which they are arrived at and the reasons underpinning their success. For such a system to be deployed in the real world, policy makers need to be reasonably confident the model produces sensible outputs. At the same time, testing, isolation and vaccination decisions have to be explainable and verifiable when audited or contested. Here, we scrutinize two central factors that determine the RL agent’s performance: quality of candidate selection and timing of detection. This analysis is extended in Appendix C, where we employ GraphLIME (Huang et al., 2022) to offer explanations for the actions taken by our agent’s GNN module, putting into perspective its competitive adaptability.

4.2 Related work and Background

4.2.1 Epidemic modelling

We simulate viral epidemics using a modified version of the framework introduced in Chapter 3, which relies on the original SEIR compartmental formulation but retains the capacity to leverage an individual’s locality information through contact graphs and mean effects (Farrahi et al., 2014; Huerta and Tsimring, 2002). For completeness, we also investigate our policies in a standard agent-based setup, similar in spirit to the network-based approaches proposed in recent works (Abueg et al., 2020; Meirom et al., 2021). In both cases, we employ the dynamics parameters published by Di Domenico et al. (2020) for the SARS-CoV-2 virus. For more details regarding these epidemic modelling methods, please refer to Section 1.3.

4.2.2 Graph neural networks

Our framework leverages ideas from GATv2 (Brody et al., 2022) and GIN (Xu et al., 2019) to attain expressive power and computational efficiency, as per Bronstein (2022). Additionally, we attempt to limit the impact of feature *oversmoothing* (Hoang and Maehara, 2019; Oono and Suzuki, 2019) and *bottlenecks* (Alon and Yahav, 2022) by

employing randomised node features, keeping the number of GNN layers reduced, and utilizing a final FA layer. More details about these approaches and other GNN architectures can be consulted in Section 2.7.

4.2.3 Reinforcement learning

Sequential decision processes are often modelled via Markov Decision Processes (MDPs) of the form $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R})$ (Puterman, 1994), where \mathcal{S} is a state space, \mathcal{A} is an action space, \mathcal{P} is a transition probability matrix, while \mathcal{R} is a reward function for the state-action pairs. Agents sample actions from their policy $a_t \sim \pi(a|s_t; \theta)$, with $a \in \mathcal{A}$, and θ the acting parametrization, then execute them, transitioning to different states s_{t+1} and earning rewards R_t , according to the environment-specific mappings \mathcal{P} and \mathcal{R} . Partially Observable MDPs (POMDPs) of the form $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \Omega, \mathcal{O})$ are a generalization of MDPs (Kaelbling et al., 1998; Åström, 1965), where the agent can only observe $o_t \sim \mathcal{O}(o|s_t, a_t)$ instead of the true state s_t , where $o \in \Omega$. In some scenarios, such as the one we examine, the observation o_t can eventually approach a reasonable estimate of s_t , enabling the agent to act appropriately despite the partial observability of the environment. In these cases, conventional MDP solvers can be used effectively by substituting s_t with o_t in their corresponding equations. For notational convenience, we will keep using s_t instead of o_t throughout this chapter.

One of the main goals of RL is to solve MDPs (and by extension some POMDPs) by predicting and/or maximizing the γ_r -discounted returns of future rewards until the end of the episode T , $G_t^{\gamma_r} = \sum_{i=1}^T \gamma_r^{i-1} R_{t+i}$ (Sutton and Barto, 2018). The prediction task is routinely achieved through supervision using a w -parameterized model of the *value function* $V(s_t|w)$ that predicts $G_t^{\gamma_r}$. The maximization task is equivalent to finding the optimal policy π^* that maximizes $G_t^{\gamma_r}$. One way to achieve this is to estimate the *action-value function* $Q_\pi(s_t, a_t|\theta)$, which is the expected return if the agent chooses action a_t in state s_t , and follows the policy $\pi(a|s_t; \theta)$ thereafter. The optimal policy π^* is typically a greedy policy, related to the action-value function through the relation:

$$\pi^*(a|s_t; \theta^*) = \begin{cases} 1 & \text{if } a = \arg \max_{a'} Q_{\pi^*}(s_t, a'|\theta^*) \\ 0 & \text{otherwise.} \end{cases}$$

To approximate the value functions V or Q , one commonly implements update steps on w or θ , correspondingly, relying on the true returns $G_t^{\gamma_r}$ (see Eq 4.1 for V), or some intermediary bootstrapped values as regression targets (see Eq 4.2 for V). The former of these methods is known as the Monte Carlo (MC) algorithm, and it has been proven to be effective despite its considerable limitations, such as the slow offline learning and high variance (Sutton, 1988; Sutton and Barto, 2018). These drawbacks stem from the fact that MC requires the full episode to be run before any parameter update. A

variation of MC featuring search trees has been used to derive competitive policies in 2-player board games (Silver et al., 2016, 2017). In contrast, the online temporal difference (TD) learning method casts the sum of R_t and the current estimate of the next-step return $G_{t+1}^{\gamma_r}$ as a regression target, lowering the variance and speeding up training at the cost of inducing more bias (van Hasselt et al., 2021). The difference between this target and the current estimate is often called the TD error, denoted as $\delta_t^{\gamma_r}$. Using a delayed parameter set w_D , where D marks the delay ($w_0 = w$), to compute the target values is a common technique used to reduce their variance by making them more stable and less correlated with the current estimates (Mnih et al., 2013). TD learning constitutes the basis for many RL algorithms to date, such as the on-policy SARSA (Rummery and Niranjan, 1994; Sutton and Barto, 2018) and the off-policy Q-learning (Watkins, 1989), which proved successful in multiple settings: reaching or outperforming human-level performance at board and video games (Mnih et al., 2013, 2015), robotic path planning (Harwin and P., 2019), autonomous car driving (Kiran et al., 2021), and many others. The bias-variance trade-off in these approaches can be controlled in the alternative TD(λ) formulation by varying the λ parameter, with the λ -return $G_t^{(\gamma_r, \lambda)}$ used as target (Sutton and Barto, 2018). The latter can also be used to obtain a Generalized Advantage Estimate (GAE), $\hat{A}_t^{(\gamma_r, \lambda)}$ in Eq 4.3, with the special cases $\lambda = 0$ in Eq 4.2, when the advantage is equal to the TD error, and $\lambda = 1$ in Eq 4.1, when the minuend of the RHS term is $G_t^{\gamma_r}$ (Schulman et al., 2018). The advantage function, in its various forms, quantifies the relative benefit of an action over the expected action for a given state, being a key component of policy gradient methods.

$$\hat{A}_t^{(\gamma_r, 1)}(a_t; w) = G_t^{\gamma_r} - V(s_t; w) \quad (4.1)$$

$$\hat{A}_t^{(\gamma_r, 0)}(a_t; w) = \delta_t^{\gamma_r}(w) = R_t + \gamma_r V(s_{t+1}; w_D) - V(s_t; w) \quad (4.2)$$

$$\hat{A}_t^{(\gamma_r, \lambda)}(a_t; w) = G_t^{(\gamma_r, \lambda)} - V(s_t; w) = \sum_{l=0}^T (\gamma_r \lambda)^l \delta_{t+l}^{\gamma_r}(w) \quad (4.3)$$

Approaches that directly optimize both θ and w are called actor-critics (Konda and Tsitsiklis, 1999), and have become the preferred algorithmic choice when faster convergence rates are sought after and sample efficiency is not required. Recent years have seen actor-critic methods, such as Proximal Policy Optimization (PPO) (Schulman et al., 2017) and Deep Deterministic Policy Gradient (DDPG) (Fujimoto et al., 2018; Lillicrap et al., 2019), achieve state-of-the-art results across a wide range of challenging tasks (Lazaridis et al., 2020; Schulman et al., 2017). A commonly adopted strategy for these models is to use the same parameter set in both components (Schulman et al., 2017), i.e. $w = \theta$, which can enrich the learned representations while also increasing the training efficiency. Although online implementations are possible, as is the case with the model we introduce in Section 4.3.5.2, actor-critics are traditionally trained using MC.

Assuming $w = \theta$, we reproduce the standard PPO equations in Eqs 4.5 to 4.8, where \mathcal{E} is the ratio clipping parameter that limits the deviation of a new policy from the old policy, $\text{clip}(\cdot)$ is a function that clips its argument to the specified range, $H_t(\theta)$ is an entropy regularizer that encourages exploration (Haarnoja et al., 2018), while c_1 and c_2 are coefficients that control the importance of the value function error and the entropy regularizer, respectively. The first term in Eq 4.7, given by Eq 4.5, is the clipped surrogate objective, which aims to maximize the expected return while bounding the policy ratio in Eq 4.4. The second term, which in Eq 4.6 is represented as the squared advantage function when $\lambda = 1$, corresponds to a squared-error loss for the value function that evaluates how closely it matches the discounted return. Finally, the third term represents the entropy bonus, which adds a small amount of randomness to the policy to prevent premature convergence to suboptimal policies.

$$r_t(\theta') = \frac{\pi(a_t|s_t;\theta')}{\pi(a_t|s_t;\theta)} \quad (4.4)$$

$$\mathcal{L}_t^{\text{CLIP}}(\theta') = \min[r_t(\theta')\hat{A}_t^{(\gamma_r,\lambda)}(a_t;\theta'), \text{clip}(r_t(\theta'), 1 - \mathcal{E}, 1 + \mathcal{E})\hat{A}_t^{(\gamma_r,\lambda)}(a_t;\theta')] \quad (4.5)$$

$$\mathcal{L}_t^{\text{VF}}(\theta') = [\hat{A}_t^{(\gamma_r,1)}(a_t;\theta')]^2 \quad \mathcal{H}_t(\theta') = - \sum_{a \in A} \pi(a|s_t;\theta') \log \pi(a|s_t;\theta') \quad (4.6)$$

$$\mathcal{L}_t^{\text{PPO}}(\theta') = \mathbb{E}_t[-\mathcal{L}_t^{\text{CLIP}}(\theta') + c_1\mathcal{L}_t^{\text{VF}}(\theta') - c_2\mathcal{H}_t(\theta')] \quad (4.7)$$

$$\theta = \arg \min_{\theta'} \mathcal{L}_t^{\text{PPO}}(\theta') \quad (4.8)$$

The training process of a standard PPO agent proceeds as follows: The actor acts with a policy parameterized by θ for the full duration of an episode, recording in an episodic replay buffer B_∞ the visited states, their evaluations by the critic, the actions taken, and the log-probabilities of those actions stemming from the policy. At the end of the episode, the recorded information is used to calculate $G_t^{\gamma_r}$ and $\hat{A}_t^{(\gamma_r,\lambda)}$, $\forall t$, then everything is batched and used within a multi-epoch optimization procedure that minimizes \mathcal{L}^{PPO} . The optimized parameter set is assigned to θ after all these epochs have elapsed, and the resulting policy is ready to be evaluated on a new episode. This process is repeated for the set amount of episodes.

4.2.4 Combinatorial action spaces

Learning policies in environments with combinatorial action spaces, such as the one we investigate here, has traditionally been considered a difficult undertaking. In spite of this, RL methods proved to be effective in instances like multiple item (Song et al., 2019) or thread popularity selection (He et al., 2019). In the context of epidemics, an RL system based on multi-armed bandits and demographics data was recently introduced by the Greek authorities to prioritize the COVID-19 testing allocations at border control (Bastani et al., 2021). For classic combinatorial problems, such as the

travel salesman (TSP) and its vehicle routing variants, RL approaches have also been shown to perform well (Bello et al., 2017; Delarue et al., 2020; Kool et al., 2018). Incorporating graph embeddings into the RL agents have generally lead to improved solvers, outcompeting other learning methods (Dai et al., 2017; Joshi et al., 2019). More recently, an RL model relying on GATs was used to solve a wide range of complex selection tasks, from varying skills to escape mazes to choosing appropriate tools for a given goal (Jain et al., 2022).

4.2.5 Influencing graph dynamics

The problem of influencing diffusion processes over networks has been extensively studied in many different settings, most notably for optimizing immunization strategies (Preciado et al., 2014), solving influence maximization (Murata and Koga, 2018), or targeting pathogen testing (Meirom et al., 2015). It has long been established that random vaccination policies tend to be suboptimal, and even simple heuristics like acquaintance sampling can outperform them (Cohen et al., 2003; Miller and Hyman, 2007). Centrality-based strategies were also explored in this context, with PageRank (Chung et al., 2009), eigenvector (Masuda, 2009) or betweenness centrality (Salathé and Jones, 2010) becoming popular choices. Recently, a modelling study showed that vaccinating younger groups, akin to heuristic-based approaches, after the over 80s could have been a superior alternative to the widely-adopted strict age-based prioritization against the highly contagious Delta variant of SARS-CoV-2 (B.1.617.2) (Keeling et al., 2023). For influence maximization, degree-based strategies were shown to render competitive results (e.g. LIR by Liu et al. (2017), degree discount by Chen et al. (2009)). Over time, however, multiple authors have identified problem instances where any centrality measure used by itself can lead to suboptimal results (Braha and Bar-Yam, 2006; Preciado et al., 2014). The question of which heuristic to use for what problem has since become a focal point in many application domains. As an alternative, RL techniques have been proposed for mixing different heuristics in an optimal manner, thus reducing the impact of the aforementioned drawbacks (Meirom et al., 2021; Tian et al., 2020). Node targeting for inferring the state of a spreading process remains less explored in the literature, but efficient heuristics that exploit the epidemic state knowledge regarding a vertex's neighborhood have proven to be more successful than centralities (Meirom et al., 2018, 2015). The domain of prioritizing contact tracing, however, remains largely uninvestigated to date, but recent work suggests that isolating subsets of individuals based on the frequency of appearing in the vicinity of positive cases can lead to similar levels of containment as naively isolating every contact (Kojaku et al., 2021).

Meirom et al. (2021) introduce an RL model that can derive general control policies for diffusion processes over networks, using test prioritization and influence

maximization as illustrations. A GNN-based controller, cast in an actor-critic framework detailed in Section 2.9.2, learns effective policies using simulated data, integrating local and long-distance information over time. The elegance of the approach stems from the fact that the training process is not conditioned on having the full epidemic state made available to the agent. The work also shows that it is possible to learn a policy on small networks (e.g. 1000) and deploy it on larger graphs featuring similar statistics (e.g. 50000, the size of a small city). Our study builds on this versatile control methodology, but differs from it in several key aspects: First, we extend the problem formulation to cover prioritizing testing, tracing, and vaccination, modifying the framework to rank nodes only from eligible subsets. This allows us to add a simple extension to our agents that significantly improves performance: removing recently-tested negative individuals from the action space. Second, we deepen the control analysis, using longer evaluation episodes, plotting epidemic curves, and revealing the key factors that underlie the RL agent’s success. Third, we use COVID-specific spreading parameters and analyze the behavior of the policies beyond agent-based modelling. Finally, we perform several algorithmic changes aimed at improving efficiency, as detailed in Section 4.3.4.

4.3 Methodology

4.3.1 Simulating epidemics

Several epidemics are simulated using an SEIR compartmental formulation together with COVID-specific parameters: a base infection rate of $\beta = 0.0791$ and an average exposed duration of $\epsilon = 3.7$ days. The viral infection diffuses over multiple interaction network configurations, of various sizes, with events getting generated by either a multi-site mean-field or an agent-based model. Irrespective of the model, we allow all disease-unrelated events to be time-discretized, that is selecting an action or updating the active links set takes place every t_u days, with t_u set to 1. To eliminate random fluctuations that might obscure the comparative effectiveness of different control strategies, most of our experiments assume that nodes remain infectious for the whole duration of an episode unless they get isolated (i.e. recovery rate $\gamma = 0$). Intuitively, the impact of this assumption becomes significant only when the control scenario becomes *oversaturated* (i.e. the testing budget k and/or the recovery rate γ are large enough for most policies to achieve containment). For completeness, however, we also report epidemic outcomes for varying γ in Section 4.4.2, while Section 4.4.4 uses the same γ range to extract generalizable insights about the RL testing policy. We remark that these parameters could also be adjusted to fit the data trends of specific regions of interest, and we exemplify one possible approach to achieve this in Appendix D.2.

Most epidemics in this chapter are simulated using the multi-site mean-field model introduced in Chapter 3, relying on exponential waiting times efficiently sampled via Gillespie’s algorithm. The SEIR-based state transition probabilities for a given node v are defined as follows:

$$\begin{aligned} p(S \rightarrow E) &= \beta K^I \Delta t = \beta \sum_u A_{v,u}^t \mathbb{I}_u(t) \Delta t \\ p(E \rightarrow I) &= \epsilon^{-1} \Delta t \\ p(I \rightarrow R) &= \gamma \Delta t, \end{aligned} \quad (4.9)$$

where A^t is the time-dependent weighted adjacency matrix, $\mathbb{I}_u(t)$ is the infection indicator function of node u evaluated at time t , $\sum_u A_{v,u}^t \mathbb{I}_u(t)$ describes the mean-field infection effect over node v , while Δt is a short time interval.

In contrast, our agent-based models loop through all the edges (v, u) at every time increment, checking whether an infection event occurs over that connection, according to the transmission probability $\beta A_{v,u}^t \mathbb{I}_u(t)$. Concurrently, every node v gets visited, and the appropriate transition events get executed if the node waiting times $d_v \sim \mathcal{N}(\epsilon, 1)$, $r_v \sim \mathcal{N}(\frac{1}{\rho}, 1)$ have depleted.

4.3.2 Network configurations

The performance of the proposed public-health interventions is evaluated across multiple network configurations derived from various sources, such as classic random graph models (Erdős–Rényi, dual Barabási–Albert, Holme–Kim – discussed in Section 2.3), interaction data originating from the Social Evolution dataset (refer to Chapter 3 for details on how the weighted dynamic network is built), and contact tracing data partially conveyed in [Meirom et al. \(2021\)](#). Since the latter source anonymized the data so that it can only be accessed indirectly through its degree and edge infection probability distributions (reproduced in Eq 4.10), we employ a configuration model to reconstruct samples of the original interactions.

$$\begin{aligned} P(\text{deg}_{G_t}(v) = x) &= 2.68 \cdot \mathcal{N}(-4.44, 11.18) + 3.2 \cdot 10^{-3} \cdot x^{-0.36} \\ P(A_{v,u}^t = x) &= .47 \cdot \mathcal{N}(0.41, 0.036) + .53 \cdot \mathcal{B}e(5.05, 20.02) \end{aligned} \quad (4.10)$$

For dynamic graphs, the diffusion process depends on the current temporal snapshot of the graph, G_t . In contrast, in a static graph, the diffusion process is independent of time and operates on the entire graph G , such that $G_t = G$ for any t . Every edge (v, u) of G_t features a transmission weight $A_{v,u}^t$, and we calculate that interaction’s transmission potential by scaling the base infection rate β with this weight. If the latter cannot be inferred or directly sampled from the available data, as is the case for the random graph models, we assume $A_{v,u}^t \sim \mathcal{U}(0.5, 1)$. Typically, the resulting artificial

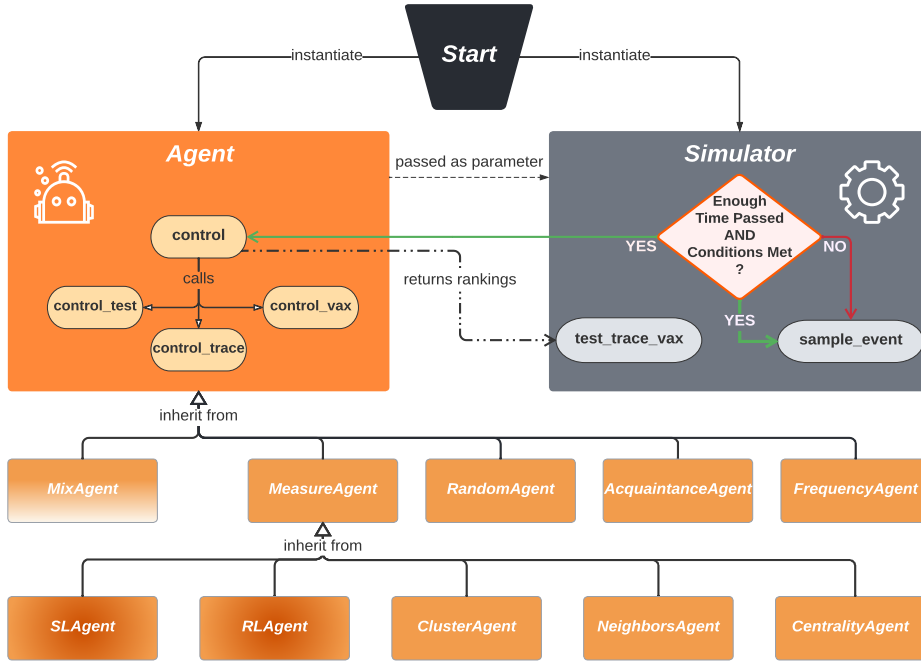


FIGURE 4.1: **Block diagram of our simulation-control framework.** The *Agent* is passed as a parameter to the *Simulator* instance, which is responsible for delegating calls to the *control(.)* method of the *Agent* once enough time has passed and other preset conditions are met (e.g. minimum infection threshold). The *control(.)* function performs preliminary node filterings, then calls *control_test(.)*, *control_trace(.)* and *control_vax(.)*. These methods are responsible for producing node rankings for testing, tracing and vaccinating, respectively, exhibiting a different logic in each *Agent* sub-type. Every *MeasureAgent* decides the final ranking via a node-level score. Both the *SLAgent* and *RLAgent* feature pluggable and trainable weights, sharing an identical evaluation mode. Combinations of agent behaviors for each of the three tasks can be selected by utilizing a *MixAgent*. Further details about each type of agent can be consulted in Sections 4.3.4 and 4.3.5.

networks are also static in nature. To simulate dynamic links in this case, a random fraction of edges is sampled at every time increment. The fraction is assumed to be uniformly distributed according to $\mathcal{U}(0.4, 0.8)$.

It should be noted that, in practice, the graph connections required for producing the fine-grained control policies described here would need to be inferred from specialized monitoring systems, such as digital tracing mechanisms (Meirom et al., 2021) or GPS-based mobility trackers (Serafino et al., 2022). Recently, large-scale contact tracing data has been used to reconstruct such transmission networks for the state of New York (Pei et al., 2022). After acquiring and anonymizing this interaction data, our method has the advantage of requiring no further personal information.

4.3.3 Control setup

Each epidemic is allowed to progress until at least c_d days have passed since the simulation began and a minimum of c_i nodes become infected before the agent

commences its interventions. In the first day of control, the agent is informed at random about the status of a proportion c_k of the *infected* population, after which it is only allowed to test k individuals and isolate k_c contacts of recently-detected positive nodes (i.e. in the previous 6 timestamps) per day. As the actor is not aware of a node’s state unless it is a part of c_k or it got tested recently, the environment is partially observable. In this work, we fix $c_a = 5$, $c_i = 5\%$ and $c_k = 25\%$, while the budgets are varied between experiments. A block diagram of our framework, which also illustrates the class hierarchy of the agents, is provided in Fig 4.1.

At evaluation, each agent is tasked to select the top- k nodes to test, top- k_c contacts to isolate, and top- k_v people to vaccinate at every time increment, according to their knowledge of the epidemic and graph states. Consequently, this constitutes an instance of the subset selection problem (Rayner et al., 2019), where nodes that are found to be positive, or that become traced/vaccinated by the system, are marked as isolating/immune, becoming incapable of infecting other participants. In principle, those individuals remain disconnected from the graph, yet we allow messages to continue flowing through their connections during the training phase of the learning-based agents. Importantly, the process of tracing is assumed to be carried with delays shorter than a day, which typically implies that a contact tracing application is deployed and running (Ferretti et al., 2020; Wymant et al., 2021). Moreover, we presume the detected positives and the traced vertices acquire immunity during self-isolation, either via exposure or vaccination. The immunity delay following vaccination without tracing is set to 21 days, the recommended timeframe between two doses of BNT162b2 (Polack et al., 2020) or mRNA-1273 (Baden et al., 2021), after which the immune response normally peaks. To evaluate a policy’s efficacy, we compare the fraction of nodes it kept healthy through the entire epidemics (i.e. complement of the attack rate), and the corresponding infection curves.

4.3.4 Baseline agents

In this work, we investigate a wide variety of baseline agents, each of which utilizes a distinct heuristic: *random* sampling (also designated *randag*), *acquaintance* sampling (or *acq*), centrality-based ranking (e.g. *degree* or *deg*, *eigenvector* or *eig*, *PageRank* or *prank*, *closeness*, *betweenness*), and *neighborhood*-based ranking (or *neigh*). The latter exploits some limited epidemic state information, targeting the nodes that have the highest number of positively-detected neighbors in their 2-hops vicinity via lexicographic ordering (Meirom et al., 2015).

In addition to the above, rule-based agents can also leverage another source of information: the occurrence frequency of nodes in the immediate proximity of identified cases. This leads to two new baselines specifically-designed for tracing: *frequency*, which randomly samples nodes with probabilities proportional to the

individual frequencies, being equivalent to the tracing mechanism studied in the mean-field approach of Chapter 3, and *backward*, which greedily picks the nodes with the highest registered frequencies (i.e. the approach of Kojaku et al. (2021)).

We also propose a simple yet powerful extension to all agents: recollection of recent negative test results. This effectively restricts the action space to untested nodes in the past t_n days, speeding up the network exploration. We set $t_n = 3$, an appropriate timeline for COVID-19 (Smith et al., 2021) that also renders good results empirically.

4.3.5 Learning agents

4.3.5.1 Overview

Our learning-based agents leverage two GNN components, a single-layered Diffusion module and a long-range Information module, followed by an MLP that computes the nodes embedding space h_t , and another MLP that produces the per-node ranking score. In spite of this, our proposed solution features several improvements or simplifications: First and foremost, we utilize a second output MLP, parallel to the first, that produces a full state score using the same embedding space as the latter. Secondly, we employ a GATv2 layer in the Diffusion module to leverage attention when aggregating embeddings from the immediate neighborhood of each node in the current snapshot G_t , and 3 GIN layers followed by an FA layer in the Information module to improve the model’s expressivity and the long-range information flow for the corresponding temporal graph $G_{t' \leq t}$ used as input. Finally, after experimenting with different normalization schemes for mitigating the previously-reported issue of exploding hidden states h_t , we propose the usage of standard scaling and clipping the gradient to a value of 0.5, which leads to sufficiently stable training behaviors. More details about our neural architecture can be consulted in Fig 4.2.

In addition to the architectural aspects described above, we carefully scrutinize different combinations of node features, choosing the following final set for training our policies: the degree and eigenvector centralities, the number of infected vertices in the 1-hop and 2-hop neighborhoods, five random features for breaking the structure symmetries, and four features related to the test state: a one-hot vector of size 3, marking the test status of node i at the previous timestamp (i.e. untested, negative or positive), and a binary value marking whether the vertex has ever tested positive. To allow for the hidden states to incorporate information from these features before the training commences, we disable gradient updates for the first $u_{interim} = 11$ passes, as suggested by Kapturowski et al. (2022) for stabilizing the learning process.

The ranking of nodes can be performed using either SL or RL agents, with little to no changes to the underlying neural network architecture. The SL agent is trained as a

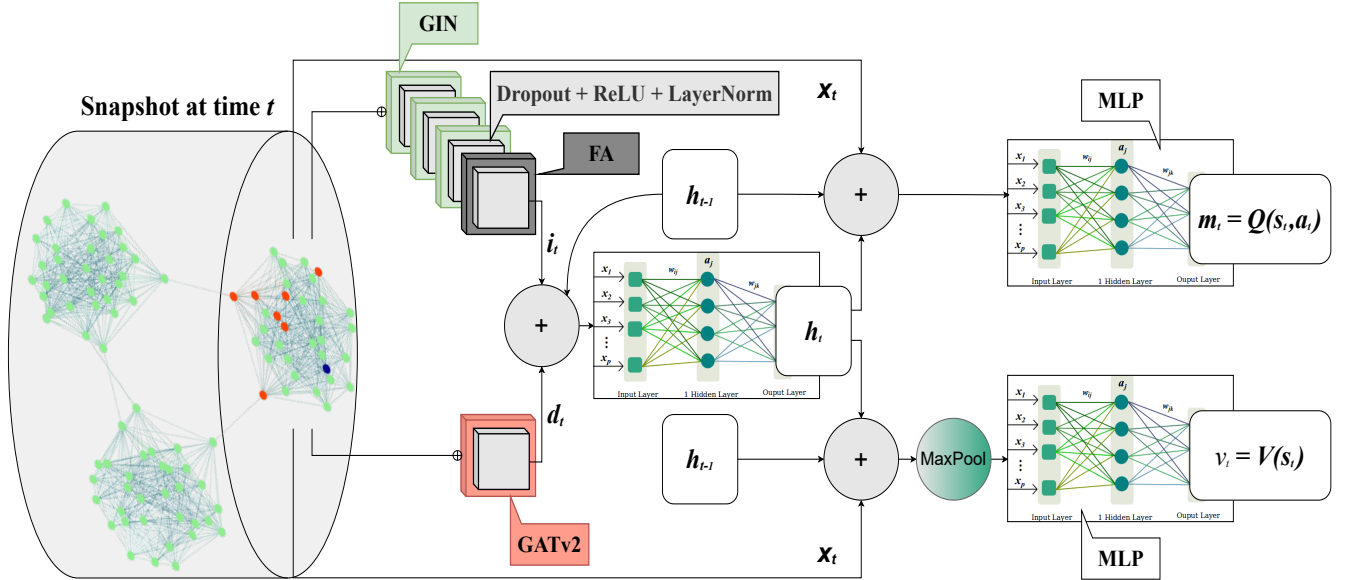


FIGURE 4.2: **Neural architecture for the learning-based control agents.** At each time increment t , the graph node features $x_t = f_v(t)$ and connections get passed to the Information and Diffusion modules, producing the outputs i_t and d_t , respectively. The Information component has 3 GIN layers and a final FA layer, interconnected by Dropout, ReLU and LayerNorm. The Diffusion module has a single GATv2 layer, followed by Dropout, ReLU and LayerNorm. i_t , d_t , and the node hidden states from the previous timestamp h_{t-1} are utilized as inputs for the MLP that outputs the new standard-scaled hidden states h_t . By using x_t , h_t , and h_{t-1} , two parallel MLPs produce the per-node action-values of the actor $m_t = Q(s_t, a_t)$, and the state value of the critic $v_t = V(s_t)$. The latter is obtained by applying a *MaxPool* per rows before passing the inputs to the afferent MLP. The *RLAgent* uses the m_t output to rank and sample nodes, and the v_t output to receive feedback, both being key components of the PPO objective. The *SLAgent* feeds m_t directly into a cross-entropy loss, using the infection indicators $\mathbb{I}_v(t)$ as labels. MLP sketch source:

<https://www.gabormelli.com/RKB/File:2NNw.png>.

simple node classifier by optimizing a binary cross-entropy loss on the infection status of each vertex, with the output space representing the next-step infection likelihood. In contrast, our RL agent gets optimized via a surrogate PPO objective, which only needs access to the total number of infected at each time point, ultimately solving for the criterion in Eq 4.11, where $E(t)$, $I(t)$ and $R(t)$ are the number of exposed, infectious and recovered, respectively, at time t , with $t \geq t_0$. We note that an agent may further benefit from receiving both the infection labels and counts during training, but investigating such a configuration remains beyond the scope of our current analysis.

$$\min \sum_{t=t_0}^T \gamma_r^{t-t_0} (E(t) + I(t) + R(t)) \quad (4.11)$$

Assuming the embedding space of a learning-based agent is informative enough, we posit that clustering it can reveal potential targets for mass public health interventions, such as vaccination campaigns. Following this logic, we propose a new baseline,

called *cluster* or *clust*, which ranks nodes according to the infection likelihood of their corresponding clusters. The likelihood can be obtained from past data by aggregating confirmed infections to date or by predicting via an SL agent, or from future data when the simulation can be rolled forward in time. This method has several benefits over directly using the SL/RL prediction values, as follows:

- It protects the privacy and security of individuals by avoiding singling them out based on their influence, preventing potential manipulations from other actors.
- It allows for versatility and flexibility in grouping nodes based on different criteria or objectives, depending on the scenario or context of the intervention.
- It reduces the impact of noise or uncertainty in the individual ranking values, which may arise from measurement errors, incomplete data, or dynamic changes in the network, thus making the intervention process more stable and reliable.
- It simplifies the network and the intervention process by grouping nodes into clusters, making the approach more scalable, especially for larger graphs.
- It provides backup and redundancy for the intervention in case certain individuals are not responsive to it (e.g. refusing vaccination), by having multiple nodes with the same ranking value in each cluster, increasing the chances of success in the real world.

4.3.5.2 Reinforcement learning routines

To simplify the notation, we use θ to denote the set of all the parameters of the ranking module, which comprises the MLP heads for both the actor and the critic, as well as the previous-step node hidden states h_{t-1} , with the mention that the latter always have a zero gradient. In contrast to the original formulation, Eqs 4.13 to 4.18 describe our proposed modification of PPO that allows for optimizing the objective in a memory-efficient online manner. In particular, we replace all advantage functions with the single-step advantage, $\hat{A}_t^{(\gamma_r, 0)}(a_t; \theta)$, and introduce an intermediary operation that accumulates the gradients of the modified loss into an unified eligibility trace (Sutton and Barto, 2018), in a similar fashion to Kobayashi (2022), ultimately obtaining a backward-view approximation of the GAE, $\hat{A}_t^{(\gamma_r, \lambda)}$. Before accumulating them in the trace, the gradients in Eq 4.16 can be adjusted by $x = \delta_t^{\gamma_r}(\theta')$, as suggested by Kobayashi (2022), alluding to the original mathematical form of online actor-critics given by Sutton and Barto (2018). However, in our environment, we observe better results when we assign $x = 1$. Next, the current value of the updated trace E_t gets multiplied in Eq 4.17 by $\delta_t^{\gamma_r}(\theta')$, and the outcome is altered by the function transform(.) in accordance with the chosen optimization routine (e.g. Adam by Kingma and Ba (2017), AdamW by Loshchilov and Hutter (2019)). Finally, Δ_t can accumulate

gradients across multiple time increments ($u_{interim} = 11$ in our experiments) to reduce the variance associated with updating the neural network after each step. We note that, by considering $r_t = r_t^{SARSA}$ from Eq 4.12, we can eliminate the requirement of storing s_t in memory for the subsequent timestamp, while retaining the benefits of the policy ratio clipping. The premise of this simplification is that substantial variations between s_{t+1} and s_t are infrequent, condition which aligns with our environment, resulting in negligible differences between the two ratios. Based on previous work and our own analysis, we choose to fix $\gamma_r = 0.99$, $\lambda = 0.97$, $\mathcal{E} = 0.2$, $c_1 = 0.5$, $c_2 = 0.01$, and update the target value network every $D = 5$ episodes across all our experiments.

$$r_t^{SARSA}(\theta') = \frac{\pi(a_t|s_{t+1};\theta')}{\pi(a_t|s_t;\theta)} \quad (4.12)$$

$$\mathcal{L}_t^{OCLIP}(\theta') = \min[r_t(\theta')\hat{A}_t^{(\gamma_r,0)}(a_t;\theta'), \text{clip}(r_t(\theta'), 1 - \mathcal{E}, 1 + \mathcal{E})\hat{A}_t^{(\gamma_r,0)}(a_t;\theta')] \quad (4.13)$$

$$\mathcal{L}_t^{OVF}(\theta') = [\hat{A}_t^{(\gamma_r,0)}(a_t;\theta')]^2 \quad \mathcal{H}_t(\theta') = - \sum_{a \in A} \pi(a|s_t;\theta') \log \pi(a|s_t;\theta') \quad (4.14)$$

$$\mathcal{L}_t^{OPPO}(\theta') = -\mathcal{L}_t^{OCLIP}(\theta') + c_1\mathcal{L}_t^{OVF}(\theta') - c_2\mathcal{H}_t(\theta') \quad (4.15)$$

$$E_t = \gamma_r\lambda E_{t-1} + \frac{\nabla_{\theta'}\mathcal{L}_t^{OPPO}(\theta')}{x}, \text{ with } x = \delta_t^{\gamma_r}(\theta') \text{ or } x = 1 \quad (4.16)$$

$$\Delta_t = \Delta_{t-1} + \text{transform}(\delta_t^{\gamma_r}(\theta')E_t) \quad (4.17)$$

$$\theta_{n+1} = \begin{cases} \theta_n - \Delta_t & \text{if } t \bmod u_{interim} = 0 \\ \theta_n & \text{otherwise.} \end{cases} \quad (4.18)$$

In each step t , the agent's ranking module outputs the actor's raw node likelihoods m_t and the critic's evaluation of state s_t , v_t , using the current parameters θ_n . An action a_t is then sampled according to π and m_t , while its corresponding log-probability, $\log \pi(a_t|s_t;\theta_n)$, is calculated. At the same time, the agent receives the reward R_{t-1} for its previous action a_{t-1} , as the *Simulator* reveals the total number of infected. Using these quantities, we outline below two alternative approaches that can be used for the online parameter iteration in Eq 4.18. Note that we increment the subscript of θ only when the parameters of the ranking module itself are updated, regardless of the fact that h_{t-1} changes at every timestamp.

1. For every $t \geq 1$, the gradients in E_{t-1} and Δ_{t-1} are accumulated according to $\mathcal{L}_{t-1}^{OPPO}(\theta')$. The latter is computed using $v(s_{t-1};\theta')$ and $v(s_t;\theta_D)$, given the old state s_{t-1} in the single-step buffer, R_{t-1} , and either r_{t-1} or r_{t-1}^{SARSA} , given the old log-probability $\log \pi(a_{t-1}|s_{t-1};\theta)$ recorded in the buffer. The current state s_t , the action taken a_t , and its corresponding log-probability then continue to replace the contents of the buffer for the subsequent timestamps in a loop, becoming their estimates for $t - 1$, for the rest of the episode. When $u_{interim}$ iterations have elapsed since the last update, but *after* the action a_t has been sampled, the agent's

policy gets updated to θ_{n+1} according to the information present in Δ_t , including the gradients pertaining to the current timestamp. As such, θ' has continued to be θ_n from the last update until the step that follows this new update, when it becomes θ_{n+1} . Consequently, the steps $t + 1$, where $t \bmod n_{interim} = 0$, have $\theta = \theta_n$ as the acting policy, but use $\theta' = \theta_{n+1}$ to compute the gradients.

2. On *even* control timestamps, the current information, including a_t , $\log \pi(a_t|s_t; \theta_n)$, and the evaluation v_t , is added to the single-step buffer, becoming the estimates for $t - 1$ in the subsequent step. The gradients, however, are only computed on the *odd* control timestamps, getting accumulated in E_{t-1} and Δ_{t-1} according to $\mathcal{L}_{t-1}^{OPPO}(\theta')$, with parameter updates happening after every $u_{interim}$ interval has elapsed. In this formulation, both the acting policy θ and the policy to be updated θ' correspond to a version of θ_n , thus having the same ranking module parameters, but differing in the previous-step hidden states h . As a consequence, we can use v_{t-1} from the buffer to compute the gradients w.r.t. \mathcal{L}_{t-1}^{OVF} , while employing the alternative ratio r_{t-1}^{SARSA} to avoid the requirement of adding the previous state to the buffer. If r_{t-1} is utilized, however, the policy needs to be re-evaluated at s_{t-1} to compute the afferent log-probabilities.

When implemented in our epidemic setting, both of these techniques generate adequate policies, surpassing the other baselines considered. However, the second approach seems to yield better and more stable results. As such, we reproduce results only for the latter training approach in this thesis, but we acknowledge that other contexts might benefit more from the former method.

To ensure sufficient exploration during training, the RL policy π passes the raw outputs m_t through a softmax(.) function with temperature parameter ϵ_{ep} , starting from $\epsilon_0 = 0.5$ and decaying every episode by 0.99. We note that more complex strategies could also be suitable here, including the transforms proposed by Mei et al. (2020) or Meirum et al. (2021), but our simple alternative has proven to be sufficiently effective at exploring the state space. During evaluation, the sampling process is turned off, greedy actions are taken instead, and the edges connected to positively-identified vertices are masked before being fed to the Information module, thus limiting feature oversmoothing. In contrast, the single-layer Diffusion GNN is allowed to ‘observe’ the aforesaid links, yet only for the last snapshot, enabling the features associated with recent positive detections to propagate through the network.

4.3.6 Simulation-control framework pseudocode

The logic behind our epidemic control framework in the continuous-time simulation scenario is outlined in Algorithm 2. The class hierarchy of the agents, together with

TABLE 4.1: Legend for the control framework pseudocode.

Name(s)	Description(s)
M, M_D	GNN-based ranking and target models (shared across epidemics).
E_{conf}	Episode configuration. Consists of tuples mapping an episode ID ep to its exploration-control variable ϵ_{ep} .
S_{conf}	Simulation configuration. Enumerated type that defines the maximum network, infection and event seeds, which in turn control the range of the loops over each seeded configuration.
s_{net}, s_{inf}, s_{ev}	Interaction network, infection and event seeds.
N_p, S_p, A_p	Interaction network, Simulator and Agent hyperparameters. A_p contains the sampling strategy st and learning rate lr .
N, S, A	Interaction network, Simulator and Agent main objects.
i_{cu}, i_{eu}	Iterators for time-discretized events: dynamic control and edge-updating interaction events.
e, t_e	Interaction event enumerated type and its corresponding time of occurrence.
k_t, k_c, k_v	Daily budgets for testing, tracing, and vaccination.
c_t, c_c, c_v	Sensible candidates to rank for testing, tracing, and vaccination.
n_t, n_c, n_v	Nodes chosen by the agent for testing, tracing, and vaccination.
d	Boolean that determines whether the action is sampled or greedily taken from top-k ranking.
sg	Sampling strategy employed by the RLAgent. This can be one of the following: 'softmax', 'escort-transform' Mei et al. (2020), 'nvidia-explore' Meirum et al. (2021).
m	Node ranking scores computed by a specific agent.
v	Epidemic state score computed by the GNN ranking model.
B_∞, B_1	Single-step for online training, and episodic replay buffer for offline training of an RLAgent.
O	Optimizer object that updates parameters according to a predefined rule (e.g. Adam, AdamW).
$a, \log \pi_a$	Sampled action and its corresponding log of probability.
R_{t-1}	Reward of previous action taken (i.e. for action sampled and executed at time $t - 1$).

their high-level logic, can be consulted in Algorithm 3. Refer to Table 4.1 for details regarding the variables involved in these routines.

Algorithm 2 Simulation-control logic.

```

1: global variables
2:    $M, M_D$             $\triangleright$  GNN ranking model and target model
3: end global variables
4: procedure RUN_EPIDEMIC( $E_{conf}, S_{conf}, N_p, S_p, A_p$ )
5:   for each ( $ep, \epsilon_{ep}$ )  $\in E_{conf}$  do            $\triangleright$  Episode ID and  $\epsilon_{ep}$ 
6:     for  $s_{net} \in \{0, \dots, S_{conf}.MAX\_NET\_SEED\}$  do
7:        $\triangleright N$  has all node states and edges over time
8:        $N \leftarrow INIT\_NET(N_p, s_{net})$ 
9:       for  $s_{inf} \in \{0, \dots, S_{conf}.MAX\_INF\_SEED\}$  do
10:         $S \leftarrow INIT\_SIMULATOR(S_p, s_{inf}, N)$ 
11:         $A \leftarrow INIT\_AGENT(A_p, M, M_D, ep, \epsilon_{ep})$ 
12:         $i_{cu} \leftarrow 0$             $\triangleright$  Iterator for control update
13:         $i_{eu} \leftarrow 0$             $\triangleright$  Iterator for edge update
14:        for  $s_{ev} \in \{0, \dots, S_{conf}.MAX\_EVENT\}$  do
15:           $e \leftarrow S.SAMPLE\_NEXT\_EVENT()$ 
16:           $t_e \leftarrow e.TIME$ 
17:           $S.RUN\_EVENT(e, N)$ 
18:          if  $S.DO\_CONTROL(N, t_e, i_{cu})$  then
19:             $i_{cu} \leftarrow \lfloor t_e \rfloor$             $\triangleright$  Floor function
20:             $(n_p, n_c, n_v) \leftarrow A.CONTROL(N, i_{cu})$ 
21:             $S.UPDATE\_STATES(N, n_p, n_c, n_v)$ 
22:             $i_{cu} \leftarrow i_{cu} + 1$ 
23:          end if
24:          if  $S.DO\_UPDATE\_EDGES(N, t_e, i_{eu})$  then
25:             $i_{eu} \leftarrow \lfloor t_e \rfloor$ 
26:             $N.UPDATE\_EDGES(i_{eu})$ 
27:             $i_{eu} \leftarrow i_{eu} + 1$ 
28:          end if
29:        end for
30:         $\triangleright$  Log and update offline parameters (if any)
31:         $A.FINISH(N)$ 
32:      end for
33:    end for
34:  end for
35: end procedure

```

Algorithm 3 Control agents' hierarchy.

```

36: struct AGENT
37:    $k_t, k_c, k_v$    ▷ Budgets for testing, tracing, and vaccination
38:    $t, t_p, t_n$    ▷ Control time; positive & negative test recollection
39:   procedure CONTROL( $N, i_{cu}$ )
40:      $t \leftarrow i_{cu}$    ▷ Set current control time
41:     ▷ Find candidates based on time  $t$ , states in  $N$ , and  $t_n$ 
42:      $c_t \leftarrow$  CANDIDATE_TEST( $N, t_n, t$ )
43:     ▷ Calls CONTROL_ALL by default; can be overridden
44:      $n_t \leftarrow$  CONTROL_TEST( $N, c_t, k_t$ )
45:     ▷ Test and split nodes into positives & negatives
46:      $(n_p, n_n) \leftarrow$  TEST_NODES( $N, n_t$ )
47:     UPDATE_KNOWLEDGE( $n_p, n_n, t_p, t_n$ )
48:      $c_c \leftarrow$  CANDIDATE_TRACE( $N, t_p, t$ )   ▷ Filter  $t_p$  contacts
49:      $n_c \leftarrow$  CONTROL_TRACE( $N, c_c, k_c$ )
50:      $c_v \leftarrow$  CANDIDATE_VAX( $N, n_p, n_c, t$ )   ▷ Except  $n_p, n_c$ 
51:      $n_v \leftarrow$  CONTROL_VAX( $N, c_v, k_v$ )
52:     return  $(n_p, n_c, n_v)$ 
53:   end procedure
54: end struct
55: struct MEASUREAGENT(AGENT)
56:    $d$    ▷ Boolean controlling if sampling or top-k ranking
57:   procedure CONTROL_ALL( $N, c, k$ )
58:     ▷ Compute score for each node in  $c$ ; RL samples  $k$  nodes
59:      $m \leftarrow$  COMPUTE_MEASURES( $N, c, k$ )
60:     if  $d$  then
61:       return  $c[\text{ARGTOPK}(m, k)]$    ▷ Heap sort for top-k
62:     else
63:       return  $m$    ▷ In this case,  $m$  is a list of sampled nodes
64:     end if
65:   end procedure
66: end struct
67: struct SAGENT(MEASUREAGENT)
68:    $lr$    ▷ Learning rate; if 0, evaluation mode is assumed
69:   procedure COMPUTE_MEASURES( $N, c, k$ )
70:     if  $lr > 0$  then
71:        $(m, v) \leftarrow$  M.FORWARD( $N$ )   ▷ Message passing
72:       BACKPROP_LOSS( $N, m$ )   ▷ BCE on infection status
73:     else
74:        $(m, v) \leftarrow$  M.FORWARD(SUBGRAPH( $N, c$ ))
75:     end if
76:     return  $m$ 
77:   end procedure
78: end struct
79: struct RLAGENT(MEASUREAGENT)
80:    $lr$    ▷ Learning rate; if 0, evaluation mode is assumed
81:    $sg$    ▷ Action sampling strategy (e.g. softmax)
82:    $\epsilon_{ep}$  ▷ Action sampling noise (i.e. softmax temperature)
83:    $B_\infty$  ▷ Episodic replay buffer; if null, conduct online learning
84:    $B_1$    ▷ Single-step replay buffer for online learning
85:   procedure COMPUTE_MEASURES( $N, c, k$ )
86:     if  $lr > 0$  then
87:       ▷ Reward of previous action (if any)
88:        $R_{t-1} \leftarrow -N.\text{NUM\_INFECTED}()$ 
89:        $(m, v) \leftarrow$  M.FORWARD( $N$ )   ▷ Message passing
90:        $(a, \log \pi_a) \leftarrow$  SAMPLE( $m, c, k, sg, \epsilon_{ep}$ ) ▷ Sample action
91:       ▷ Existence of  $B_\infty$  determines training online/offline
92:       if  $B_\infty$  is null then
93:         ONLN_TRAIN( $R_{t-1}, N, a, \log \pi_a, m, v$ )
94:       else
95:         ▷ Add  $(R_{t-1}, s_t, a_t, \log \pi_{a_t}, V_t)$  to replay buffer
96:          $B_\infty.\text{ADD}(R_{t-1}, N, a, \log \pi_a, v)$ 
97:       end if
98:        $m \leftarrow a$ 
99:     else
100:       $(m, v) \leftarrow$  M.FORWARD(SUBGRAPH( $N, c$ ))
101:    end if
102:    return  $m$ 
103:  end procedure
104: end struct

```

Algorithm 4 RL online routines.

```

105: struct RLAGENT.ONE(RLAGENT)
106:    $O$    ▷ Optimizer object that updates parameters
107:   procedure ONLN_TRAIN( $R_{t-1}, N, a, \log \pi_a, m, v$ )
108:     if  $B_1$  is not empty then
109:       ▷ Evaluate state  $s_t$  using the target policy
110:        $(m_D, v_D) \leftarrow$  M_D.FORWARD( $N$ )
111:       ▷ Evaluate  $s_{t-1}$  with current policy
112:        $(m_A, v_A) \leftarrow$  M.FORWARD( $B_1.N$ )
113:       ▷ Infer policy ratio, entropy, TD-error
114:        $r_{t-1} =$  GET_RATIO( $m_A, sg, B_1.a, B_1.\log \pi_a$ )
115:        $\mathcal{H} =$  GET_ENTROPY( $m_A, sg$ )
116:        $\delta_{t-1} \leftarrow R_{t-1} + v_D - v_A$ 
117:       ▷ Calculate and accumulate gradients
118:       BACKPROP_LOSS( $r_{t-1}, \delta_{t-1}, \mathcal{H}$ )
119:       if  $t \bmod u_{interim} = 0$  then
120:          $O.\text{step}()$ 
121:       end if
122:        $B_1.\text{CLEAR}()$ 
123:     end if
124:      $B_1.\text{ADD}(N, a, \log \pi_a)$ 
125:   end procedure
126: end struct
127: struct RLAGENT.TWO(RLAGENT)
128:    $i_{cont}$  ▷ Special iterator for control that starts at 0
129:    $r_{alt}$    ▷ Boolean for using  $r_t^{\text{SARSA}}$  or standard  $r_t$ 
130:   procedure ONLN_TRAIN( $R_{t-1}, N, a, \log \pi_a, m, v$ )
131:     if  $i_{cont} \bmod 2 = 1$  then
132:        $(m_D, v_D) \leftarrow$  M_D.FORWARD( $N$ )
133:     if  $r_{alt}$  then
134:       ▷ Use likelihoods for  $s_t$  and old  $V$ 
135:        $m_A \leftarrow m$ 
136:        $v_A \leftarrow B_1.v$ 
137:     else
138:       ▷ Evaluate  $s_{t-1}$  with current policy
139:        $(m_A, v_A) \leftarrow$  M.FORWARD( $B_1.N$ )
140:     end if
141:     ▷ Infer policy ratio, entropy, TD-error
142:      $r_{t-1} =$  GET_RATIO( $m_A, sg, B_1.a, B_1.\log \pi_a$ )
143:      $\mathcal{H} =$  GET_ENTROPY( $m_A, sg$ )
144:      $\delta_{t-1} \leftarrow R_{t-1} + v_D - v_A$ 
145:     ▷ Calculate and accumulate gradients
146:     BACKPROP_LOSS( $r_{t-1}, \delta_{t-1}, \mathcal{H}$ )
147:     if  $t \bmod u_{interim} = 0$  then
148:        $O.\text{step}()$ 
149:     end if
150:   else
151:      $B_1.\text{ADD}(R_{t-1}, N, a, \log \pi_a, v)$ 
152:   end if
153:    $i_{cont} \leftarrow i_{cont} + 1$ 
154: end procedure
155: end struct

```

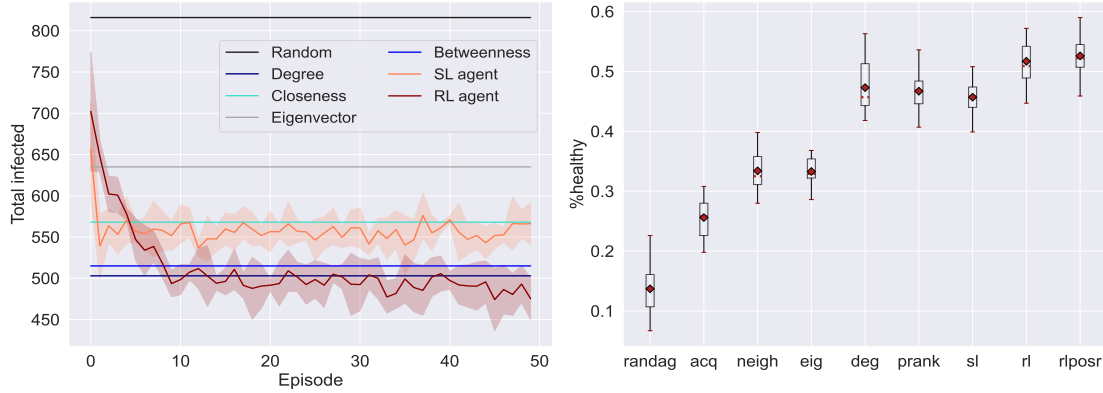


FIGURE 4.3: **Number of infected during training (left) and fraction kept healthy during evaluation (right) of learning- and heuristic-based agents.** This is reported for dual Barabási-Albert networks with $N = 1000$ and mean degree ≈ 3 . Test budget set to $k = 2$. The left panel shows results for a single network for consistency, while the right panel provides boxplots from 7 different networks and 7 infection seeds, marking the Q1, Q2 and Q3 quartiles, the means, and the extremity points without outliers from these runs. All agents have recollection of negative test results. Since the training behavior of the SL and RL agents tends to fluctuate across episodes and configurations, we plot the average curves and standard deviations from multiple seeded runs.

4.4 Results and discussion

4.4.1 Initial comparison of learning-based variants

Comparing the containment levels achieved by the SL and RL agents during training with the centrality-based actors with recollection reveals a significant performance difference, as evidenced by Fig 4.3. While the RL policy outperforms all baselines in several episodes, even before entering evaluation mode (i.e. when exploration would be disabled), the SL policy falls short. The former’s inadequacy is also apparent in the evaluation plot of Fig 4.3, as well as in the ample comparison of [Meirom et al. \(2021\)](#). This may be a consequence of not optimizing directly against the infection counts, while also lacking exploration mechanisms in training. As such, we opt to confine our analysis to the policies generated by RL, comparing them against the other baselines.

In our environment, two reward functions are possible for the RL agent: the negative of the infected count or the susceptible count at time t (denoted in plots as *rl* and *rlpos*, respectively). The performance between the two can vary due to numerical reasons, but these differences tend to be small (see Fig 4.4). As a consequence, we only display the former in our tables of results for providing a better readability.

4.4.2 Prioritizing testing in static graphs

Next, we investigate our agents’ performance in the context of targeted testing campaigns. To that end, we record the fraction of nodes kept healthy throughout




FIGURE 4.4: Infection control performance on different static network architectures and sizes, with a budget of $k = 2$. Uncertainties are shown as boxplots.

several epidemics, diffused across numerous static network topologies, when the budget of daily testing k is fixed, while k_c and k_v are set to 0. As stated before, most of our setups assume nodes do not spontaneously become uninfected (i.e. $\gamma = 0$), but, for completeness, we also present results for different recovery rates in Table 4.2.

Despite being trained for only 50 episodes on a single epidemic configuration spanning a preferential attachment network of 1000 nodes, our RL agent consistently outperforms the other baselines across a wide range of different network sizes (see Table 4.3), wiring configurations (Fig 4.4), and budgets (Fig 4.5). Interestingly, the learning-based agents pose a great generalization capability when the daily budgets scale with the number of nodes, enabling their deployment in larger networks, irrespective of the training graph size, and without a significant loss in efficacy.

TABLE 4.2: Fraction kept healthy with budget $k = 1\%$ and different recovery rates. Average over 5 seeded runs for each of the considered 5 realizations of dual Barabási-Albert networks with $N = 1000$ nodes and a mean degree of approximately 3. ‘w/ R’ denotes agents with recollection of recent negative test results.

Agents	$\gamma = 0$	$\gamma = 0.01$	$\gamma = 0.02$	$\gamma = 0.03$
Degree	0.555 ± 0.027	0.616 ± 0.034	0.662 ± 0.039	0.697 ± 0.039
Degree (w/ R)	0.744 ± 0.032	0.769 ± 0.028	0.801 ± 0.028	0.847 ± 0.025
PageRank (w/ R)	0.720 ± 0.026	0.755 ± 0.023	0.792 ± 0.037	0.834 ± 0.039
RL	0.822 ± 0.033	0.846 ± 0.026	0.876 ± 0.026	0.897 ± 0.026




FIGURE 4.5: Infection control performance on different static network architectures with varying budgets. The uncertainties are shown as boxplots.

TABLE 4.3: Fraction kept healthy with budget $k = 1\%$ and different population sizes. Average over 5 seeded runs for each of the considered 5 realizations of dual Barabási-Albert networks with a mean degree of approximately 3. ‘w/ R’ denotes agents with recollection of recent negative test results. Here, a single model is trained for 50 episodes on a network of size 1000, but its policy is able to generalize to appreciably larger graphs.

Agents	N = 1000	N = 2000	N = 5000	N = 20000
Degree	0.555 ± 0.027	0.552 ± 0.017	0.567 ± 0.027	0.557 ± 0.005
Degree (w/ R)	0.744 ± 0.032	0.736 ± 0.027	0.737 ± 0.025	0.746 ± 0.009
PageRank (w/ R)	0.720 ± 0.026	0.724 ± 0.021	0.729 ± 0.021	0.725 ± 0.008
RL	0.822 ± 0.033	0.811 ± 0.024	0.821 ± 0.025	0.803 ± 0.026

4.4.3 Prioritizing testing in dynamic graphs

In practice, a human’s interaction patterns are prone to change over time. While the previous section presents scenarios in which the links between nodes remain fixed throughout the simulation, Fig 4.6 features boxplots summarizing the average results obtained by our agents on several preferential attachment networks where the set of *active* edges changes every day. The RL agent is retrained to adapt to this dynamic context, restricting the GNN message passing routine to the most recent connections only. A parameter t_b controls how far back in time the restriction applies, and from our analysis, better results are achieved when this is set to less than 3 days, i.e. the input temporal graph takes the form $G_{t_b \leq t' \leq t}$.




FIGURE 4.6: Infection control performance on dynamic network architectures. Suffix ‘nor’ shows variants with no recollection.

TABLE 4.4: Fraction kept healthy for a 1000-node network built from real tracing statistics. Results are averaged over 5 runs for each of the 5 realizations of a configuration model using these statistics.

Agents	$k = 20$	$k = 50$
Acquaintance (w/ R)	0.465 ± 0.086	0.736 ± 0.085
Degree (w/ R)	0.406 ± 0.020	0.746 ± 0.025
Eigenvector (w/ R)	0.186 ± 0.013	0.409 ± 0.026
PageRank (w/ R)	0.363 ± 0.016	0.668 ± 0.039
RL	0.506 ± 0.029	0.831 ± 0.047

The top performing policies are also evaluated on dynamic networks built using the contact tracing statistics published by [Meirom et al. \(2021\)](#). The resulting mean containment rates are shown in Table 4.4.

4.4.4 Testing performance analysis

Consistent with the findings of [Meirom et al. \(2021\)](#), we note that the RL approach does not always test the most positive nodes among the agents, yet it still surpasses the other methods in terms of the epidemic outcome. We believe this seemingly paradoxical result can be attributed to two determining factors: individuals that are more likely to initiate threatening transmission cascades get targeted first, and the timing of detection is faster on average. Using the most competitive heuristic method as a reference, the degree centrality agent with negative test recollection, we investigate their mechanisms and occurrence.

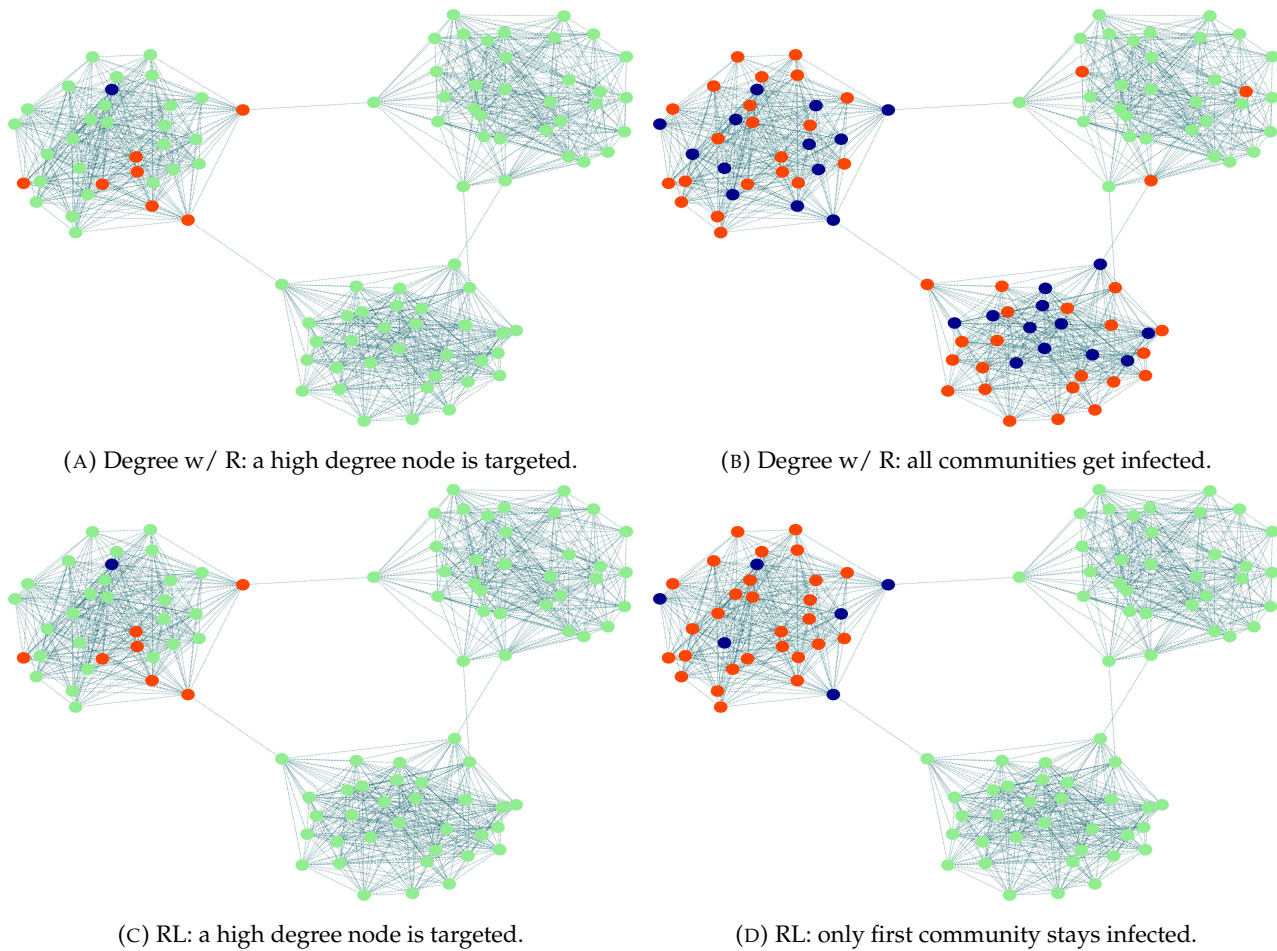


FIGURE 4.7: **Visualization of the spread for the degree w/ R and the RL agents.** This corresponds to a stochastic-block network [Holland et al. \(1983\)](#) with three communities. Susceptibles are green (light), infectious red (darker), and detected blue (the darkest). Initially, the two policies target the same nodes, but later on the RL agent preferentially tests the bridges.

Being a direct consequence of the node feature set, the first determinant is arguably the most easily grasped by intuition: a node with a high degree and/or eigenvector centrality that is in the vicinity of several positive cases is more likely to be both a superspreader and among the top-ranked vertices by our learning-based agents. This goes further for the RL agents, however, as they can also infer transmission risk patterns that are not readily quantifiable, possibly spanning over long temporal and graph ranges. Fig 4.7 illustrates a scenario where an RL tester initially targets the same node as the degree-directed policy, but subsequently proceeds to testing the bridging vertices, despite never computing any time-consuming betweenness centrality. By focusing on the more probable initiators of long transmission chains, the pathogen is effectively confined to the first cluster, whereas the degree agent appears incapable of averting the infection of every community.

While the second factor can be reliably quantified numerically, its relative impact on transmission can often be obscured by higher recovery rates. We plot in Fig 4.8 the




FIGURE 4.8: **Detection-infection time lag histograms.** Results are summed up across 49 simulated epidemics, and displayed per each recovery rate γ .

histogram of time lags between infection and detection by each testing agent recorded across 49 epidemic seeds under different recovery regimes, setting $N = 1000$, $k = 5$, and $k_c = k_v = 0$. While the degree agent is shown to be effective at detecting individuals within the first 4 days of infection, its performance drops significantly beyond this timeframe. In contrast, the RL agent keeps identifying high numbers of infections within 10 days of exposure, after which point a lighter tail emerges. As such, the two approaches end up detecting a similar number of positive nodes within the first 2 weeks of infection, despite the fact that the RL policy tests and isolates less people on average (see Fig 4.9). As the recovery rate increases, however, these differences become less pronounced, yet the gain in isolation efficiency remains above 10% for $\gamma \leq 0.05$.

4.4.5 Targeted test and trace programmes

Next, we investigate the extent to which different combinations of agents tasked with conducting testing and contact tracing, under the constraints of a fixed budget, can reduce the spread of a pathogen. For this problem, we train RL agents for either 50 or 200 episodes on the same testing task as before, comparing the resulting policies against the proposed baselines. Tables 4.5 and 4.6 confirm the RL tester improves the overall quality of ‘test and trace’ programmes, irrespective of the chosen tracer. That being said, employing the same agent to perform the ranking of contacts as well generally improves the containment.

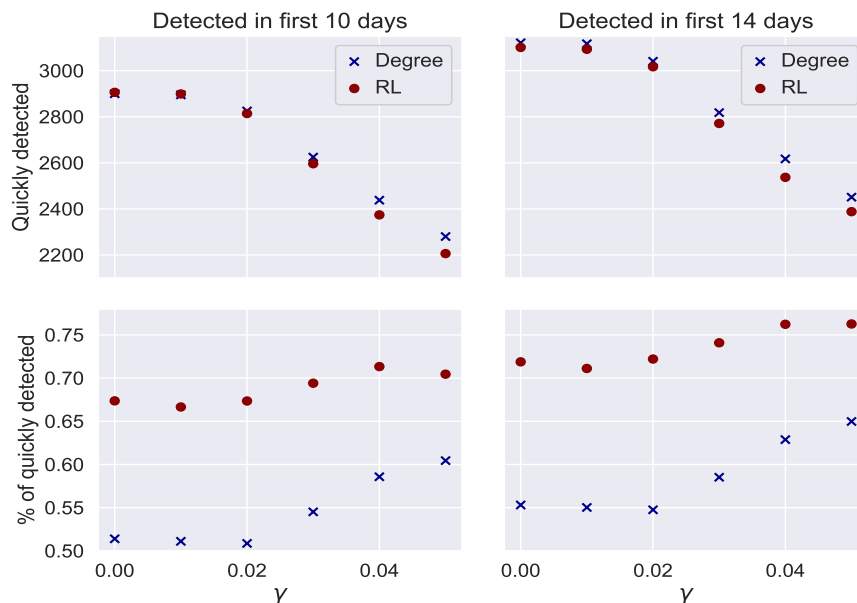


FIGURE 4.9: **Cumulative detection within X days since exposure.** Second row displays results normalized by the average total number of detected.

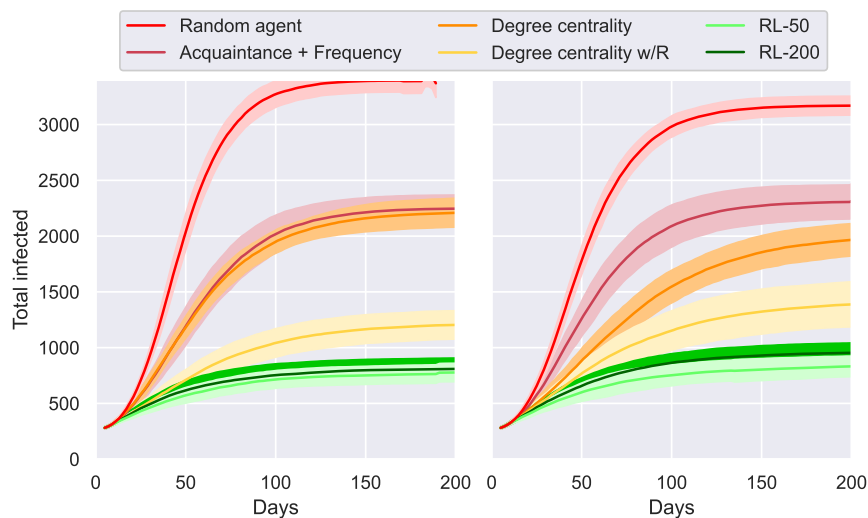


FIGURE 4.10: **Averaged epidemic curves and standard deviations during test and trace control.** These are for 5000 nodes dual Barabási-Albert graphs featuring a mean degree of approximately 3, with a daily testing budget of $k = 1\%$ and no tracing on the left, and $k = 10$ with a limit of $k_c = 25$ traced contacts on the right. Results displayed for two RL agents: one trained for 50, and the other for 200 episodes.

We also inspect the averaged epidemic curves associated with these targeted test and trace campaigns when $N = 5000$. The results obtained by each agent in the full environment is shown on the second column of Fig 4.10, with the first serving as a test-only reference. As stated before, heuristics with recollection bring large improvements over random policies, yet the RL agents outcompete them in most setups. Note the performance of $k = 50$ tests is similar to $k = 10$ tests, but tracing up to $k_c = 25$ contacts every day. While the balance between these processes will depend on various factors (see Appendix C), these results highlight the effectiveness of tracing.

TABLE 4.5: **Percentage of nodes kept healthy for graphs of size 1000 and an approximate mean degree of 3 (or $m_{HK} = 3$ for Holme-Kim), with budgets set to $k = 2$, $k_c = 5$. Averages over 5 runs for each of the considered 5 realizations of the following: dual Barabási-Albert with $m_1 = 5$, $m_2 = 1$ (BA 5-1) and $m_1 = 10$, $m_2 = 1$ (BA 10-1), Holme-Kim (HK), and Erdős-Rényi (ER).**

Agents (Test + Trace)	BA 5-1	BA 10-1	HC	ER
Random + Random	0.389 ± 0.044	0.446 ± 0.060	0.131 ± 0.023	0.199 ± 0.023
Random + Frequency	0.387 ± 0.033	0.465 ± 0.059	0.202 ± 0.030	0.195 ± 0.022
Acquaintance (w/ R) + Random	0.541 ± 0.054	0.657 ± 0.054	0.212 ± 0.031	0.215 ± 0.017
Acquaintance (w/ R) + Frequency	0.582 ± 0.055	0.674 ± 0.059	0.228 ± 0.040	0.217 ± 0.021
Acquaintance (w/ R) + Backward	0.591 ± 0.056	0.769 ± 0.080	0.213 ± 0.039	0.208 ± 0.019
Acquaintance (w/ R) + RL	0.644 ± 0.048	0.806 ± 0.058	0.248 ± 0.038	0.217 ± 0.018
Degree (w/ R) + Degree	0.764 ± 0.038	0.915 ± 0.032	0.528 ± 0.053	0.333 ± 0.037
RL + Random	0.818 ± 0.034	0.882 ± 0.026	0.542 ± 0.050	0.438 ± 0.043
RL + Frequency	0.832 ± 0.035	0.890 ± 0.033	0.567 ± 0.054	0.448 ± 0.048
RL + Backward	0.849 ± 0.033	0.923 ± 0.023	0.590 ± 0.058	0.434 ± 0.047
RL + Degree	0.853 ± 0.034	0.928 ± 0.014	0.614 ± 0.055	0.453 ± 0.039
RL + RL	0.876 ± 0.025	0.936 ± 0.009	0.620 ± 0.050	0.451 ± 0.039

TABLE 4.6: **Percentage of nodes kept healthy when controlling epidemics over a dynamic real interaction network.** Graph consists of 74 vertices, and is reconstructed from the Social Evolution dataset Madan et al. (2012). Averages over 5 runs for each of the considered 5 infection seeds. Test budget is $k = 2$.

Agents (Test + Trace)	$k_c = 2$	$k_c = 4$
Random + Frequency	0.511 ± 0.130	0.659 ± 0.114
Acquaintance (w/ R) + Frequency	0.494 ± 0.113	0.649 ± 0.089
Acquaintance (w/ R) + Backward	0.522 ± 0.115	0.654 ± 0.126
Neighborhood (w/ R)	0.620 ± 0.108	0.704 ± 0.107
Degree	0.614 ± 0.107	0.741 ± 0.084
Degree (w/ R)	0.636 ± 0.104	0.750 ± 0.084
RL	0.711 ± 0.089	0.773 ± 0.069

4.4.6 Agents interacting with different spreading dynamics

To assess the generalization capability of the agents, we compare their achieved containment rates when controlling epidemics generated by both a multi-site mean-field and an agent-based model run with similar hyperparameters. In this case, the RL agent reuses the parameters it learned from the previous experiments.

Despite the fact that the control mechanism in the mean-field case requires discretizing a continuous-time process, we observe minor differences between the two simulation approaches (Tables 4.7 and 4.8). This result confirms not only that the two models are analogous, but also that our agents transfer well across these spreading dynamics.

TABLE 4.7: **Fraction kept healthy for 2000 nodes and an average degree of approximately 3.** Results represent averages over 5 runs for each of the considered 5 instances of a dual Barabási-Albert model ($m_1 = 10, m_2 = 1$). Testing budget is $k = 2$ and no contact tracing is conducted.

Agents	Multi-site mean-field	Agent-based
Random	0.164 ± 0.037	0.195 ± 0.034
Acquaintance (w/ R)	0.251 ± 0.033	0.263 ± 0.035
Degree	0.390 ± 0.032	0.394 ± 0.029
Degree (w/ R)	0.443 ± 0.032	0.457 ± 0.034
RL	0.468 ± 0.035	0.477 ± 0.034

TABLE 4.8: **Fraction kept healthy for 2000 nodes and an average degree of approximately 3.** Results represent averages over 5 runs for each of the considered 5 instances of a dual Barabási-Albert model ($m_1 = 10, m_2 = 1$). Budgets are $k = 2$ and $k_c = 10$.

Agents (Test + Trace)	Multi-site mean-field	Agent-based
Random + Random	0.372 ± 0.035	0.371 ± 0.042
Acquaintance (w/ R) + Backward	0.633 ± 0.046	0.627 ± 0.053
Degree (w/ R) + Degree	0.841 ± 0.034	0.809 ± 0.028
RL + Backward	0.867 ± 0.029	0.851 ± 0.030
RL + Degree	0.889 ± 0.026	0.856 ± 0.025
RL + RL	0.911 ± 0.020	0.882 ± 0.018

4.5 Targeting vaccination campaigns

By coordinating vaccination with population testing, we enable our agents to explore the state space under realistic conditions. Besides ranking the nodes in the general population by varying k_v , we take inspiration from our tracing results and propose alternative formulations for each control policy that restrict the ranking pool to the contacts of newly-detected infected (using k_c , denoted with prefix ‘ct’ in Fig 4.11). As such, we set $k = 5$, and interchange between ($k_c = 0, k_v = B_v$) and ($k_c = B_v, k_v = 0$), where B_v is the vaccination budget per time increment (fixed at 5 days in this study). We consider the tester to be an RL actor operating over a dual Barabási-Albert network of 1000 nodes, and introduce a hard limit for the possible vaccinations to model stricter budgeting requirements. Fig 4.11 shows the boxplots of the containments obtained by our vaccination agents under two distinct budget schemes. Strikingly, the contacts-restricted policies appear to consistently outcompete their counterparts, demonstrating the benefits of a finer-grained control strategy. The RL actor continues to be the top performer in both settings, however, while its latent space quality is further confirmed by the competitiveness of the *clust* approach, which becomes a viable alternative to the degree- and learning-based solutions.

By targeting and isolating high-risk individuals, testing effectiveness reduces the benefits of a vaccination policy to a similar degree as it does for a targeted contact tracing strategy. As shown in Fig 4.12, the attack rate differences among three heuristic




FIGURE 4.11: **Infection control performance when prioritizing vaccination.** First 7 agents utilize k_c , while the rest use k_v . The total budget of vaccines is capped at 100.




FIGURE 4.12: **Attack rate when varying the vaccination fraction.** The attack rate is the percentage of nodes getting infected from a given population, representing the complement of the fraction kept healthy analyzed in the rest of this chapter. The differences between the heuristic policies become smaller when testing with an RL agent (second column), as opposed to conducting no testing (first column).

agents noticeably decrease when an RL actor tests $k = 5$ individuals daily, as opposed to no testing. The vaccination fraction has a negligible effect on this outcome, which is magnified when a single policy governs both processes, as the testing quality and the information gathered by it deteriorate over time due to the competing immunization.

4.6 Execution time analysis for test prioritization

Finally, we compare the mean execution time for running epidemics using each of our testing agents in Table 4.9. These results corresponds to the wall clock time recorded on an average Windows machine equipped with an Intel i7-7700 CPU, an NVIDIA RTX 3060 GPU and 32GB of RAM. As can be observed, the learning agents maintain a

TABLE 4.9: **Average wall clock time per epidemic during evaluation.** Configuration: dual Barabási-Albert networks of 2000 nodes, an average degree ≈ 3 , and a daily testing budget of $k = 2$.

Agents	Wall time (s)	Agents	Wall time (s)
Random	1.12	Acquaintance (w/ R)	1.12
Degree	3.23	Degree (w/ R)	3.19
Closeness (w/ R)	787.33	Betweenness (w/ R)	1176.32
Eigenvector (w/ R)	7.8	Pagerank (w/ R)	6.39
Neighborhood (w/ R)	1.49	RL/SL	15.92

low evaluation time of a few seconds, which is comparable to ad-hoc agents like eigenvector centrality and PageRank. In contrast, the policies backed by the closeness or the betweenness centrality have slower runtimes, being infeasible for big graphs.

4.7 Conclusion and future work

In this chapter, we show how policies for controlling an epidemic through targeted testing, tracing and vaccination in a resource-limited environment can be learned using expressive graph neural networks that can integrate both local and long range infection dynamics. Across many different scenarios, a policy inferred by a reinforcement learning agent outperforms a wide range of ad-hoc rules drawing from the connectivity properties of the underlying interaction graph, achieving containment rates of up to 15% higher than degree-based solutions with recollection, and more than 50% higher than random samplers. Interestingly, our agent also exhibits strong transferability, with one model trained on small preferential attachment networks being able to control the viral diffusion on several graphs of tens of thousands of vertices and diverse linkage patterns. While building on previous efforts (Meirom et al., 2021), we explore the role of contact tracing and vaccination, compare different ways of modelling the infection spread (multi-site mean-field versus agent-based), and scrutinize a varied set of heuristics. Analyzing more complex epidemic configurations and assessing the *EpiCURB* framework on further realistic scenarios represent natural extensions to this work. Moreover, investigating whether the SL policy can be improved by also feeding it information about the infection counts during training, or whether the RL actor benefits from knowing the infection labels represents an interesting venue for future exploration.

Building on the above, Appendix C outlines how the *EpiCURB* framework, coupled with appropriate visualization tools, can be used to inform decision-making in epidemiology. In essence, the aforementioned supplement illustrates how policy makers can approach the problem of finding the optimal level of public health interventions when these are deployed in packages, guiding the allocation of

resources for each measure according to the balance of budgets and epidemic outcomes. Additionally, it examines how the learning-based agents we introduce here can act as a platform for decision making in realistic settings where auditing is necessary, presenting several methods that can be used to evaluate their performance, such as mapping and grouping their latent space, or scrutinizing their decisions through applying local explainable models to each sampled node. By partitioning this embedding space, the authorities could further devise smarter public health policies that can effectively target the clusters of the population most at risk.

A logical next step for this research area would be to test the suggested methods in a more diverse range of epidemic control scenarios. Some of these should ideally aim to evaluate the agent's ability to make real-time decisions, which would require a careful examination of the suitable infrastructural solutions necessary to support this.

Chapter 5

Conclusions and Future Perspectives

5.1 Summary of principal contributions

Diffusion processes appear as a central theme in a broad range of different subjects, and this fact is not surprising, considering that many real-world phenomena, from the atomic to the human level, can be characterized by cascades of ‘spreading’ over time from a specific source to one or multiple targets. One crucial aspect in studying any form of diffusion is the placement of the target points relative to the source in the model of the space where it occurs. In physics and chemistry, many problems involving diffusion phenomena can be formulated and solved in a Hilbert space, which is endowed with a metric that determines how the Laplacian operator and the diffusion equation are defined and computed, ultimately allowing the calculation of the diffused quantity in each region. In contrast, analyzing diffusion processes in the sociology or public health domains can be challenging due to the lack of a standard metric space to measure their operation. In this context, interaction graphs represent a natural choice for depicting the relative location of individuals without relying on a specific metric. These can facilitate the development of effective applications that have the potential to significantly impact our society, such as producing precise predictions of future epidemic states, simulating varied counterfactual scenarios, or assessing the benefits of different public health intervention strategies.

A case study we extensively analyze throughout this thesis concerns the spread of the SARS-CoV-2 virus, responsible as of August 2023 for more than 6.9 million deaths. One of the major contributions we bring forward in this field is a novel individual-based mean-field model that represents contact tracing as an antagonist diffusion process to the viral spread. Building on the seminal work of [Farrahi et al. \(2014\)](#), our model allows us to capture the dynamics of infection and tracing in a realistic way, and to evaluate the impact of different outbreak control strategies. Unlike the interaction network over which the virus diffuses, contact tracing can only

access and intervene in the limited subset of it that is observable. This limitation arises primarily from the low adoption rates of tracing applications and the poor effectiveness of tracing interviews, which are widespread in our society. On the other hand, many people are known to ignore the self-isolation recommendations issued by authorities after testing positive or being in contact with an infected person. In spite of these challenges, we use simulations to show that contact tracing, especially when supported by digital solutions, can effectively lower the peak and the effective reproduction number R_e of a COVID-19 outbreak, notwithstanding suboptimal adoption rates and pervasive interview inefficiencies. In addition, these simulations corroborate the claim of [Mancastroppa et al. \(2021\)](#) by revealing that the epidemic outcome is frequently more profoundly impacted by improving the precision of the manual tracing network than that of digital methods, as the latter are unavoidably restricted by unreachable individuals. Our methodology, explained in detail in Chapter 3, can address from a modelling perspective several of the open problems about tracing raised by [Anglemyer et al. \(2020\)](#), such as the need for models to account for internet or phone access issues and the prevalence of tracing false positives, while also being highly adaptable and easily compatible with other kinds of diffusion processes that can be expressed through compartmental formulations.

The signals stemming from the interaction network over which the diffusion occurs feature learnable patterns that ought to be leveraged to grasp and manage its dynamics. Of particular interest to this thesis is the reinforcement learning technique of [Meirom et al. \(2021\)](#), which we extend in Chapter 4 to devise *EpiCURB*, a simulation-control framework for studying heuristic- and learning-based epidemic mitigation strategies. These come in the form of prioritized testing, tracing and vaccination programmes, having as goal the limitation of economic and societal costs (e.g. false positives) with the maximum diffusion suppression efficiency. Using the COVID-19 pandemic as an example, we reveal that various RL-backed policies outperform ad-hoc methods by up to 15% in containment rates, while far surpassing standard random strategies by more than 50%. This result is attributed to the superior timing of detection and candidate selection that the learned strategy exhibits, factors which we analyze in depth. Interestingly, control policies derived by RL agents are also transferable between network configurations, with models trained on graphs of one thousand nodes shown to perform similarly on larger networks containing tens of thousands of vertices. On the other hand, we demonstrate that incorporating recent negative test results into any targeted testing strategy can significantly enhance its performance. Finally, we show that restricting the pool of vaccination based on contact tracing further improves the containment rates of a targeted immunization campaign.

In conclusion, our current investigation aims to relay the significance of systematically examining diffusion processes, in their various forms, over interaction networks. In particular, a major objective of this thesis is to demonstrate that the outcomes of

spreading phenomena, such as viral epidemics, contact tracing, or information dissemination, are strongly influenced not only by the diffusion dynamics, but also by the underlying interaction network structure. By jointly exploiting these two factors, we show the one can obtain more realistic models of these phenomena, and devise efficient strategies for controlling their evolution. This marks the end of our discussion regarding the main findings and contributions of this document. That being said, we note that our work opens up new possibilities for future research that could explore more facets of the relationship between these factors, as well as address important open questions pertaining to future epidemics or other analogous diffusion processes. To support this further, the following sections offer a comprehensive overview of the broader impact and the possible future directions arising from our work.

5.2 Broader impact

5.2.1 Impact on pathogen epidemiology

In previous chapters, we have introduced several approaches that can be used to study, predict, and control viral epidemics. Most importantly, we have shown how epidemics can be simulated when a range of different control strategies needs to be scrutinized. Given the inherent limitations of contact tracing, taking the form of suboptimal application adoption rates and pervasive interview inefficiencies, it is desirable to obtain realistic approximations of the impact this intervention has on the spread. Using the recent pandemic as an example, we have demonstrated how our new compartmental formulation, SEIR-T, can be integrated with individual-based models to investigate the dynamics of SARS-CoV-2 under different testing and tracing strategies, while acknowledging the limitations mentioned above. Our proposed individual-based mean-field framework, which separates the traced status from the disease progression state, is directly applicable to any other epidemic, past and future, provided that it can be appropriately described via a compartmental model.

In addition to contact tracing, we have also explored other types of public health interventions that can help contain the spread of a pathogen: mask mandates, partial lockdowns, reduction of mixing or targeted ‘test, trace and vaccinate’ programmes. Adequately balancing these measures is paramount for containing a pathogen when the least amount of impact on the society is sought after. Using simulated COVID-19 outbreaks, we have outlined some of the options that authorities have at their disposal when these interventions are to be implemented cost-effectively, evaluating their effects using suitable visualization tools that depict the resulting containment rates. To address the issue of false positives that unavoidably arises when node-level interventions are deployed, we have advocated for the use of reinforcement learning agents, exploiting information from the interaction network structure through graph

neural networks to prioritize individuals for testing, tracing and vaccination. Additionally, we have revealed that using a learned control policy for these tasks outperforms heuristic-based methods, while demonstrating how policy makers could explore the former's latent space and its decision-making process to further optimize its outcomes. Analogously to the individual-based setup described above, the methods we have proposed for controlling and visualizing outbreaks are also general and independent of the type of pathogen studied.

5.2.2 Impact on information diffusion

Going beyond pathogen epidemics, the model presented in Chapter 3 can also be employed for studying information diffusion by considering testing and contact tracing as analogies for monitoring and feedback. Monitoring and feedback are the processes of tracking and evaluating how many people have received or accessed the information of interest, and how they have reacted or behaved after receiving or accessing it. These processes help measure the reach and impact of the information disseminated by different sources, and how they have influenced the opinions, reactions, or behaviors of the people who have come in contact with it. In this context, contact tracing represents the process of identifying and reaching out to the people who have been exposed to or influenced by the information from a given set of seed nodes. Based on this knowledge, one can target personalized incentives or messages to further propagate the desired content. Modelling contact tracing can therefore help one target the most relevant and receptive audience for his/her information diffusion strategy, increasing its effectiveness. As such, by simulating our 'test and trace' model with an appropriate compartmental formulation, such as SEIZ, we can estimate the expected number of people who would be exposed to or influenced by a given piece of information. Moreover, we can thoroughly scrutinize the impact of utilizing different levels of testing and contact tracing, while considering the pervasive effect of incompletely reconstructing the network of interactions. Finally, we can compare and contrast different diffusion strategies by analyzing the metrics captured by the monitoring and feedback processes, choosing the most suitable plan of action for any given scenario. By applying our model in this manner, one can design and optimize the information diffusion pertaining to any context that can be accurately represented via a compartmental formulation.

Influence maximization problems are a further domain of application that our work could address beyond its original scope. This could be realized through utilizing the epidemic control strategies we propose in Chapter 4, following some minor amendments. First of all, the information diffusion process must be describable through a graph-based model, as discussed in Section 1.4.2. Secondly, since the objective of influence maximization is to identify the optimal set of interventions that

can facilitate the diffusion process rather than hinder it, the formulation of the goal needs to change accordingly: $\max \sum_{t=t_0}^{\infty} \gamma_r^{t-t_0} \mathcal{I}(t)$, where $\mathcal{I}(t)$ is the total number of influenced individuals by time t , and γ_r is a discount factor. Thirdly, if the underlying diffusion formulation does not include an exposed/indecisive node state, the model needs to account for the possibility that the intervention fails (i.e. adoption occurs with a certain probability). Lastly, a separate mechanism for testing (i.e. observing whether the information was adopted) has to be established so that the agents can (partially) perceive the current state. This can be either stochastic or governed by a testing prioritization policy. [Meirom et al. \(2021\)](#) provide an excellent example of applying such a setup for solving influence maximization. In particular, they employ the IC and LT models to simulate the social influence process, a random tester to explore the state space, while allocating interventions, with a fixed probability of success, according to various policies. In this context, they demonstrate that RL agents still outperform standard heuristic strategies, regardless of the network configuration utilized. Their setup could easily be improved, however, by using a targeted testing process, thus ensuring that vital information about the diffusive state is captured early, and by analyzing alternative spreading formulations, such as SEIZ, which could offer a more realistic picture of the information dissemination process in online social networks.

5.3 Limitations and future perspectives

5.3.1 Future work needed to address current limitations

Apart from our main findings, this thesis also identifies some important limitations and challenges of applying epidemiological models over graphs for examining and optimizing diffusion processes, such as data availability and quality, parameter estimation and uncertainty, model calibration and validation, among others.

Therefore, future research is needed to tackle these issues and enhance the applicability and reliability of the approaches presented here.

First and foremost, our findings and conclusions are based on the premise that simulation data can be extrapolated to the real world. Nonetheless, despite our best efforts to simulate realistic and diverse scenarios, it is possible that our environments, and by extension the results, do not always match the reality. The network and parameter configurations studied here, for instance, are by no means exhaustive. As such, future research is required to establish the connectivity patterns and parameter ranges that more faithfully describe the transmission chains within each region of interest. Collecting more contact tracing data to reconstruct the true social interaction pathways would constitute a solid initial step, followed by fitting the epidemiological parameters to the hospitalization and death counts recorded in the target area in order to provide more reliable estimates of the impact of various public health interventions,

as well as a more representative training setup for the learning-based agents. For enabling the full potential of the latter, future work should also consider a stronger decoupling of the control agent's logic from the simulation engine, which is currently limited in our own implementation. Furthermore, for real-time decision making, the agent's infrastructure should ideally be scalable and resilient, while being able to rapidly ingest new information about test results and connectivity shifts to adjust its policy, by leveraging suitable technologies such as stream processing frameworks, NoSQL databases, and cloud computing services.

Aside from the epidemiological parameter values curated from the literature, we have had to make several assumptions about the appropriate ranges for previously undefined or loosely-defined parameters. A prominent example is represented by the τ_r and τ_t rates, which were first introduced in [Huerta and Tsimring \(2002\)](#), later adopted by [Farrahi et al. \(2014\)](#), but seldom fully reconciled with real epidemic data. As we highlight in Section 3.4, the effectiveness of a "test and trace" strategy within our IBMF simulation framework heavily depends on the difference in value between β and both τ_r and τ_t , with the latter having the potential to mislead if they are set to excessively high values, creating the false impression that these interventions can single-handedly contain any outbreak. Consequently, we have chosen to present results for a wide range of bounded values in Chapter 3, targeted for each individual scenario, aiming to provide a comprehensive and unbiased understanding of their effects. That being said, our analysis could be improved by calibrating these rates according to measurements of contact frequency or by fitting them to observed epidemic curves, alongside other unknown diffusion parameters. Future endeavors could also explore further configurations for the relative probability of getting tested positive, r^T , and the non-compliance rate η , parameters that we introduce, but for which we only consider a single assignment for simplicity.

Our framework also has the limitation of employing a stochastic process for creating the visible tracing subviews. In contrast, the patterns of contact recollection, tracing effectiveness and application uptakes in the real world are not completely random, but influenced by many factors, such as socio-economical ones, pathogen prevalence, perceived risk, data collection awareness, privacy concerns, and others. As such, our findings regarding the network overlap and adoption levels in Chapter 3 should not be interpreted in terms of absolutes, but rather they should serve as indicators of their relative impact on the spread under the assumption that the aforementioned factors become negligible upon averaging. Future endeavors could test whether the latter is true by verifying whether epidemic outcomes significantly shift under more accurate subview derivation regimes (e.g. derived from contact tracing data or calibrated according to mobile tracking information). Considering the high complexity of the control problem under scrutiny, we make a further simplification in Chapter 4 by assuming a complete recovery of the transmission network, which is largely consistent

with previous work. It would be desirable, however, to extend the evaluation of the control agents beyond this premise, by introducing a dual/triad topology configuration, where the agents can only observe a partial and noisy infection graph. This would make the problem more realistic and challenging, while testing the robustness of the proposed agents.

Focusing on our key modelling choices, we additionally highlight some areas of potential improvement where systematic ablation studies could further reveal the impact of different modules and parameters on the epidemic outcome. To begin with, we suggest that a comparison between exponentials and more general Gamma distributions for the event waiting times within our IBMF models would be very valuable, as the assumption of exponentials, although justified by the literature in our setup, may be too optimistic for certain pathogen outbreaks. Going forward, we remark that our conclusions in Chapter 4 on the efficiency of SL and RL agents in controlling outbreaks are based on the presupposition that using the same architecture and diffusion setup ensures a fair benchmarking. However, one could also argue that a fairer comparison would require the same input signals for both agents. This could be explored in future ablation studies by augmenting the SL agent objective to concomitantly optimize against the total infection counts or the RL actor loss with a cross-entropy term on the set of infection labels. Finally, we acknowledge our neural network architectural choices have been based on literature recommendations and preliminary experimentation with general node classification and outbreak control tasks, rather than a more rigorous ablation analysis. As such, future research would be needed to assess the performance of other popular GNN configurations in the *EpiCURB* framework, including but not limited to GRAND-generated models (see Appendix A.2) and Graph Transformers (Dwivedi and Bresson, 2021).

As a final point, we would like to emphasize the potential of the policy visualization techniques we propose in Appendix C to enhance agent-guided outbreak control, while also recognizing the limited scope of our current experimentation approach. The latter significantly relies on several assumptions, such as a constant value for infectiousness reduction in mask wearers, or the ability of authorities to deploy interventions in certain subpopulations or to achieve a fixed contact reduction percentage (relative to a past reference point). For this reason, we believe subsequent work should explore these simulation-based visualization methods across more diverse and realistic epidemic configurations, analyzing the impact of the parameters associated with each public health measure (e.g. by quantifying how the infectiousness varies due to mask wearing or establishing reliable criteria for deterministically dividing the population into percentiles), performing cost-benefit assessments to root the scaling of interventional efforts into actionable decisions, expanding our clustering analysis to determine the features that contribute to each grouping, and performing a

temporal analysis of the feature importances obtained from GraphLIME in order to find hidden patterns that can better describe an agent’s decision-making process.

5.3.2 Future directions inspired by preliminary findings and related works

As outlined in Section 5.2, the epidemiological models studied in this thesis are suitable for analyzing the dynamics of diffusion processes across different domains. While the preceding chapters have exemplified the merits and utility of our proposed modelling and control methods using the COVID-19 pandemic as a case study, we remark that these approaches are flexible enough to accommodate other contexts that involve comparable dissemination phenomena over graphs. For instance, they can be applied to other infectious diseases, past and future, as well as to information propagation, viral marketing, cybersecurity, or computer network analysis.

GNNs and graph kernels have imposed themselves as two of the most versatile and efficient tools for exploiting the structure of networks in predictive problems. In Chapter 2, and continuing in Appendix A, we discuss both methods at length, highlighting the theoretical concepts behind the most popular variants of these algorithmic frameworks and their current limitations. GNNs are widely regarded as superior to kernel-based methods in graph learning tasks, but there are few studies that reliably compare both approaches across a wide range of settings to be confident of such a conclusion. Based on our preliminary comparison in Table E.1, GATs and GCNs appear to slightly outperform the other baselines in vertex classification. However, SVMs with the simple diffusion kernel and a linear kernel on the features perform unexpectedly well on the Cora dataset (Yang et al., 2016), raising the question of whether more complex Gram matrices can compete with GNNs. The context of diffusion processes, in particular, remains relatively unexplored by kernelized methods, possibly due to the challenge of capturing temporal signals from the data. It has been claimed, however, that graph transformations based on temporal random walks could address this issue (refer to Appendix A.3.2). As such, given multiple connectivity patterns, based on both artificial and real interaction networks (e.g. Social Evolution, Haggel or Infectious datasets – see Appendix E.2), and simulations of stochastic compartmental models over these connections, we propose that future work could evaluate whether the most promising graph kernels to date, such as the WL, WL-OA or propagation kernels, can be combined with the aforementioned transformations to outperform GNN-based solutions, similar to our SL agent in Chapter 4, in answering epidemiological questions of the form: ‘How many people will get infected in total?’; ‘Can the prediction be improved as the process unfolds?’; ‘What types of signals do the models need, and how often?’; ‘If a model offers adequate predictions, can we verify whether it faithfully captures the underlying

outbreak dynamics by refitting a compartmental model to its outputs, followed by quantifying the estimation errors?'.

The response to future epidemics could further benefit from a deeper investigation into aggregated mobility datasets that capture the inter-regional flux of people (see examples in Appendix E.2). In earlier studies (Kapoor et al., 2020; Wang et al., 2020b; Panagopoulos et al., 2021), these datasets have been used to construct spatio-temporal graph structures that can facilitate disease prediction, as outlined in Appendix A.3.1. Unlike prior work, however, we suggest a more general approach to this setting, framing it as a semi-supervised learning task where communities with no reliable infection forecasting depend on GNN models to aggregate such information from the “neighboring” regions based on their associated mobility. Although GCNs and their variants have been shown to be particularly effective for semi-supervised problems (Kipf and Welling, 2017; Zhou et al., 2020) or partially observable environments (see Chapter 4), most epidemiological studies involving GNNs and mobility datasets assume the complete historical infection data is available to predict future outcomes. This scenario may be too optimistic with respect to the reliability of each region’s estimates. Moreover, transferring such models to countries where only some regions report daily epidemic data could be challenging, even if those areas are scarce and possess rich mobility information. As such, a partially-labelled setting would be more suitable and realistic in this context, allowing for zones with no reliable estimate to be masked and predicted for. Furthermore, once a graph-based learning algorithm is trained, we posit that adjusting the parameters of an appropriate epidemiological meta-population model to its predictions could prove fruitful, leveraging the inferred dynamics to forecast future epidemic trends in an easily reproducible manner.

Another domain that offers a notable example of further applications for the methods we introduce in this thesis is that of information diffusion. Ranging from label propagation in graphs featuring non-human actors to news dissemination over online social networks, the applications of information diffusion models are both numerous and of great importance. In particular, predicting and partially controlling the information flow over social media can often make the difference between a successful or a failing campaign, sometimes even in the absence of profile-specific data (Bagrow et al., 2019). Be it for enriching political discourses (Anderson, 2017), maximizing the reach of viral marketing (Even-Dar and Shapira, 2007), or stopping malicious actors from distributing harmful content, modelling the nodes and their interactions in a social network has many far-reaching implications. In this thesis, we have conducted a preliminary survey of various information diffusion models, and have observed that the epidemiological approaches generally exhibit a remarkable efficacy in this context, especially for discriminating between true and false news or for investigating social influence. However, we have also noted that these models tend to be confined to a complete graph setting in the literature. Taking all the above into account, we consider

that a fruitful future research direction in this domain would be to leverage data from Twitter APIs on various types of news, both authentic and fabricated, related to highly contentious topics (e.g. COVID-19 vaccines) to create graph representations, and analyze their properties using our network-based simulation-control framework. In this regard, the Twitter Tennis graph datasets (see Appendix E.2) can serve as benchmarks for comparing the performance of different learning methodologies. Likewise, the IBMF approach outlined in Chapter 3 could also offer valuable insights into the problem of information diffusion by drawing analogies between the testing and contact tracing processes in viral outbreaks and the monitoring and feedback mechanisms in the latter, thus representing another potential direction for subsequent work. Finally, we envisage that the reinforcement learning agents we rigorously examine in Chapter 4, together with the complementary visualization techniques presented in Appendix C, could enable the generation and evaluation of effective influence maximization campaigns in complex networks, leading to more optimal solutions than conventional hill-walking techniques were able to offer in the past.

Appendix A

Additional Review Material on Graph-based Machine Learning

A.1 Graph-level kernels

One of the most common approaches for constructing graph-level kernels using information from lower-level kernels is the R-Convolution method (Haussler, 1999). Let \in_r be the binary relation ‘is subcomponent of’. Then the standard R-Convolution kernel on the graphs X and Y can be written as Eq A.1. One also has the option to swap the sums in Eq A.1 to other aggregating operators, giving rise to different mapping functions ϕ (e.g. *max* in optimal assignment kernels (Fröhlich et al., 2005)).

$$k_G(X, Y) = \sum_{x \in_r X} \sum_{y \in_r Y} k(x, y) \quad (\text{A.1})$$

Not all graph-level kernels use the R-Convolution model, however. These methods are often unable to leverage the kernel trick, relying on explicitly deriving the ϕ mapping instead. We note that this does not necessarily lead to worse performances, as in fact, many such approaches are computationally faster than implicit schemes (Kriege et al., 2019). In particular, instead of using Kronecker δ base kernels, it is usually more advantageous to bin the node labels (or the discrete hashes of the feature vectors), and produce graph embeddings from counting the bin sizes (Neumann et al., 2016). Important mentions here are the Graphlet kernel (Shervashidze et al., 2009), which creates network representations $\phi(\cdot)$ that contain the number of occurrences of common subgraph patterns – known as graphlets (Przulj et al., 2004), the WL kernel (Shervashidze et al., 2011), which stacks in each graph’s $\phi(\cdot)$ the histogram of labels resulted from running i iterations of k -WL updates, the propagation kernel (Neumann et al., 2016), which replaces 1-WL with a custom propagation scheme like random walks, or the WL-OA (WL optimal assignment) kernel of Kriege et al. (2016). The latter

has been shown to perform particularly well across a vast range of graph classification tasks, albeit being slightly slower to compute than some of the alternatives (Kriege et al., 2020). By letting $\mathcal{B}(X, Y)$ be the set of all bijections between the nodes of some networks X and Y , the optimal assignment graph kernel can be written as:

$$k_{G-OA}(X, Y) = \max_{B \in \mathcal{B}(X, Y)} \sum_{(x, y) \in B} k(x, y) \quad (\text{A.2})$$

The above formulation is not always p.s.d. unless the base kernel k is *strong* (Kriege et al., 2016). WL-OA uses a strong node-level comparator based on 1-WL updates, given by the equation $k(x, y) = \sum_{i=0}^h \delta(l^i(x), l^i(y))$, where $l^i(x)$ is the 1-WL assignment of node x at iteration i . Computing k_{G-OA} with this Kronecker delta base kernel has been shown to be equivalent to the histogram intersection kernel in Eq A.3, where Σ is the image of the node label function (i.e. the set of all labels), while X_h and Y_h are the lists of nodes in graphs X and Y , respectively, that have been assigned the label h in one of the 1-WL iterations. With this equivalence in mind, k_{G-WLOA} can be linearly calculated from the ϕ representations of the WL kernel defined above, because these directly encode the histograms of WL labels stemmed from every iteration.

$$k_{G-WLOA}(X, Y) = \sum_{h \in \Sigma} \min\{|X_h|, |Y_h|\} \quad (\text{A.3})$$

Adding further complexity to any of these representations is also possible by applying a hyperbolic kernel on top (e.g. RBF). In this case, the Euclidean distance in the original RBF formulation gets replaced by a distance metric defined by the lower-level kernel (see Eq A.4), giving rise to the formulation in Eq A.5 (Kriege et al., 2020).

$$d_k(x_i, x_j) = \|\phi(x_i) - \phi(x_j)\|_{\mathcal{H}} = \sqrt{k(x_i, x_i) + k(x_j, x_j) - 2k(x_i, x_j)} \quad (\text{A.4})$$

$$k_{G-RBF}(X, Y) = \exp(-d_{k_G}(X, Y)^2 / 2\sigma^2) \quad (\text{A.5})$$

A.2 Alternative approaches for building GNNs

Here we discuss the GRAND approach for constructing GNNs with a continuous depth that limits the impact of oversmoothing and bottlenecks. In the spirit of Neural ODEs (Chen et al., 2018), GRAND utilizes the time parameter in the heat diffusion equation (refer to Eq 2.3) as a continuous analogy to the GNN layers. Given a network G with N vertices, the node features $X(t)$ at depth t , and \mathbf{A} an $N \times N$ attention matrix mirroring the structure of the adjacency of G (i.e. $\mathbf{A}_{ij} = 0$ for $(i, j) \notin E(G)$), the

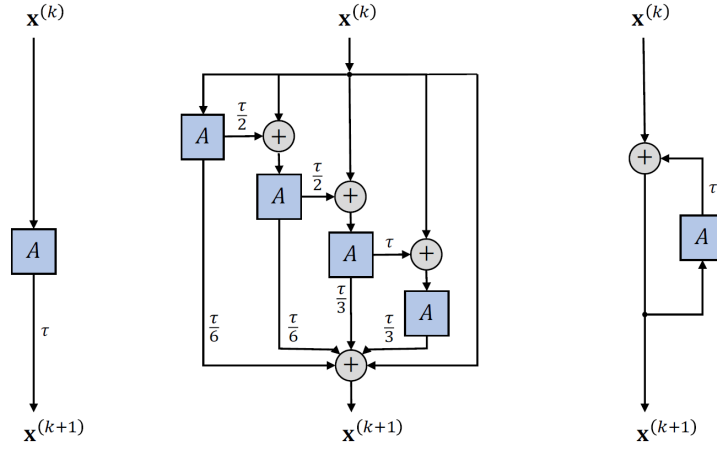


FIGURE A.1: **Illustration of different discretization schemes for Eq A.6.** From left to right, the schemes are explicit Euler, 4th order Runge-Kutta, and implicit Euler. Source: Chamberlain et al. (2021).

diffusion partial differential equation w.r.t. the GNN depth becomes:

$$\frac{\partial}{\partial t} X(t) = (\mathbf{A}(X(t)) - I)X(t) = \bar{\mathbf{A}}(X(t))X(t) \quad (\text{A.6})$$

$$\mathbf{A}(X(t))_{i,j} = \frac{1}{h} \sum_h \text{softmax} \left(\frac{(W_K x_i(t))^T W_Q x_j(t)}{d^k} \right) \quad (\text{A.7})$$

The attention matrix in the general case of the GRAND-nl formulation can be obtained as an expectation over multiple attention heads, as in Eq A.7, where W_k and W_Q are learned matrices that are time-independent (i.e. shared among layers), d_k is a hyperparameter, while $x_i(t)$ and $x_j(t)$ are node features at depth t . The attention mechanism utilized here is the scaled dot product attention (Vaswani et al., 2017), which the authors claim outperforms the Bahdanau attention variant of GAT. As a special case of GRAND-nl, GRAND-l utilizes a constant $\mathbf{A}(X(t)) = \mathbf{A}$, making the objective in Eq A.6 solvable with eigendecomposition: $X(t) = e^{\bar{\mathbf{A}}t} X(0)$.

The final GRAND model is summarized by Eq A.8. In this formulation, $\text{Linear}_\phi(X)$ is a linear encoder layer of the initial node features X which outputs the first time-point value $X(0)$. Starting from $X(0)$, the integral term gets solved iteratively by temporal discretization using any established ODE solver, with $\frac{\partial X(t)}{\partial t}$ approximated at each timestamp t using the GNN layer outlined by Eq A.6 (also check Fig A.1). Finally, a prediction \hat{Y} is obtained by feeding the sum between $X(0)$ and this approximated integral to $\text{Linear}_\psi(\cdot)$, a linear decoder layer. The total depth T of the architecture need not be specified a priori when using adaptive time step ODE solvers, offering GRAND a significant advantage compared to more common techniques for building GNNs.

$$\hat{Y} = \text{Linear}_\psi \left(X(0) + \int_0^T \frac{\partial X(t)}{\partial t} dt \right) \quad X(0) = \text{Linear}_\phi(X) \quad (\text{A.8})$$

The authors experiment with multiple types of time discretization in GRAND by considering different ODE solvers and varying the step sizes: explicit and implicit Euler, explicit Adams-Bashforth, implicit Adams-Moulton (Hairer et al., 1993; Butcher, 2000), and adaptive Runge-Kutta 4/5 (Press and Teukolsky, 1992). Their analysis suggests that implicit schemes with large step sizes ($\tau > 0.25$) remain stable and achieve high test accuracies considerably faster than adaptive explicit algorithms, while the latter schemes with fixed step sizes greater than 0.005 diverge away from appropriate solutions.

Given that W_K and W_Q are shared among all the layers, GRAND uses less parameters than GCN, GAT, or other similar models. At the same time, performance-wise, GRAND achieves competitive results when compared with the aforementioned architectures in a series of node classification benchmarks, being the highest performer across the majority of the datasets considered by Chamberlain et al. (2021). As the authors concede, however, this design choice may constrain the model's capacity to attain its maximum potential, while foregoing it would significantly increase the computational demands.

A.3 More graph-based ML applications to diffusion processes

A.3.1 Modelling the region-to-region spread of COVID-19

Since the COVID-19 pandemic started, many companies have open-sourced data concerning mobility trends of mobile phone users in order to help the research community and various government officials better understand and plan against its spread. Among the main sources of such data (for more information please refer to Ilin et al. (2021)), we note Google¹, Facebook², SafeGraph³, and Cuebiq⁴. This type of data is typically anonymized and aggregated on geographical regions, showing either differences in mobility through various public places or places of interest (POI), sometimes relative to a value considered normal, or daily/weekly flows of population from a source place to a target region. The exact locations and interactions of individuals that are part of such datasets are usually masked, an overview of the mobility across an entire neighborhood/area being provided instead. This aspect gives rise to generalized epidemic networks, which no longer have individuals as nodes and interactions as edges, but rather entire communities as vertices, and the mobility flows between them as links. For such networks, simple graph-based

¹Google COVID-19 Community Mobility Reports: <https://www.google.com/covid19/mobility/>

²Facebook Data for Good: <https://dataforgood.fb.com/docs/covid19/>

³SafeGraph COVID-19 Data: <https://www.safegraph.com/covid-19-data-consortium>

⁴Cuebiq Data for Good: <https://www.cuebiq.com/about/data-for-good>




FIGURE A.2: **Spatio-temporal input graph.** White node from Day 10 gets an aggregated representation from its neighbors at the current timestamp, as well as its images from previous days (i.e. red-filled vertices). Source: Kapoor et al. (2020).

compartmental models are no longer suitable, and meta-population ODE-based formulations are often used instead (Rahmadani and Lee, 2020; Chang et al., 2021).

Recently, many works have turned to GNNs to leverage this type of data for improving the forecasts related to the spread of SARS-CoV-2 (Kapoor et al., 2020; Wang et al., 2020b; Panagopoulos et al., 2021). In this setting, the prediction problem represents a regression task that takes as input a time series of graphs G_{t-k}, \dots, G_t describing the extent of the epidemic in each particular location, and the varying mobility between these regions, and outputs forecast values for indexes $\{t+1, \dots\}$. To model the spatio-temporal nature of this sequence, Kapoor et al. (2020) use a GCN that operates on networks with temporal edge connections to previous timestamps, while Panagopoulos et al. (2021) propose three approaches: GNNs with shared weights for each time index, GNNs combined with LSTMs (Hochreiter and Schmidhuber, 1997), and GNNs with shared weights and transfer learning between regions.

The skip-connection model proposed by Kapoor et al. (2020) is summarized in Eq A.9, where $H^{(0)}$ is obtained by feeding the concatenation of all temporal node features X_t to an MLP, $H^{(l+1)}$ is the GCN layer's output with $H^{(0)}$ concatenated to it, while P is the final prediction. This model is trained with a mean squared logarithmic error (MSLE), which receives as input a graph that manifests as 100 stacked layers of mobility networks recorded at each time index (i.e. 100 days), with each node at time t having links to its image at the previous 7 time steps (see visualization in Fig A.2).

$$\begin{aligned}
 H^{(0)} &= \text{MLP}(X_t || X_{t-1} || \dots || X_{t-d}) \\
 H^{(l+1)} &= \sigma(D^{-\frac{1}{2}} A D^{-\frac{1}{2}} H^{(l)} W_l) || H^{(0)} \\
 P &= \text{MLP}(H^{(L)})
 \end{aligned} \tag{A.9}$$

A.3.2 Temporal graph kernels for diffusion processes

In this section, we review the methods proposed by Oettershagen et al. (2021) for transforming dynamic networks into alternative representations that enable the use of graph kernels to model diffusion processes. The authors employ these techniques to perform two classification tasks based on multiple time series of network states of the form G_0, \dots, G_T : (1) predict whether the labels of an input sequence got assigned via node-to-node spreading or not; (2) predict the most likely diffusion process, from a predefined set, to have generated the labels in that time series. The first of these approaches entails reducing the temporal graph to a static representation as follows: Reduce all temporal links between nodes u and v to a single edge labelled by the time index of its first appearance. Next, create a second labelling map l' which assigns to v the time index of its first label update, or a value 0 if the latter stays constant over time. The time indexes need to be global for the whole operation. Finally, the compression step is followed by the kernel computation. The representations in this case remain small throughout the process, $O(|V(G)|^2)$, but they are generally lossy.

Many popular static graph kernels, including random walk, shortest path and WL, can be formulated in terms of walks (or traversals) along the vertices and edges of the input networks. The main idea behind the other two methods is to replace these static traversals with *temporal walks*. The latter produce several label sequences of the form $L_s(w_v) = (l(v_1, t_1), l(e_{(v_1, v_2)}, t_1), l(v_2, t_1 + 1), l(v_2, t_2), \dots, l(v_{l-1}, t_{l-1}), l(e_{(v_{l-1}, v_l)}, t_{l-1}), l(v_l, t_l))$, where l is a labelling function for both nodes and links along w_v . This formulation accommodates some waiting time at each vertex, with labels for a given node appearing both at $t_{i-1} + 1$ and t_i . Note the temporal walk w_v must be valid for the corresponding $L_s(w_v)$ to exist (i.e. w_v can follow only hops that are available at each individual timestamp). Comparing counts of walks indexed by $L_s(w_v)$ effectively renders a temporal random-walk kernel. The authors show that stochastically sampling w_v , followed by computing the corresponding Gram matrices and training SVMs with them, performs well empirically without being too resource-intensive.

The second representation method, Directed Line graph expansion (DL), builds on the ideas above by constructing a new graph formed of pairs of vertices (n_{vu}^t, n_{uv}^t) for every edge (v, u) present in G_t , with links dictated by possible temporal walks between them, i.e. (n_{vu}^t, n_{uz}^{t+i}) for any $i > 0$. A special labelling function gets created $l'(n_{vu}^t) = (l(v, t), l(u, t + 1))$, respecting the temporal walk conditions mentioned before. The authors prove that any static kernel operating on DL-transformed networks is equivalent to using a temporal walk kernel on their original structures. The main limitation of this method is that the size of the resulting graph grows quadratically in the number of edges: $O(|E(G)|^2)$.

The final approach creates a representation formed of time-vertices (v, t) , where the labels are preserved and the connections mirror the flow of time (i.e. node (v, t)

connected to (u, t') only if $t < t'$). Any edge (v, u) in the original graph generates at most four time-vertices connected by exactly two edges: $((v, t), (u, t + 1))$ and $((u, t), (v, t + 1))$. If waiting times at nodes do exist, further links are added to the new graph structure to make the temporal walk complete. The method, called Static Expansion (SE), requires appreciably less memory than DL, $O(|E(G)|)$, without being as lossy as the first approach.

The experimental setup used by [Oettershagen et al. \(2021\)](#) involves numerous network datasets onto which the SI compartmental model is simulated. In problem (1), several SVMs with WL and random walk kernels are applied to the proposed alternative representations of the graph sequences, learning to distinguish between valid simulations and random system states (i.e. stochastic infections without the point-to-point condition). In contrast, exercise (2) predictors are trained to classify instances generated by separate SI diffusion processes, featuring different transmission rates. The authors show that DL and SE achieve competitive accuracies in both tasks across the whole range of considered datasets, outperforming both static kernels and general temporal walk kernels.

That being said, the biggest limitation of the work of [Oettershagen et al. \(2021\)](#) is the restricted scope of the experiments conducted with these graph transformation methods. In the context of epidemics, it is rarely of practical use to perform binary classification of diffusion dynamics. Nonetheless, the ability to accurately perform the aforementioned task may indicate that complex signals from the temporal data are being learned. An interesting venue for further exploration would be to utilize state-of-the-art kernels, such as the WL, WL-OA or propagation kernels presented in Appendix A.1, in conjunction with either the DL or the SE transformation, to solve more important problems in diffusion, such as infection forecasting.

Appendix B

Modelling Contact Tracing: Supporting Information

B.1 Open-source model and data

The open-source implementation of the simulator can be consulted at:

<https://github.com/andrei-rusu/contact-tracing-model>.

The statistics our simulations captured can be analyzed in full by following:

<https://doi.org/10.6084/m9.figshare.14101946>.

B.2 Simulation statistics

Aside from the metrics analyzed in the main chapters, our individual-based model can readily be used to analyze various other statistics about the simulated epidemic: the total number and peak of hospitalization and death, the total of people that have isolated or have recovered, tracing false positives (Fig B.1) and false negatives (Fig B.2), the number of non-compliant nodes, the tracing *efforts* (Farrahi et al., 2014), the incidence and growth rates registered over a variable window size etc.

Both Fig B.1 and Fig B.2 give an alternative view of the repercussions a country can face if contact tracing is too zealous or too slow. If excessively many people get incorrectly isolated (false positives), the resulting socio-economic burden may significantly disrupt a community. If, on the other hand, tracing is very inefficient, the infectious population (false negatives) will spread the disease uncontrollably, ultimately leading to many hospitalizations and deaths. The right balance needs to be struck between these two for a ‘test and trace’ strategy to be deemed successful.

B.3 Further analysis of noteworthy trends

As previously discussed, the values ascribed to the overlap Γ and the uptake r dictate whether a contact tracing rate τ_t is effective. The trends imposed by these quantities on τ_t can be further scrutinized in Fig B.3 and Fig B.4. At the 0.5 level, both show the ability to noticeably influence the infection curves obtained by $\tau_t \geq 0.04$. The peak differences reached by the same tracing rates become substantial at the extreme points.

When studying the combined effects of manual tracing at different Γ and digital tracing at various r on the effective reproduction number R_e , it is often desirable to visualize the corresponding three-dimensional trends. To that end, we plot in Fig B.5 the 3D surface of the aforesaid variables belonging to the experiment in Section 3.4.5.

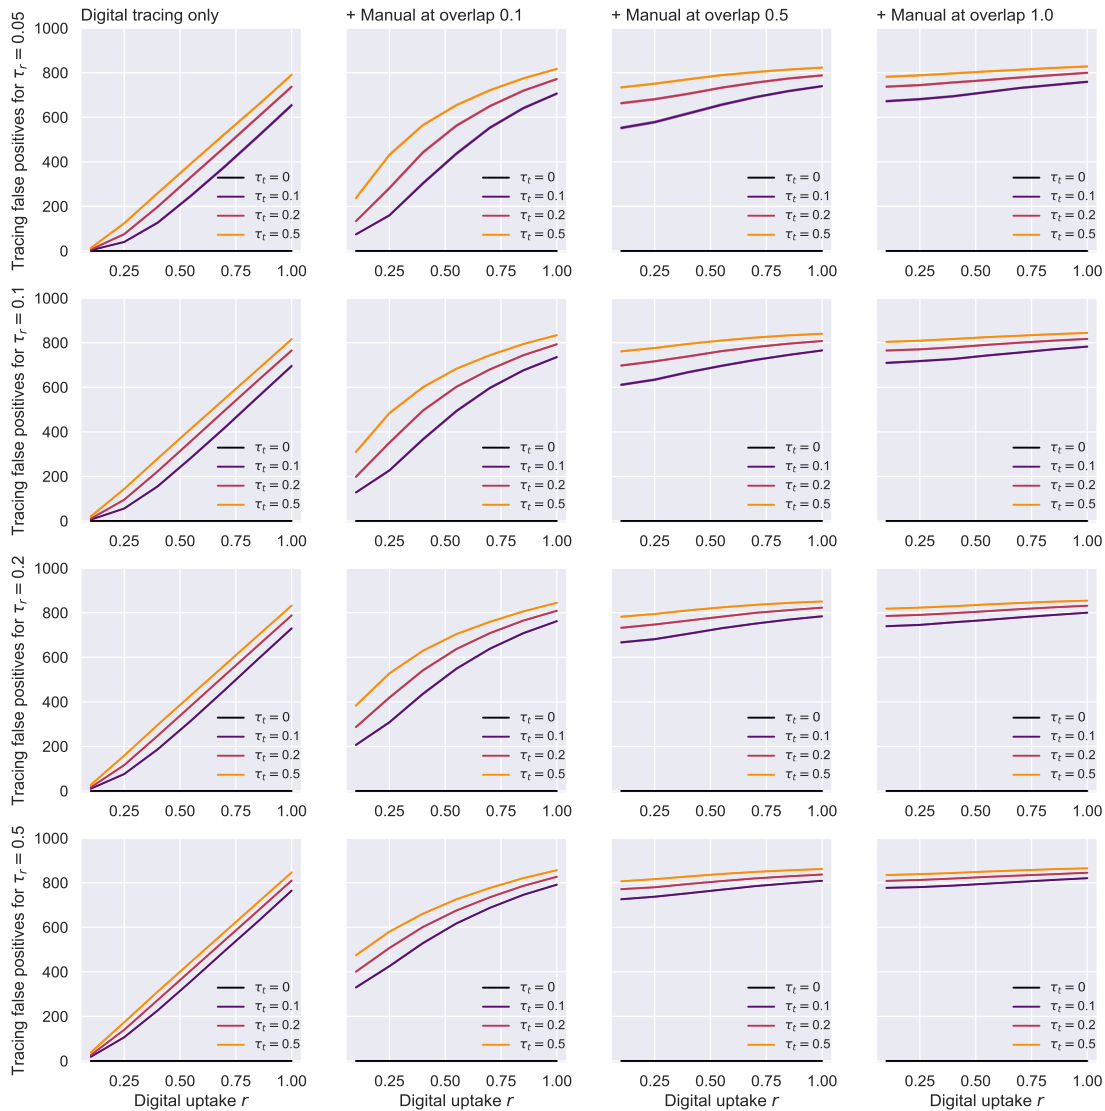


FIGURE B.1: **Tracing false positives.** The amount of susceptibles being incorrectly traced and isolated. Results here correspond to the last experiment, with $N = 1000$, Holme-Kim graph topology and 10% initial infected.

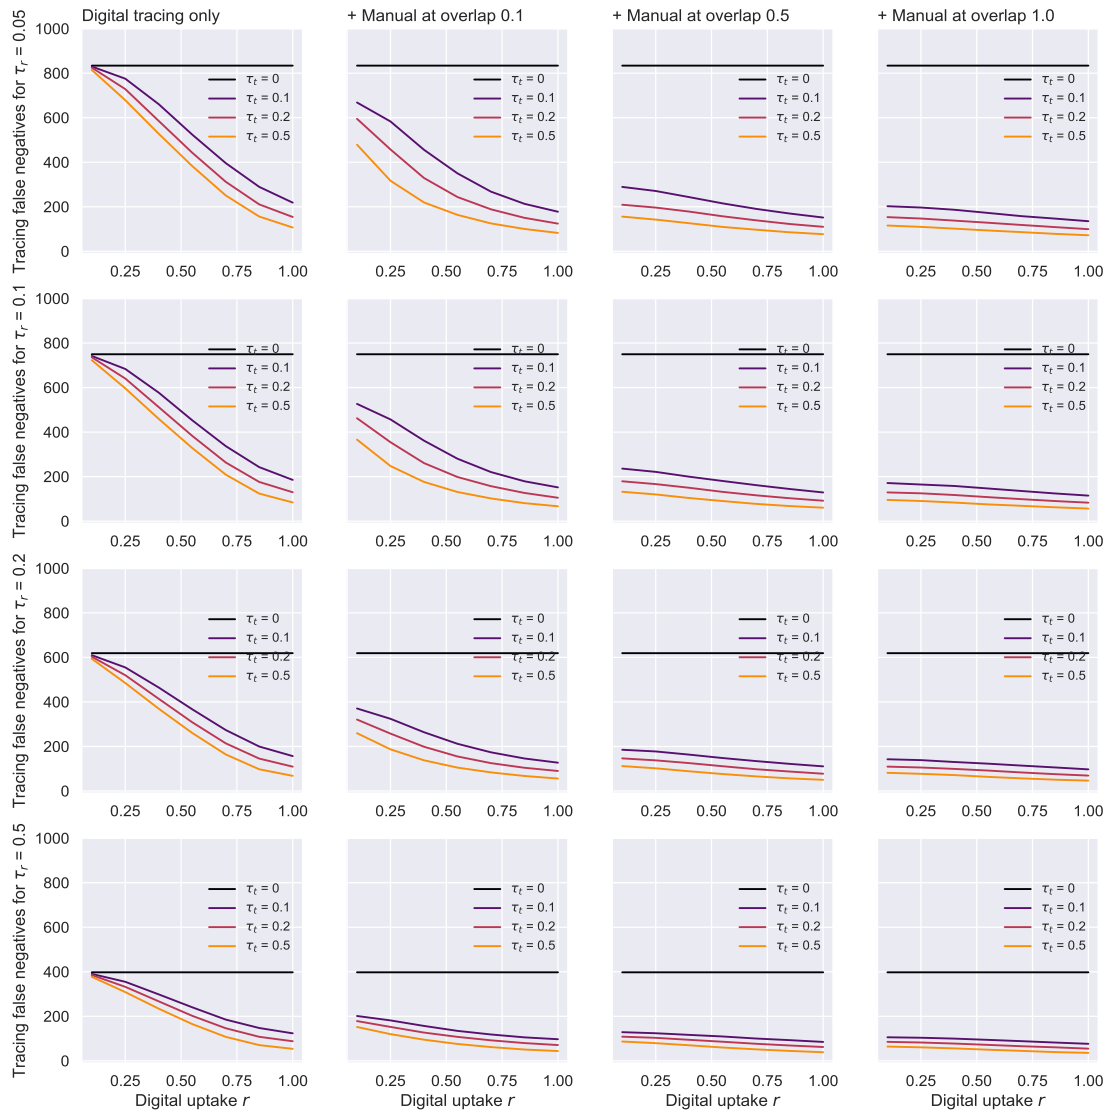


FIGURE B.2: **Tracing false negatives.** Number of *infectious* people not traced. Results correspond to a Holme-Kim graph topology with $N = 1000$ and 10% initial infected.




FIGURE B.3: **The overlap influencing the efficacy of contact tracing rates.** Results for $N = 10000$, random graph topology with mean degree $K = 10$. τ_r fixed at 0.04.




FIGURE B.4: Uptake influencing the efficacy of contact tracing rates. Results for $N = 1000$, random graph topology with mean degree $K = 10$. τ_r fixed at 0.05.

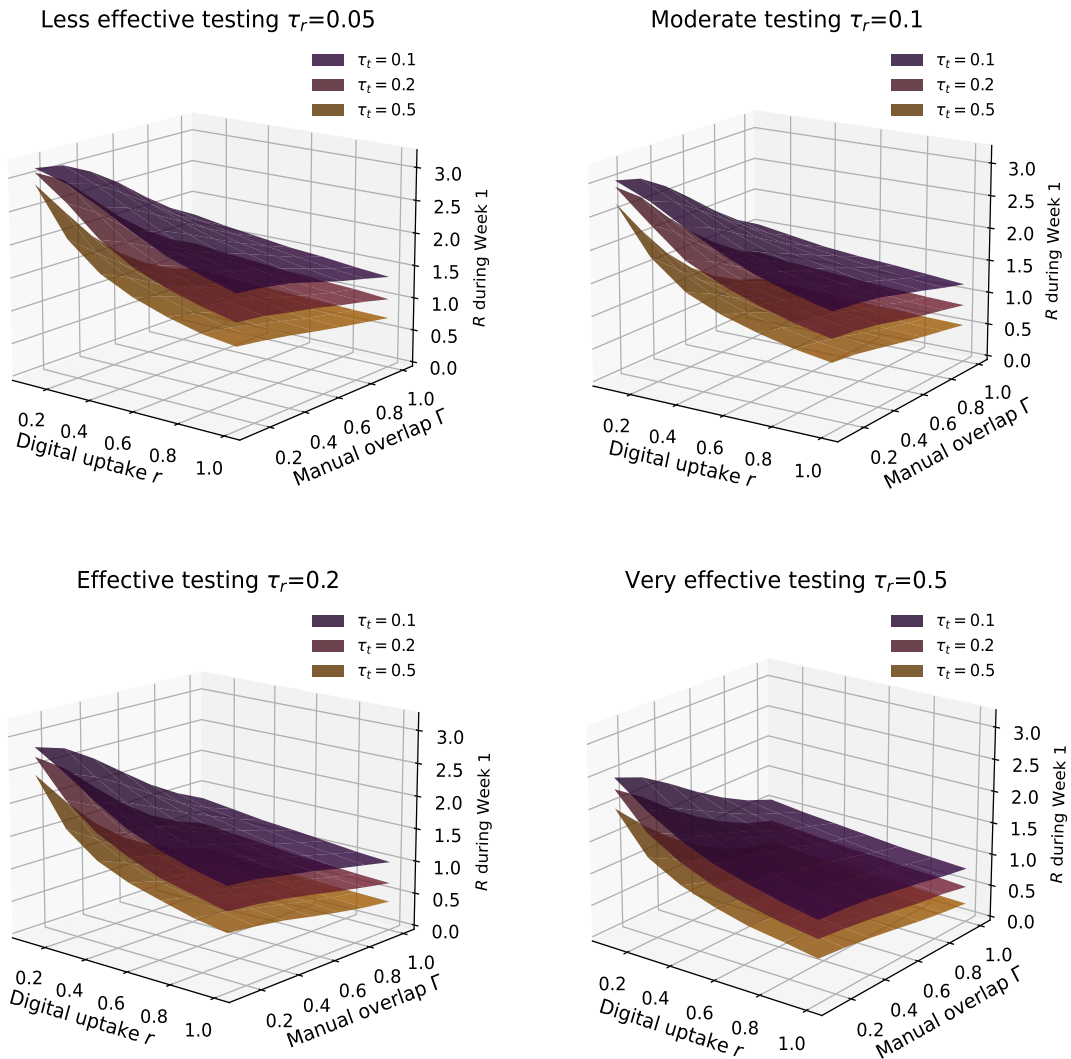


FIGURE B.5: Reproduction number R_e vs. uptake r vs. overlap Γ . Results correspond to a Holme-Kim graph topology with $N = 1000$ and 10% initial infected.

Appendix C

Informing Policy using *EpiCURB*

C.1 Introduction

The COVID-19 pandemic has prompted many countries to implement various non-pharmaceutical public health interventions (NPIs) to prevent the spread of the virus and protect their populations. Effective as they may have been, these policies have rarely been popular among the population, and have faced intense scrutiny ever since. Given the choice between implementing interventions that entail significant economic and social costs or allowing the virus to spread uncontrollably, authorities faced difficult decisions that could have impacted many lives. Balancing NPIs between their individual effectiveness and limitations, their co-occurrence and potentially inflicting major unintended consequences has been a persistent challenge for authorities around the globe (Herby et al., 2023; Hopman and Mehtar, 2020). In this chapter, we demonstrate how simulation-based methods, relying on the *EpiCURB* framework presented in Chapter 4, along with suitable visualization tools can support policy makers in optimizing these decisions and the budgets allocated for each intervention. Additionally, we explore different techniques that they can apply to manage interventions decided by learning-based control agents, providing assurances for audit and the public alike. As such, here we put forward three key advancements:

1. We design a simulation setup that policy makers can easily utilize to examine the impact of individual NPIs when combined with other interventions. We also suggest appropriate visualization tools that can inform them where additional resources would need to be deployed to achieve a given infection threshold. Equipped with a cost-based model of increasing the strength/stringency of each intervention, authorities would ultimately be able on the basis of these results to determine which action is more cost-effective.

2. When interventions are prioritized via learning-based agents, their corresponding embedding space often contains insightful information about the epidemic state. We propose visualizations for this learned space that policy makers can employ to identify at-risk regions where mass interventions with limited budgets, such as vaccination campaigns, should be deployed. These plots can also aid users in establishing when an agent starts underperforming.
3. We illustrate how policy makers can investigate the factors underpinning the decisions taken by the learning-based control agents, as well as devise or justify their strategy to auditing bodies, using a deep learning explanation technique designed for graph neural networks: GraphLIME (Huang et al., 2022).

C.2 Policy choices

According to IOM (2020), by June 2020, 219 countries, territories or areas had implemented at least one form of mobility restriction, affecting billions of people in total. These stringent measures represent some of the most controversial decisions that authorities have taken, as they have limited the public's social interactions, as well as affected the economy and society (Brodeur et al., 2021; Donnelly, 2023). Stay-at-home orders (also known as lockdowns) are NPIs that restrict the movement and interaction of people within a community, often with the exception of essential activities, such as obtaining food, healthcare, work, or physical activity. The extent and duration of these orders can widely differ depending on the context (e.g. availability of vaccines) and the authorities' objectives (e.g. bring the effective reproduction number R under 1). For instance, some stay-at-home orders may target specific segments of the population, such as the elderly or those with underlying health conditions, while others may encompass the entire population. Some lockdowns may be partial, affecting only certain areas or sectors, while others may be total, impacting the whole country or region (Haider et al., 2020; Islam et al., 2020). The effectiveness and impact of stay-at-home orders have been widely debated and studied in the literature. Several works have found that these measures were pivotal for reducing the caseload of COVID-19, and thus implicitly preventing numerous deaths (Flaxman et al., 2020; Hsiang et al., 2020; Ferguson et al., 2020). What is more, some studies have suggested that the timing of introducing lockdowns is crucial for their effectiveness and that delaying or relaxing them too early can lead to disproportionately more deaths (Ferguson et al., 2020; Davies et al., 2021; Mishra et al., 2021; Booth, 2021). In stark contrast, a few authors have challenged these findings and argued that stay-at-home orders had little to no effect on the COVID-19 mortality (Chaudhry et al., 2020; Bendavid et al., 2021; Allen, 2022; Herby et al., 2023). Claiming to be the first comprehensive analysis of NPIs across 180 countries using data from 2020, Herby et al. (2023) have argued that lockdowns had no clear impact on COVID-19 excess deaths or

mortality rates compared with countries that imposed them later, less stringently, or not at all. By the authors' claim, one of the main contributing factors to this result has been the emergence of voluntary behavioral changes among the population, which have significantly contributed to the reported reduction in mixing, irrespective of enforced measures. The work has also pointed out that lockdowns had profoundly negative effects on economic growth, mental health, education, civil liberties, and democracy. Despite being praised by some media outlets and commentators for challenging the conventional wisdom on stay-at-home orders and providing evidence for alternative approaches to managing viral outbreaks (Donnelly, 2023), the study has been criticized by several experts for its methodological flaws, data quality issues, causal inference problems, and ideological bias (Sample and Geddes, 2023). Most notably, the meta-analysis fails to consider the timing of mandate enactment (Sample and Geddes, 2023), which is a critical factor for their effectiveness according to the literature (Ferguson et al., 2020; Davies et al., 2021; Mishra et al., 2021; Booth, 2021). A salient example is the UK government's lockdowns policy, which has faced frequent criticism for its suboptimal timing, as well as its laxity and inconsistency, that may have ultimately undermined its overall effectiveness (Ferguson et al., 2020; Booth, 2021). In this chapter, we highlight how a simulation-based method, underpinned by the *EpiCURB* framework introduced in Chapter 4, can be applied to determine the optimal stringency levels and timing for interventions like lockdowns, taking into account the observed reductions in social mixing, driven by either voluntary behavioural change or imposed limits on gatherings.

Mask mandates have also been the subject of extensive debates throughout the COVID-19 pandemic (Picheta, 2021; Tufekci, 2023), with several countries choosing a more relaxed yet inconsistent approach to this issue (Smith et al., 2020; Baker, 2021; Rolander, 2021). Early studies have indicated masks have a significant positive effect on the spread of SARS-CoV-2, with influential meta-analysis placing the pooled relative risk (RR, proportion of infected among the masked group divided by the proportion of infected among the unmasked group) or the odds ratio (OR, divided odds of getting infected versus not getting infected from both groups) at 0.34 (Chu et al., 2020) or 0.38 (Li et al., 2021). The works later summarized in the comment paper of Brooks and Butler (2021) have added to the body of evidence regarding their effectiveness, with notable mentions including the natural experiments of Mitze et al. (2020), which report a 47% decrease in the infection growth rate after the introduction of mask mandates in Germany, and Karaivanov et al. (2021), which estimate a 25%-40% decline in the weekly diagnoses following these mandates in Canada. Lab studies have also confirmed that masks significantly reduce the transmission of viral-infused droplets, with the common surgical and cotton masks shown to block around 50% when the spreader wears it, and 20%-40% when the receiver uses it (Ueki et al., 2020). More recently, however, a Cochrane review that sparked public debate has pointed to the lack of reliable randomized control trials in the aforementioned

works, concluding that existing evidence on mask efficiency and of higher certainty is limited, and points to an insignificant effect of wearing them (Jefferson et al., 2023). The review's controversial conclusions have been met with considerable criticism (Tufekci, 2023; McDonald, 2023), mostly targeted at the inclusion of a majority of pre-pandemic studies in the meta-analysis underpinning them. Moreover, as the authors admit themselves, the relatively low adherence observed in the scrutinized trials 'hampers drawing firm conclusions'. Indeed, compliance with mask recommendations is known to have been strikingly low in previous viral outbreaks, and it has continued to be low during the COVID-19 pandemic in countries like the UK and Netherlands, while others have seen significant surges in their uptake (Royal Society, 2020). These compliance rates have been significantly influenced by the level of perceived risk and trust among the population of each country, which has often been adversely affected by the lack of clear and consistent guidance from the authorities and scientific bodies in those regions, and the insufficient evidence on the isolated benefits of masking compared to other preventive measures, such as hand hygiene or reducing social contacts (Royal Society, 2020). Undoubtedly, the task of disentangling the effect of each intervention remains very challenging (Donnelly, 2023), and without this knowledge, consistency and compliance with mandates may be compromised. Our simulation-based method can shed some light on this issue, enabling policy makers to evaluate the impact of masking in combination with other interventions, as well as to identify the optimal balance between their coverage or adoption. Furthermore, the visualizations we propose can facilitate more coherent public health policies that can ultimately enhance the civic compliance with them.

For an epidemic control strategy to be successful in the absence of vaccines or total lockdowns, testing and contact tracing have to be efficiently conducted by the authorities. As exemplified in Chapter 3, however, 'test and trace' programmes often suffer from potentially crippling inadequacies. First of all, manual tracing is upper bounded by staff numbers and their efficiency, leading to intrinsic delays (Ferretti et al., 2020), with memory fallacies of positively-tested individuals often ensuing (Garry et al., 2021; Mancastroppa et al., 2021). Digital tracing, on the other hand, is rarely optimally adopted by populations, with issues such as smartphone access, Bluetooth reliability or privacy concerns being important determining factors (Anglemyer et al., 2020; Mancastroppa et al., 2021; Chen and Thio, 2021). Anglemyer et al. (2020) highlight the importance of understanding how manual and digital contact tracing complement each other in controlling the spread of SARS-CoV-2, and suggest further research on their combined effects is required. Other studies introduce modelling frameworks aimed at exploring how different levels of tracing efficiency affect the epidemic outcomes (see Mancastroppa et al. (2021) and Chapter 3). Using the same simulation framework and visualization setup as for other interventions, we investigate the trade-off between manual and digital tracing that policy makers can take into account when distributing resources for contact tracing. In contrast, the

second inadequacy stems from the costs associated with running widespread testing and tracing programmes, both in terms of monetary (Mason et al., 2022) and human resources (Pandit et al., 2022; Rimmer, 2021). In Chapter 4, we advocate for the use of reinforcement learning-based agents to maximize these interventions' impact under limited budgets. In this chapter, we take the idea a step forward by illustrating how the balancing of budgets can be optimized to achieve a desired containment level, in a similar fashion to the other studied interventions.

When these policy choices become informed by learning-based control agents, their internal knowledge can offer policy makers an important overview of the epidemic state, as well as provide indications of their current performance status. First of all, t -SNE projections of their embedding space h_t can be used for real-time tracking of the disease transmission, aiding in understanding where and when additional public health efforts are needed. Compared to analyzing raw numerical infection data, this approach allows for individual-level interventions to be monitored, while compared to full network visualizations, such as Figs 3.2, 3.3 and 4.7, it can be utilized for tracking significantly larger populations. Secondly, dendrograms describing the clustering of h_t , as well as the infection likelihood of each cluster, can inform policy makers where group-level interventions would be more effective. For a specific example of enacting such interventions, refer to the usage of the *clust* baseline for controlling group-based vaccination in Chapter 4. Finally, the performance of the control agent itself can be scrutinized by inspecting the t -SNE structures described above across time, as well as the cluster cardinalities displayed in the aforementioned dendrograms. When new infections constantly fill this space, nodes rarely get grouped according to their known infection status, or the hierarchical clusters display abnormal characteristics (e.g. most vertices shown as belonging to one cluster, infections appearing in most regions etc.), the agent exhibits signs of suboptimal performance, and thus might require further training.

Despite their efficacy and potential to offload policy makers of some of their duties, we believe learning-based control agents should not have the final say in these sensitive decisions. As such, for any given day in the controlled epidemic, policy makers should have the option to analyze explanations for the control decisions taken by the GNN-based rankers. In Appendix C.6, we showcase how the GraphLIME algorithm (Huang et al., 2022) can be used to explain our reinforcement learning policy's decisions by fitting multiple interpretable nonlinear models to the raw action-values that the underlying deep learning model outputs. The resulting feature importances demonstrate the agent focuses on the information conferred by the network heuristics at the beginning of the simulation, when limited other information exists, then switches its focus to the epidemic state features to test nodes more effectively. Conversely, when the process approaches equilibrium, the emphasis reverts back to the degree and eigenvector centralities.

C.3 Methodology

We utilize the *EpiCURB* framework introduced in Chapter 4 to simulate several epidemics over weighted interaction networks with an average degree of approximately 3, studying a range of different outbreak control and intervention assessment strategies that policy makers may consider. The edge weights are sampled from $\mathcal{U}(0.5, 1)$, and correspond to various degrees of transmissibility an interaction can have, based on its duration and distancing. The pathogen spreads according to the individual-based mean-field mechanics described in Chapter 3, following an SEIR compartmental formulation with $\beta = 0.0791$, $\epsilon = \frac{1}{3.7}$, and $\gamma = 0.05$.

When interventions are prioritized by an RL agent, the epidemic is allowed to progress unhinged until $c_a = 5$ days and $c_i = 5\%$ infections have been fulfilled. We train these agents according to the routines described in Section 4.3.5, and study their behavior under different budgeting schemes. The rest of the studied interventions occur stochastically across the network, and begin after 5 individuals have been exposed to the pathogen. To simulate different levels of mixing reduction, we effectively remove a varying proportion of the edges from the infection network. For mask mandates, we assume a percentage of compliant wearers among the population, considering their interactions to have a 50% lower infection weight than the original sampled value (or 75% if both contacts use one). Stay-at-home orders presume the targeted fraction of the nodes and their household immediately discontinue their social patterns, stopping the disease from being transmitted to/from these hubs. Finally, modelling the trade-off between digital and manual tracing is based on the overlap Γ and uptake r parameters introduced in Chapter 3, occurring under the setting $\tau_r = \tau_t = 0.05$.

To draw the latent space of the RL agent, we use the *t-SNE* mappings (van der Maaten and Hinton, 2008) produced by *scikit-learn* (Pedregosa et al., 2011), using PCA as the embedding initialization. To cluster the node embeddings, we employ the *HierarchicalClustering* routine of *scikit-learn*. Lastly, to explain the predictions made by the RL agent, we utilize the GraphLIME method with k_{RBF} kernels for both the feature and the output spaces.

C.4 Cost-based balancing of public health interventions

Outbreak simulations are a useful tool for determining the impact of different public health interventions. However, when modelling various policy options simultaneously, there is no clear indication of which interventions need more efforts or the expected costs associated with enhancing them. Here, we showcase how a simple yet powerful simulation-based method can guide the authorities in making these difficult decisions.




FIGURE C.1: **Contour plots of the containment levels (i.e. percentage of population kept healthy) achieved by combinations of different public health interventions.** This figure illustrates how policy makers can assess and allocate resources to each measure, according to the desired level of containment. The plots on the first row have marked the reference points mentioned in text, with arrows delineating the space of effort-increase actions needed to reach the next infection threshold. The Digital-Manual tracing plot also highlights the most important region of the space, where uptakes are between 20-50%, overlaps are between 10-60%, and 10-15% increases in each direction leads to $\approx 8\%$ less infections. Contours are obtained by averaging 50-100 simulations, and applying a Gaussian filter to remove remnant stochastic noise.

Our approach entails creating a visual space whereby the policy makers can easily study the impact of increasing/decreasing the effort/expenditure for each type of intervention. This space can take many different shapes, such as contour or 3D plots, depicting the influence that combinations of interventions at different strengths have on the simulated epidemic outcome. For example, Fig C.1 illustrates the average fraction of individuals kept healthy from 50-100 epidemic runs over several Holme-Kim networks of $N = 2000$ nodes, $m_{HK} = 3$ and $p_{\Delta} = 0.2$, when the following interventions occur simultaneously: reducing mixing and stay-at-home orders; reducing mixing and mask mandates; classic digital and manual tracing processes, underpinned by a random testing procedure; or RL-targeted testing and tracing.

A policy maker would thus be presented with several choices of the form: if the goal is to achieve a 10-15% reduction in the spread, would it be preferred to move along the X axis or the Y axis? For example, assume that various factors, such as voluntary behavioral changes, closing entertainment venues, schools etc., have reduced the social mixing by about 33% of the pre-epidemic level, and that about 20% of the investigated community have been advised to stay at home (e.g. over 70s). The first plot of Fig C.1 shows that a 10% average reduction in the spread can be achieved either by further reducing the social mixing by 7.5%, or by increasing the fraction of people staying at home by 12% of the total population. Similarly, the second plot shows that increasing the mask wearing fraction from 27% to 41% has a comparable effect to decreasing the social mixing by more than 25%, resulting in 15% less infections overall. By applying a transparent cost-benefit analysis that accounts for the economic and social implications of scaling up interventions, authorities can use such simulated outcomes to determine which of them are most effective to implement or expand.

A key consideration for allocating resources for testing and tracing is the presence of diminishing returns: as the desired infection reach becomes more stringent, the required budget increase has to be larger to achieve it. The third plot of Fig C.1 illustrates this point clearly. Additionally, the latter also reveals that, for the parameter configuration under consideration, a reduction of about 8% in the pathogen attack rate is achieved by increasing either the digital tracing adoption rates within the acceptable and feasible region (20-50%, see Section 3.4.5) or the manual tracing coverage in the moderate range (10-60%) by approximately 10-15% (with smaller effects for the uptake, however). Here, a cost-based assessment can establish which of these actions is more advantageous in each particular circumstance. For example, at lower adoption levels (i.e. $< 30\%$), the uptake can be more viable to improve, since simple usability or privacy enhancements could be enough to attract more users. In contrast, beyond a certain threshold, increasing the adoption of any application is far more challenging, while improving the staffing of the manual process could become more achievable.

Finally, the last plot of Fig C.1 presents the daily budgeting trade-off between testing and digital tracing when an RL agent is tasked to prioritize both processes. Using such visualizations maximizes the benefits of utilizing our targeted approach for epidemic control since policy makers can optimize the budgets allocated for each intervention in a direct manner. When tests in the studied community are insufficient, rationalizing them without compromising on the epidemic outcome is possible by increasing the number of contacts that are to be isolated accordingly. Conversely, if the testing budget is less strict, more contacts can be allowed to continue their normal behavior. The analysis reveals that the two processes have a similar dependence on the budget allocation, highlighting not only the significance of contact tracing, but also the importance of balancing the two for an effective pathogen containment.

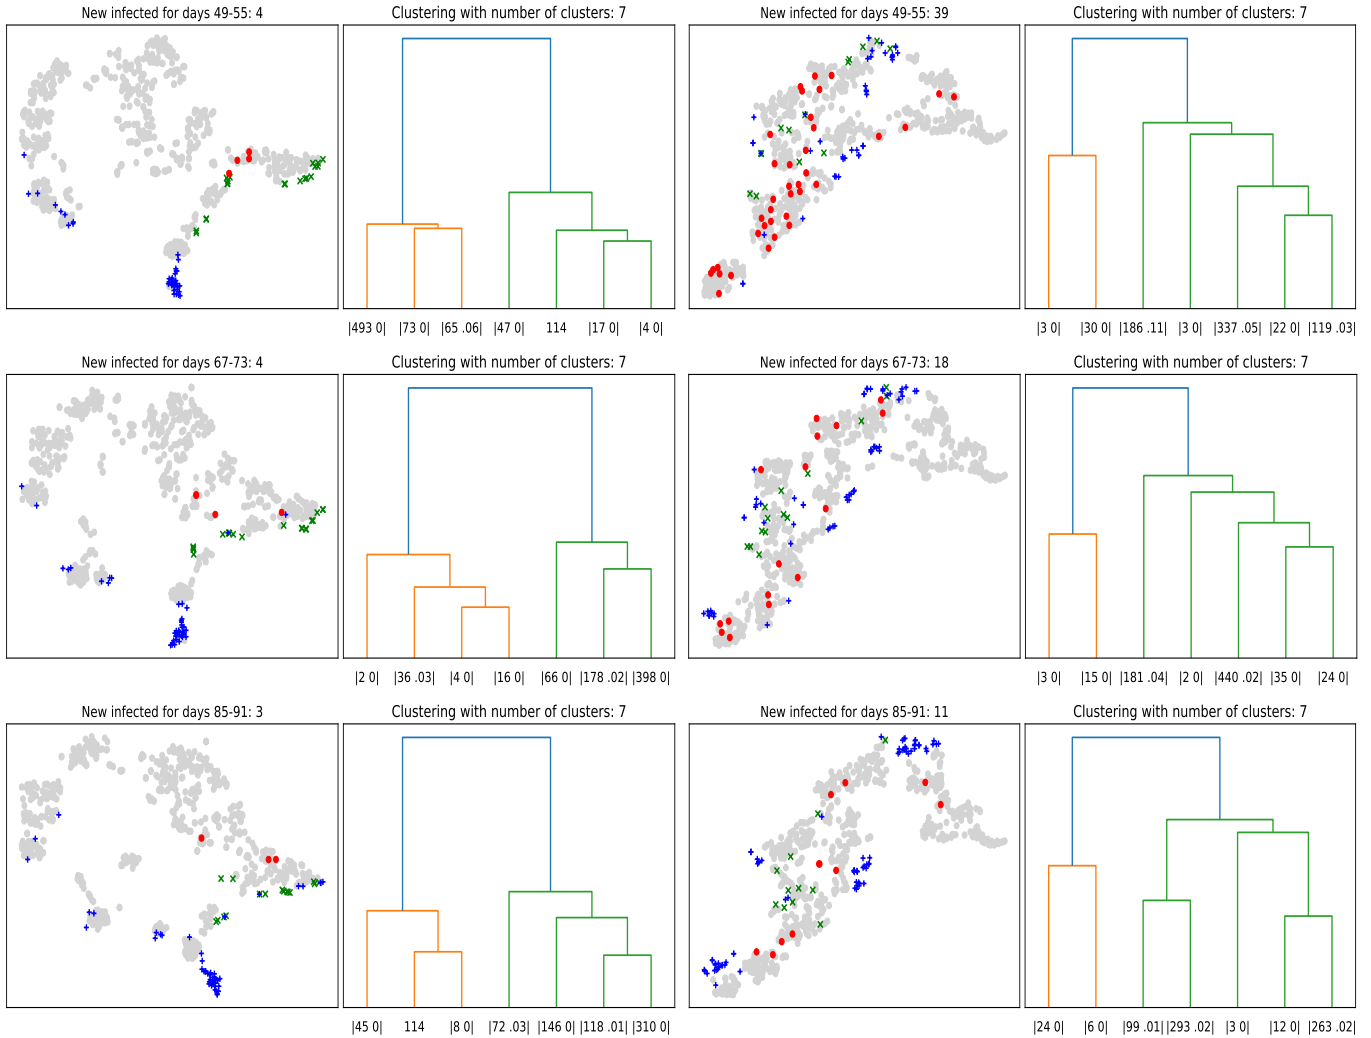


FIGURE C.2: t -SNE plots of the node hidden states h_t and dendrograms of their hierarchical clustering into 7 groups. The first column shows a GNN ranker after training, while the second column shows the same model with initial parameter values. We plot multiple day intervals for a network of $N = 700$ nodes, fixing the daily testing budget to $k = 5$. The latent space of the trained ranker is informative, often grouping the new infected (red) close to the recent negatives (green), while pushing most detected individuals (blue) to specific regions. In contrast, the untrained model is unable to cluster the nodes in a meaningful manner, with infections constantly soaring.

C.5 Inspecting the latent space of learning-based agents

Given the superior performance exhibited by the learning-based control agents presented in Chapter 4, it is sensible to assume the afferent node hidden states h_t encode insightful patterns that underpin their success. The first column of Fig C.2 showcases multiple t -SNE mappings and dendrograms corresponding to the node embeddings h_t of an RL agent during different time periods of an epidemic simulation over a Barabási-Albert network with 700 nodes. We note that the detected positives (colored in blue) have a tendency to be grouped together, while the new infections (colored in red) get pushed to a handful of clusters within the same region as the


recent negative tests (colored in green). Judging by the infection likelihood pertaining to each of the dendrogram's clusters, which can be obtained by averaging past data or by predicting with an SL agent, policy makers can identify groups at risk that should be prioritized for focused health interventions, such as localized mask mandates or mass vaccination campaigns.

By comparison, the second column displays similar plots, but this time for an untrained model with randomly initialized parameters. As a consequence of using the message passing routines of GNNs, the afferent embedding space presents some structure even before model training. However, one can observe there are no boundaries that clearly delimit a specific vertex category in this space, while the infection counts remain very high throughout all intervals. In fact, by controlling testing with the trained model, only 147 out of the simulated 700 people eventually get infected, while the untrained model is unable to prevent more than 600 overall infections. These visualizations, together with the quantitative assessment of the simulation performance, can thus act as strong indicators for the necessity of (re)training the model.

C.6 Explaining the decisions taken by the GNN ranker

For any given day in the controlled epidemic, one can generate GraphLIME explanations for the decisions taken by the learning-based agents. Fig C.3a presents the GraphLIME feature importances for day 9 of an epidemic simulation over a Barabási-Albert network, revealing that the RL agent preferentially attends to the centrality features in the early stages of the outbreak when it possess limited information about the diffusion state. However, as soon as the tester records several positive individuals in the vicinity of a vertex, the rank of the latter markedly increases. After neighborhoods become filled with detected infections, the agent starts targeting the affected sectors, thus switching its focus to the epidemic state features (see Fig C.3b). As the results in Chapter 4 also suggest, the degree remains an effective predictor of a node's importance during the entire diffusion process. Interestingly, the *untested* flag is often correlated with the action scores during the first 30 days of the simulation, which may indicate the agent favors exploring unknown sectors in certain scenarios, while in others it reinforces the testing of recently-targeted regions.


In the later stages, the flag indicating whether a node has ever tested positive gains more importance for its neighbors' rank. Conversely, the number of infected within the immediate hops becomes highly correlated with the degree centrality and this flag, thus providing less additional information compared to earlier days. GraphLIME recognizes this by reducing the relative importance of the corresponding features (see Fig C.4). Fig C.4a also illustrates that the random features intermittently play a role in



(A) Initial predictions.

(B) Early midway predictions.

FIGURE C.3: **Feature importances β_{ip} in the earlier stages.** (a) Initially, the agent does not possess enough information about the epidemic state, and therefore mostly focuses on the centrality features. (b) After a while, the agent starts attending to the epidemic state features as well, such as the previously untested or the positive flags, and the number of infected neighbors. Top row displays each node’s feature values, while neighborhood averages are shown underneath.



(A) Later midway predictions.

(B) Final-days predictions.

FIGURE C.4: **Feature importances β_{ip} in the final stages.** (a) The epidemic state features continue to be important predictors of a node’s rank in the final segment of the simulation, while the degree centrality remains the top predictor. (b) In the final days, both the degree and the eigenvector centrality become again the most important features for ranking vertices, with the epidemic state information either smoothing out or becoming highly correlated with the latter.

the agent’s decision making, by encoding information that is often missing from the other features. As the simulation advances, the epidemic state features begin converging to similar values due to the underlying diffusion process and the message passing iterations, reducing their discriminability potential. The random features, in contrast, encode positional and structural information that is specific to each network configuration, enhancing the expressiveness of the agent’s state representation. By augmenting the set of features with these additional values, the agent can single out different nodes and identify their roles in the graph, adjusting its policy accordingly.

As the epidemic process nears an equilibrium state, in the final days of the simulation, when most nodes are either recovered or isolated, both centrality measures emerge as the dominant predictors once again. Fig C.4b corroborates this observation.

C.7 Conclusion and future work

This chapter demonstrates how *EpiCURB* and appropriate visualization techniques can be used to inform various aspects of decision-making in epidemic control. As such, network-based simulations can reveal the suitable stringency/strength level of each public health intervention required for attaining the desired degree of containment. When standard packages of measures are deployed, our approach can inform policy makers where additional resources should be allocated. Balancing these interventions in terms of the associated budgets and outcomes is paramount for the effective functioning of a society during severe pathogen outbreaks, and therefore implementing and assessing our proposed framework on specific real-world scenarios could make us better prepared for future epidemics.

When learning-based autonomous agents are utilized as a platform for public health decision making, their performance has to be closely monitored by the authorities. Here, we bring forward different visualization techniques for addressing this requirement, enabling policy makers to audit every node-level decisions through local perturbation-based models or inspect the epidemic state and the agents' status through clustering their latent space. The latter could also facilitate the identification of groups most at risk for the purpose of targeting mass health interventions, such as vaccination campaigns. Future work could expand on our clustering analysis, determining the features that contribute to each grouping, as well as enhancing the proposed approaches for targeting clusters (e.g. adding demographic components). Furthermore, we hypothesize that a temporal analysis of the feature importances obtained from GraphLIME could reveal hidden patterns that better account for the agents' decision-making process, which could ultimately enhance our understanding of their internal mechanisms, and thus their utility.

Appendix D

EpiCURB: Supporting Material

D.1 Open-source model

The link for the simulator implementation can be found in Appendix B.1. The open-source code for the control agents under scrutiny can be analyzed via accessing: <https://github.com/andrei-rusu/control-diffusion>.

D.2 Markov chain Monte Carlo epidemic model fitting

By using Markov chain Monte Carlo (MCMC) algorithms, the *EpiCURB* framework can be adjusted to the specific epidemiological context of different countries or regions, ultimately facilitating the design of realistic targeted programmes for testing, tracing, and vaccinating. The actual state of the pathogen diffusion process is typically reflected more accurately by the total number of hospitalizations and deaths in the target area than by the number of confirmed infections (Hyafil and Moriña, 2021; Wang et al., 2022), as the latter often represents only a small fraction of the real infection count (Giattino, 2020). Hence, the former are customary used as dependent variables for parameter estimation. For instance, the COVID-19 base transmission rate that we use throughout this thesis is derived from hospitalization data from the Île-de-France region (Di Domenico et al., 2020).

In Fig D.1, we present multiple illustrative examples of fitting the SEIR individual-based mean-field model outlined in Section 4.3.1 to the number of COVID-19 deaths reported between 12 January 2020 and 11 July 2020 in different US states.¹ The underlying spreading networks are assumed to be Barabási-Albert graphs

¹COVID-19 infection data from the United States: <https://covidtracking.com>

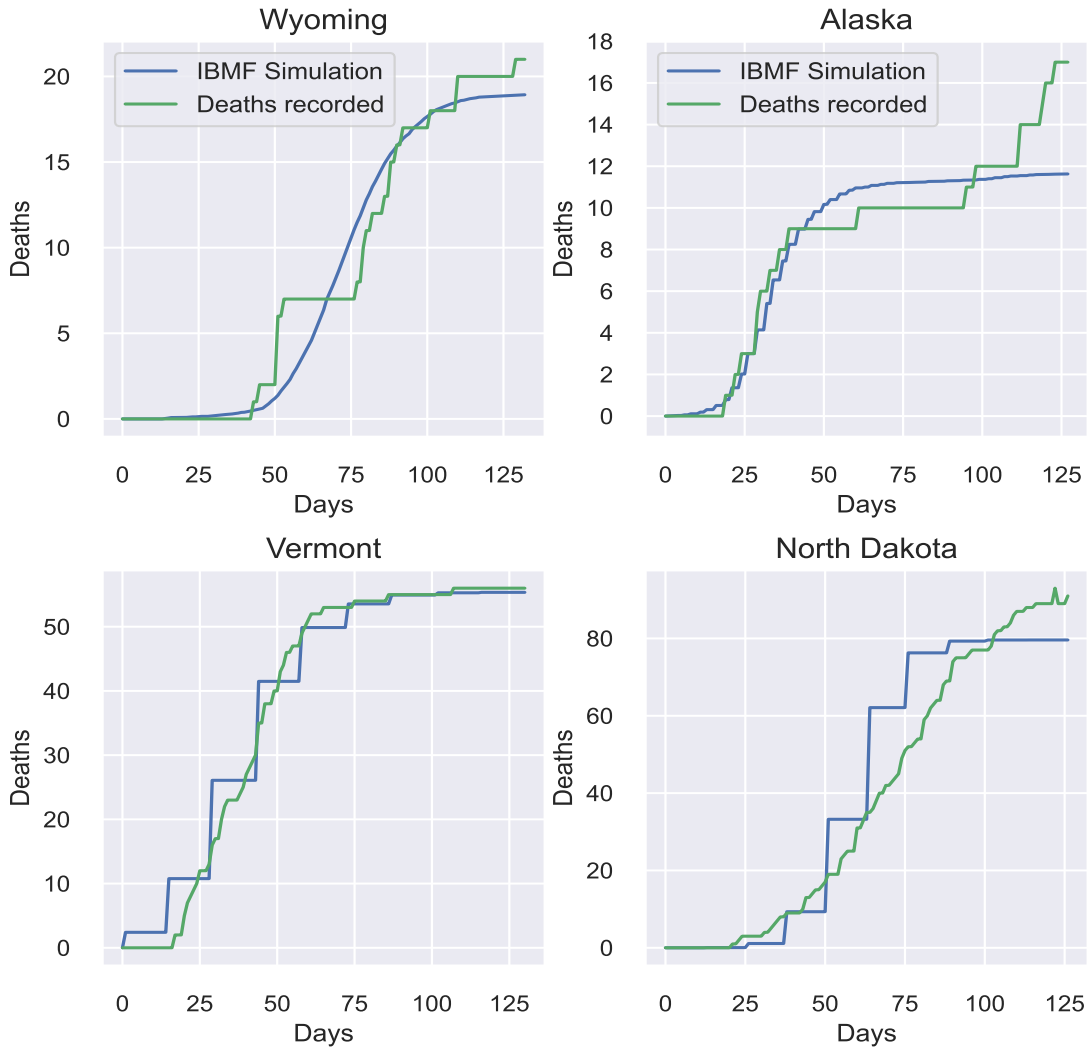


FIGURE D.1: **Mortality curve fitting using MCMC.** Data represents COVID-19 deaths recorded in 4 US states. Simulations performed via the IBMF model of Section 4.3.1.

featuring 10000-20000 nodes and 3-10 average degree, empirically adjusted based on population estimates from each region. To obtain predictions of the number of deaths from the SEIR formulation, we assume that a fraction ρ_D of the R compartment dies. While keeping the exposed duration fixed at $\epsilon^{-1} = 3.7$ days, the set of parameters we fit consists of: the transmission rate β , the recovery rate γ , the number of infections at time t_0 , and the death fraction ρ_D . We obtain the final parameter estimates after running several drawing and tuning steps using standard MCMC routines provided by the *PyMC* library. We use the Metropolis-Hastings algorithm (Hastings, 1970) to sample the posterior of the first three parameters, but we use the No-U-Turn Sampler (NUTS) method (Homan and Gelman, 2014) to sample the posterior of ρ_D , since its gradient can be calculated given the simulated cardinality of the R compartment.

As can be noted from the results above, our model provides a reasonable fit to the given data, largely matching the trends observed in it. Nevertheless, the model exhibits some limitations in dealing with highly irregular regions of the space, which

may stem from data inaccuracies, such as the absence of change between days 52 and 75 in the Wyoming case, or the low death toll recorded between days 34 and 90 in the case of Alaska, surrounded by sharp increases. For reference, we also present an example where the fit is less precise overall. We posit that the curve misalignment in the case of North Dakota is attributable to two factors: the long initial period with no recorded deaths makes it difficult for the model to adjust to the subsequent steep rise, and the size of the network simulated is too small to accurately reflect the more complex viral diffusion patterns recorded for this state.

Appendix E

Tools and Preliminary Experiments

E.1 Tools

Generating network data in this thesis has been carried out using the *networkx* library (Hagberg et al., 2008). For normalizing data, clustering routines, and training simple predictors like SVMs, we have used scikit-learn (Pedregosa et al., 2011) from the classic machine learning toolbox. The GNN models have been implemented using PyTorch (Paszke et al., 2019) and PyTorch Geometric (Fey and Lenssen, 2019). Training has been completed using PyTorch and PyTorch Lightning routines (Falcon et al., 2021). Heavier compute tasks have been performed via submitting SLURM jobs to the Iridis 5 HPC. Finally, experiment logging has been done using Tensorboard (Abadi et al., 2016) and the Neptune AI logger (neptune.ai, 2020).

E.2 Employed or explored network datasets

- Static graph benchmark datasets for node classification: *Cora*, *Citeseer*, *Pubmed*, *ogbn-arxiv*. Citation networks proposed by Yang et al. (2016) and Wang et al. (2020a), widely-used in the graph-based ML literature. Nodes are documents, while edges are citations. Available in PyTorch Geometric and the Open Graph Benchmark (Hu et al., 2020).
- Static graph benchmark datasets for graph classification: *AIDS*, *COX2*, *MUTAG* contain networks derived from small molecules, where class labels encode a biological property (e.g. toxicity). *ENZYMES*, *PROTEINS*, *DD* represent various macromolecules rendered by different graph models. *REDDIT-BINARY*, *IMDB-BINARY* are derived from social networks. All were collected and curated by Kersting et al. (2016), being directly accessible from PyTorch Geometric.

- Raw aggregated mobility data and co-location maps during the COVID-19 pandemic: *Google COVID-19 Community Mobility Reports*, *Facebook Data For Good*, *Cuebiq Data For Good*, *SafeGraph*. Refer to Appendix A.3.1 for more details.
- Dynamic graph datasets describing various diffusion processes: Mobility data and per-region infection estimates: *Covid19 England*, *Chickenpox Hungary*. Twitter data of ‘tweets’ sent between tennis players and news companies: *Twitter Tennis RG and UO*. For more details, please refer to [Rozemberczki et al. \(2021\)](#).
- Small-scale social networks data derived by wearable devices / phone GPS sensing: *Social Evolution* (see Section 3.3.3 and [Madan et al. \(2012\)](#)) and *Reality Mining* ([Eagle and Pentland, 2006](#)) involve students from MIT and their approximated interactions. *Haggle* represents contacts between people measured by carried wireless devices ([Kunegis, 2013](#)). *Highschool* contains the temporal network of contacts between students in a high school in Marseilles ([Fournet and Barrat, 2014](#)).
- Larger social networks of contacts: *Infectious* ([Isella et al., 2011](#)) contains the daily dynamic contact networks collected during the Infectious SocioPatterns event that took place at the Science Gallery in Dublin.

E.3 Preliminary benchmarking of graph-based ML methods

To get started with graph-based machine learning, we have trained several popular GNN architectures and SVM classifiers on the Planetoid datasets: Cora, Citeseer and PubMed ([Yang et al., 2016](#)). For SVMs, we have precomputed node-level heat diffusion kernels, as presented in Section 2.6. For computing the GNN accuracies, we have used either 2 or 3 layers of ReLU- or ELU-activated GCN, GAT, GIN and GraphSAGE with mean aggregation formulations, utilizing an FA layer after these. All hyperparameters have been empirically chosen, except for the initial learning rate that was tuned using PyTorch Lightning ([Falcon et al., 2021](#)). The optimization has been performed over a cross-entropy criterion with a Lasso penalty, using the AdamW optimizer ([Loshchilov and Hutter, 2019](#)), a varying weight decay, early-stopping based on validation accuracy, and a cosine annealing learning rate scheduler. The standard train-validation-test splits has been utilized in each case.

The test results we have obtained from 15 different GNN weight initializations, together with the best performing solution found by the SVM, can be consulted in Table E.1. For GNNs, we provide both the average accuracies, and the standard errors of the means as uncertainty bounds. We observe that the SVM can effectively extract some structural features from the Cora dataset, even though it employs a simple kernel, and it achieves this much more rapidly than any GNN. However, it is unable

TABLE E.1: Test accuracies on node classification.

Model	Cora	Citeseer	PubMed
GCN	81.20 \pm 0.38	68.86 \pm 0.27	77.93 \pm 0.15
GAT	81.71 \pm 0.33	69.26 \pm 0.50	77.67 \pm 0.12
GIN	80.77 \pm 0.32	65.61 \pm 0.27	77.42 \pm 0.51
GraphSAGE (mean agg.)	80.56 \pm 0.27	69.68 \pm 0.32	77.02 \pm 0.16
SVM (diffusion kernel)	72.70	50.70	OOM

to replicate this result for Citeseer, probably due to the latter being a significantly harder dataset (as evidenced by the outcomes of the other models). For PubMed, due to its large number of nodes, the matrix exponentiation needed for computing the heat diffusion kernel could not be performed within our memory bounds.

Bibliography

- M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mane, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viegas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems. *arXiv:1603.04467 [cs]*, March 2016. URL <http://arxiv.org/abs/1603.04467>.
- S. Abdullah and X. Wu. An Epidemic Model for News Spreading on Twitter. In *2011 IEEE 23rd International Conference on Tools with Artificial Intelligence*, pages 163–169, November 2011. .
- M. Abueg, R. Hinch, N. Wu, L. Liu, W. J. M. Probert, A. Wu, P. Eastham, Y. Shafi, M. Rosencrantz, M. Dikovsky, Z. Cheng, A. Nurtay, L. Abeler-Dörner, D. G. Bonsall, M. V. McConnell, S. O’Banion, and C. Fraser. Modeling the combined effect of digital exposure notification and non-pharmaceutical interventions on the COVID-19 epidemic in Washington state. In *medRxiv*, page 2020.08.29.20184135. Cold Spring Harbor Laboratory Press, September 2020. .
- M. Abueg, R. Hinch, N. Wu, L. Liu, W. Probert, A. Wu, P. Eastham, Y. Shafi, M. Rosencrantz, M. Dikovsky, Z. Cheng, A. Nurtay, L. Abeler-Dörner, D. Bonsall, M. V. McConnell, S. O’Banion, and C. Fraser. Modeling the effect of exposure notification and non-pharmaceutical interventions on COVID-19 transmission in Washington state. *npj Digital Medicine*, 4(1):1–10, March 2021. ISSN 2398-6352. .
- D. C. Adam, P. Wu, J. Y. Wong, E. H. Y. Lau, T. K. Tsang, S. Cauchemez, G. M. Leung, and B. J. Cowling. Clustering and superspreading potential of SARS-CoV-2 infections in Hong Kong. *Nature Medicine*, 26(11):1714–1719, November 2020. ISSN 1546-170X. .
- Z. Adamczyk. Diffusion of Particles. In T. Tadros, editor, *Encyclopedia of Colloid and Interface Science*, pages 247–247. Springer, Berlin, Heidelberg, 2013. ISBN 978-3-642-20665-8. .

- A. Adiga, D. Dubhashi, B. Lewis, M. Marathe, S. Venkatramanan, and A. Vullikanti. Mathematical Models for COVID-19 Pandemic: A Comparative Analysis. *Journal of the Indian Institute of Science*, 100(4):793–807, 2020. ISSN 0970-4140. .
- L. E. Aik, L. C. Kiang, T. W. Hong, and M. S. Abu. The SIR model of Zika virus disease outbreak in Brazil at year 2015. *AIP Conference Proceedings*, 1847(1):020004, May 2017. ISSN 0094-243X. .
- R. Albert and A.-L. Barabasi. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74(1):47–97, January 2002. ISSN 0034-6861, 1539-0756. .
- D. W. Allen. Covid-19 Lockdown Cost/Benefits: A Critical Assessment of the Literature. *International Journal of the Economics of Business*, 29(1):1–32, January 2022. ISSN 1357-1516. .
- L. J. S. Allen. A primer on stochastic epidemic models: Formulation, numerical simulation, and analysis. *Infectious Disease Modelling*, 2(2):128–142, May 2017. ISSN 2468-0427. .
- U. Alon and E. Yahav. On the Bottleneck of Graph Neural Networks and its Practical Implications. In *International Conference on Learning Representations*, February 2022. URL <https://openreview.net/forum?id=i800Ph0CVH2>.
- B. Anderson. The Rise of the Weaponized AI Propaganda Machine, March 2017. URL <https://medium.com/join-scout/the-rise-of-the-weaponized-ai-propaganda-machine-86dac61668b>.
- N. Andrews, J. Stowe, F. Kirsebom, S. Toffa, T. Rickeard, E. Gallagher, C. Gower, M. Kall, N. Groves, A.-M. O’Connell, D. Simons, P. B. Blomquist, A. Zaidi, S. Nash, N. Iwani Binti Abdul Aziz, S. Thelwall, G. Dabrera, R. Myers, G. Amirthalingam, S. Gharbia, J. C. Barrett, R. Elson, S. N. Ladhani, N. Ferguson, M. Zambon, C. N. Campbell, K. Brown, S. Hopkins, M. Chand, M. Ramsay, and J. Lopez Bernal. Covid-19 Vaccine Effectiveness against the Omicron (B.1.1.529) Variant. *New England Journal of Medicine*, 386(16):1532–1546, April 2022. ISSN 0028-4793. .
- A. Anglemeyer, T. H. Moore, L. Parker, T. Chambers, A. Grady, K. Chiu, M. Parry, M. Wilczynska, E. Flemyng, and L. Bero. Digital contact tracing technologies in epidemics: A rapid review. *Cochrane Database of Systematic Reviews*, 1(8), 2020. ISSN 1465-1858. .
- L. Ansell and L. Dalla Valle. A new data integration framework for Covid-19 social media information. *Scientific Reports*, 13(1):6170, April 2023. ISSN 2045-2322. .
- S. Arik, C.-L. Li, J. Yoon, R. Sinha, A. Epshteyn, L. Le, V. Menon, S. Singh, L. Zhang, M. Nikoltchev, Y. Sonthalia, H. Nakhost, E. Kanal, and T. Pfister. Interpretable Sequence Learning for Covid-19 Forecasting. In *Advances in Neural Information*

- Processing Systems*, volume 33, pages 18807–18818. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/hash/d9dbc51dc534921589adf460c85cd824-Abstract.html.
- V. Arnaboldi, M. Conti, M. La Gala, A. Passarella, and F. Pezzoni. Ego network structure in online social networks and its impact on information diffusion. *Computer Communications*, 76:26–41, February 2016. ISSN 0140-3664. .
- N. Aronszajn. Theory of Reproducing Kernels. *Transactions of the American Mathematical Society*, 68(3):337–404, 1950. ISSN 0002-9947. .
- K. J. Åström. Optimal Control of Markov Processes with Incomplete State Information I. *Journal of Mathematical Analysis and Applications*, 10:174–205, 1965. ISSN 0022-247X. URL <http://lup.lub.lu.se/record/8867084>.
- L. R. Baden, H. M. El Sahly, B. Essink, K. Kotloff, S. Frey, R. Novak, D. Diemert, S. A. Spector, N. Roupheal, C. B. Creech, J. McGettigan, S. Khetan, N. Segall, J. Solis, A. Brosz, C. Fierro, H. Schwartz, K. Neuzil, L. Corey, P. Gilbert, H. Janes, D. Follmann, M. Marovich, J. Mascola, L. Polakowski, J. Ledgerwood, B. S. Graham, H. Bennett, R. Pajon, C. Knightly, B. Leav, W. Deng, H. Zhou, S. Han, M. Ivarsson, J. Miller, and T. Zaks. Efficacy and Safety of the mRNA-1273 SARS-CoV-2 Vaccine. *New England Journal of Medicine*, 384(5):403–416, 2021. ISSN 0028-4793. .
- J. P. Bagrow, X. Liu, and L. Mitchell. Information flow reveals prediction limits in online social activity. *Nature Human Behaviour*, 3(2):122–128, February 2019. ISSN 2397-3374. .
- D. Bahdanau, K. Cho, and Y. Bengio. Neural Machine Translation by Jointly Learning to Align and Translate. *arXiv:1409.0473 [cs, stat]*, May 2016. URL <http://arxiv.org/abs/1409.0473>.
- S. Baker. Sweden, noted for its lax COVID-19 response, never mandated face masks. Now it’s dropping its vague recommendation to wear one at all. *Business Insider*, July 2021. URL <https://www.businessinsider.com/covid-19-sweden-drops-face-mask-recommendation-never-mandated-them-2021-7>.
- H. Bao, L. Dong, and F. Wei. BEiT: BERT Pre-Training of Image Transformers, June 2021.
- A.-L. Barabási. Network Science Random Networks, March 2015. URL <https://barabasi.com/f/624.pdf>.
- A.-L. Barabási and R. Albert. Emergence of Scaling in Random Networks. *Science*, 286(5439):509–512, October 1999. .
- F. M. Bass. A New Product Growth for Model Consumer Durables. *Management Science*, 15(5):215–227, January 1969. ISSN 0025-1909. .

- H. Bastani, K. Drakopoulos, V. Gupta, I. Vlachogiannis, C. Hadjichristodoulou, P. Lagiou, G. Magiorkinis, D. Paraskevis, and S. Tsiodras. Efficient and targeted COVID-19 border testing via reinforcement learning. *Nature*, 599(7883):108–113, 2021. ISSN 1476-4687. .
- D. Beaini, S. Passaro, V. Létourneau, W. L. Hamilton, G. Corso, and P. Liò. Directional Graph Networks, April 2021.
- I. Bello, H. Pham, Q. V. Le, M. Norouzi, and S. Bengio. Neural Combinatorial Optimization with Reinforcement Learning. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/pdf?id=Bk9mx1SFx>.
- E. Bendavid, C. Oh, J. Bhattacharya, and J. P. A. Ioannidis. Assessing mandatory stay-at-home and business closure effects on the spread of COVID-19. *European Journal of Clinical Investigation*, 51(4):e13484, April 2021. ISSN 0014-2972. .
- T. Berge, J. M.-S. Lubuma, G. M. Moremedi, N. Morris, and R. Kondera-Shava. A simple mathematical model for Ebola in Africa. *Journal of Biological Dynamics*, 11(1): 42–74, January 2017. ISSN 1751-3758. .
- L. M. A. Bettencourt, A. Cintrón-Arias, D. I. Kaiser, and C. Castillo-Chavez. The power of a good idea: Quantitative modeling of the spread of ideas from epidemiological models. *Physica A: Statistical Mechanics and its Applications*, 364: 513–536, May 2006. ISSN 0378-4371. .
- C. Bodnar, F. Di Giovanni, B. P. Chamberlain, P. Liò, and M. M. Bronstein. Neural Sheaf Diffusion: A Topological Perspective on Heterophily and Oversmoothing in GNNs, May 2022.
- B. Bollobás. A Probabilistic Proof of an Asymptotic Formula for the Number of Labelled Regular Graphs. *European Journal of Combinatorics*, 1(4):311–316, December 1980. ISSN 0195-6698. .
- B. Bollobás. *Random Graphs*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, Cambridge, 2 edition, 2001. ISBN 978-0-521-80920-7. .
- P. Bonacich. Factoring and weighting approaches to status scores and clique identification. *The Journal of Mathematical Sociology*, 2(1):113–120, January 1972. ISSN 0022-250X, 1545-5874. .
- P. Bongini, M. Bianchini, and F. Scarselli. Molecular generative Graph Neural Networks for Drug Discovery. *Neurocomputing*, 450:242–252, August 2021. ISSN 0925-2312. .
- R. Booth. Delaying England’s winter lockdown ‘caused up to 27,000 extra Covid deaths’. *The Guardian*, March 2021. ISSN 0261-3077. URL

- <https://www.theguardian.com/uk-news/2021/mar/18/delaying-englands-winter-lockdown-caused-up-to-27000-extra-covid-deaths>.
- K. M. Borgwardt and H.-P. Kriegel. Shortest-Path Kernels on Graphs. In *Fifth IEEE International Conference on Data Mining (ICDM'05)*, pages 74–81, Houston, TX, USA, 2005. IEEE. ISBN 978-0-7695-2278-4. .
- W. E. Boyce, R. C. DiPrima, and D. B. Meade. Elementary differential equations and boundary value problems. In *Elementary Differential Equations and Boundary Value Problems*, pages 363–400. Wiley, Hoboken, twelfth edition edition, 2022. ISBN 978-1-119-77767-0 978-1-119-82053-6.
- D. Braha and Y. Bar-Yam. From centrality to temporary fame: Dynamic centrality in complex networks. *Complexity*, 12(2):59–63, 2006. ISSN 1099-0526. .
- S. Brin and L. Page. The anatomy of a large-scale hypertextual Web search engine. *Computer Networks and ISDN Systems*, 30(1):107–117, April 1998. ISSN 0169-7552. .
- A. Brodeur, A. E. Clark, S. Fleche, and N. Powdthavee. COVID-19, lockdowns and well-being: Evidence from Google Trends. *Journal of Public Economics*, 193:104346, January 2021. .
- S. Brody, U. Alon, and E. Yahav. How Attentive are Graph Attention Networks? In *International Conference on Learning Representations*, January 2022. URL <https://openreview.net/forum?id=F72ximsx7C1>.
- M. Bronstein. Expressive power of graph neural networks and the Weisfeiler-Lehman test, October 2020. URL <https://towardsdatascience.com/expressive-power-of-graph-neural-networks-and-the-weisfeiler-lehman-test-b883db3c7c49>.
- M. Bronstein. Deep learning on graphs: Successes, challenges, and next steps, January 2022. URL <https://towardsdatascience.com/deep-learning-on-graphs-successes-challenges-and-next-steps-7d9ec220ba8>.
- J. T. Brooks and J. C. Butler. Effectiveness of Mask Wearing to Control Community Spread of SARS-CoV-2. *JAMA*, 325(10):998–999, March 2021. ISSN 0098-7484. .
- T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei. Language Models are Few-Shot Learners, July 2020.
- J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun. Spectral Networks and Locally Connected Networks on Graphs. *arXiv:1312.6203 [cs]*, May 2014. URL <http://arxiv.org/abs/1312.6203>.

- K. J. Bruxvoort, L. S. Sy, L. Qian, B. K. Ackerson, Y. Luo, G. S. Lee, Y. Tian, A. Florea, M. Aragonés, J. E. Tubert, H. S. Takhar, J. H. Ku, Y. D. Paila, C. A. Talarico, and H. F. Tseng. Effectiveness of mRNA-1273 against delta, mu, and other emerging variants of SARS-CoV-2: Test negative case-control study. *BMJ*, 375:e068848, December 2021. ISSN 1756-1833. .
- J. C. Butcher. Numerical methods for ordinary differential equations in the 20th century. *Journal of Computational and Applied Mathematics*, page 29, 2000.
- Q. Cappart, D. Chételat, E. Khalil, A. Lodi, C. Morris, and P. Veličković. Combinatorial optimization and reasoning with graph neural networks. *arXiv:2102.09544 [cs, math, stat]*, April 2021. URL <http://arxiv.org/abs/2102.09544>.
- C. Cattuto, W. V. den Broeck, A. Barrat, V. Colizza, J.-F. Pinton, and A. Vespignani. Dynamics of Person-to-Person Interactions from Distributed RFID Sensor Networks. *PLOS ONE*, 5(7):e11596, July 2010. ISSN 1932-6203. .
- CDC. Health Departments, February 2020. URL <https://www.cdc.gov/coronavirus/2019-ncov/php/contact-tracing/contact-tracing-plan/appendix.html>.
- D. Cereda, M. Tirani, F. Rovida, V. Demicheli, M. Ajelli, P. Poletti, F. Trentini, G. Guzzetta, V. Marziano, A. Barone, M. Magoni, S. Deandrea, G. Diurno, M. Lombardo, M. Faccini, A. Pan, R. Bruno, E. Pariani, G. Grasselli, A. Piatti, M. Gramegna, F. Baldanti, A. Melegaro, and S. Merler. The early phase of the COVID-19 outbreak in Lombardy, Italy. Technical report, March 2020. URL <http://arxiv.org/abs/2003.09320>.
- Y. H. Chae, C. Lee, M. K. Choi, and P. H. Seong. Evaluating attractiveness of cyberattack path using resistance concept and page-rank algorithm. *Annals of Nuclear Energy*, 166:108748, February 2022. ISSN 0306-4549. .
- B. Chamberlain, J. Rowbottom, M. I. Gorinova, M. Bronstein, S. Webb, and E. Rossi. GRAND: Graph Neural Diffusion. In *Proceedings of the 38th International Conference on Machine Learning*, pages 1407–1418. PMLR, July 2021. URL <https://proceedings.mlr.press/v139/chamberlain21a.html>.
- S. Chang, E. Pierson, P. W. Koh, J. Gerardin, B. Redbird, D. Grusky, and J. Leskovec. Mobility network models of COVID-19 explain inequities and inform reopening. *Nature*, 589(7840):82–87, January 2021. ISSN 1476-4687. .
- R. Chaudhry, G. Dranitsaris, T. Mubashir, J. Bartoszko, and S. Riazi. A country level analysis measuring the impact of government actions, country preparedness and socioeconomic factors on COVID-19 mortality and related health outcomes. *eClinicalMedicine*, 25, August 2020. ISSN 2589-5370. .

- A. T.-Y. Chen and K. W. Thio. Exploring the drivers and barriers to uptake for digital contact tracing. *Social Sciences & Humanities Open*, 4(1):100212, January 2021. ISSN 2590-2911. .
- L. Chen, Z. Chen, and J. Bruna. On Graph Neural Networks versus Graph-Augmented MLPs. *arXiv:2010.15116 [cs, math, stat]*, December 2020. URL <http://arxiv.org/abs/2010.15116>.
- R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud. Neural Ordinary Differential Equations. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL <https://papers.nips.cc/paper/2018/hash/69386f6bb1dfed68692a24c8686939b9-Abstract.html>.
- W. Chen, Y. Wang, and S. Yang. Efficient influence maximization in social networks. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '09*, page 199, Paris, France, 2009. ACM Press. ISBN 978-1-60558-495-9. .
- M. D. Choudhury, Y.-R. Lin, H. Sundaram, K. S. Candan, L. Xie, and A. Kelliher. How Does the Data Sampling Strategy Impact the Discovery of Information Diffusion in Social Media? *Proceedings of the International AAAI Conference on Web and Social Media*, 4(1):34–41, May 2010. ISSN 2334-0770. .
- D. K. Chu, E. A. Akl, S. Duda, K. Solo, S. Yaacoub, H. J. Schünemann, D. K. Chu, E. A. Akl, A. El-harakeh, A. Bognanni, T. Lotfi, M. Loeb, A. Hajizadeh, A. Bak, A. Izcovich, C. A. Cuello-Garcia, C. Chen, D. J. Harris, E. Borowiack, F. Chamseddine, F. Schünemann, G. P. Morgano, G. E. U. M. Schünemann, G. Chen, H. Zhao, I. Neumann, J. Chan, J. Khabza, L. Hneiny, L. Harrison, M. Smith, N. Rizk, P. G. Rossi, P. AbiHanna, R. El-khoury, R. Stalteri, T. Baldeh, T. Piggott, Y. Zhang, Z. Saad, A. Khamis, M. Reinap, S. Duda, K. Solo, S. Yaacoub, and H. J. Schünemann. Physical distancing, face masks, and eye protection to prevent person-to-person transmission of SARS-CoV-2 and COVID-19: A systematic review and meta-analysis. *The Lancet*, 395(10242):1973–1987, June 2020. ISSN 0140-6736, 1474-547X. .
- F. Chung, P. Horn, and A. Tsiatas. Distributing Antidote Using PageRank Vectors. *Internet Mathematics*, 6(2):237–254, 2009. ISSN 1542-7951, 1944-9488. .
- R. Clair, M. Gordon, M. Kroon, and C. Reilly. The effects of social isolation on well-being and life satisfaction during pandemic. *Humanities and Social Sciences Communications*, 8(1):1–6, January 2021. ISSN 2662-9992. .
- M. A. Clark, D. Matthew, and C. Jung. Passive Transport. In *Biology 2e*, Biology 2e. OpenStax, March 2018. URL <https://openstax.org/books/biology-2e/pages/5-2-passive-transport>.

- R. Cohen, S. Havlin, and D. ben-Avraham. Efficient Immunization Strategies for Computer Networks and Populations. *Physical Review Letters*, 91(24):247901, December 2003. ISSN 0031-9007, 1079-7114. .
- H. Dai, E. Khalil, Y. Zhang, B. Dilkina, and L. Song. Learning Combinatorial Optimization Algorithms over Graphs. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/hash/d9896106ca98d3d05b8cbdf4fd8b13a1-Abstract.html.
- M. Dashtbali and M. Mirzaie. A compartmental model that predicts the effect of social distancing and vaccination on controlling COVID-19. *Scientific Reports*, 11(1):8191, April 2021. ISSN 2045-2322. .
- N. G. Davies, R. C. Barnard, C. I. Jarvis, T. W. Russell, M. G. Semple, M. Jit, and W. J. Edmunds. Association of tiered restrictions and a second lockdown with COVID-19 deaths and hospital admissions in England: A modelling study. *The Lancet Infectious Diseases*, 21(4):482–492, April 2021. ISSN 1473-3099, 1474-4457. .
- E. L. Davis, T. C. D. Lucas, A. Borlase, T. M. Pollington, S. Abbott, D. Ayabina, T. Crellen, J. Hellewell, L. Pi, G. F. Medley, T. D. Hollingsworth, and P. Klepac. Contact tracing is an imperfect tool for controlling COVID-19 transmission and relies on population adherence. *Nature Communications*, 12(1):5412, September 2021. ISSN 2041-1723. .
- A. Delarue, R. Anderson, and C. Tjandraatmadja. Reinforcement Learning with Combinatorial Actions: An Application to Vehicle Routing. In *Advances in Neural Information Processing Systems*, volume 33, pages 609–620. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/06a9d51e04213572ef0720dd27a84792-Abstract.html>.
- J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv:1810.04805 [cs]*, May 2019. URL <http://arxiv.org/abs/1810.04805>.
- L. Di Domenico, G. Pullano, C. E. Sabbatini, P.-Y. Boëlle, and V. Colizza. Impact of lockdown on COVID-19 epidemic in Île-de-France and possible exit strategies. *BMC Medicine*, 18(1):240, July 2020. ISSN 1741-7015. .
- O. Diekmann, J. A. P. Heesterbeek, and J. A. J. Metz. On the definition and the computation of the basic reproduction ratio R_0 in models for infectious diseases in heterogeneous populations. *Journal of Mathematical Biology*, 28(4):365–382, June 1990. ISSN 1432-1416. .
- A. Dighe, L. Cattarino, G. Cuomo-Dannenburg, J. Skarp, N. Imai, S. Bhatia, K. A. M. Gaythorpe, K. E. C. Ainslie, M. Baguelin, S. Bhatt, A. Boonyasiri, N. F. Brazeau, L. V.

- Cooper, H. Coupland, Z. Cucunuba, I. Dorigatti, O. D. Eales, S. L. van Elsland, R. G. FitzJohn, W. D. Green, D. J. Haw, W. Hinsley, E. Knock, D. J. Laydon, T. Mellan, S. Mishra, G. Nedjati-Gilani, P. Nouvellet, M. Pons-Salort, H. A. Thompson, H. J. T. Unwin, R. Verity, M. A. C. Vollmer, C. E. Walters, O. J. Watson, C. Whittaker, L. K. Whittles, A. C. Ghani, C. A. Donnelly, N. M. Ferguson, and S. Riley. Response to COVID-19 in South Korea and implications for lifting stringent interventions. *BMC Medicine*, 18(1):321, October 2020. ISSN 1741-7015. .
- E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1):269–271, December 1959. ISSN 0945-3245. .
- L. Dong, Y. Li, H. Yin, H. Le, and M. Rui. The Algorithm of Link Prediction on Social Network. *Mathematical Problems in Engineering*, 2013, September 2013. .
- L. Donnelly. Lockdown benefits ‘a drop in the bucket compared to the costs’, landmark study finds. *The Telegraph*, June 2023. ISSN 0307-1235. URL <https://www.telegraph.co.uk/news/2023/06/04/first-lockdown-prevented-1700-deaths-landmark-study-finds/>.
- V. P. Dwivedi and X. Bresson. A Generalization of Transformer Networks to Graphs. In *AAAI 2021 Workshop on Deep Learning on Graphs: Methods and Applications*, January 2021. .
- V. P. Dwivedi, C. K. Joshi, T. Laurent, Y. Bengio, and X. Bresson. Benchmarking Graph Neural Networks. *arXiv:2003.00982 [cs, stat]*, July 2020. URL <http://arxiv.org/abs/2003.00982>.
- N. Eagle and A. Pentland. Reality mining: Sensing complex social systems. *Personal and Ubiquitous Computing*, 10(4):255–268, March 2006. ISSN 1617-4909. .
- A. Endo, S. Abbott, A. J. Kucharski, and S. Funk. Estimating the overdispersion in COVID-19 transmission using outbreak sizes outside China. *Wellcome Open Research*, 5, July 2020. ISSN 2398-502X. .
- P. Erdős and A. Rényi. On random graphs I. *Publicationes Mathematicae Debrecen*, 6:290, 1959.
- Etalab. Données hospitalières relatives à l’épidémie de COVID-19, May 2020. URL <https://www.data.gouv.fr/en/datasets/donnees-hospitalieres-relatives-a-lepidemie-de-covid-19/>.
- E. Even-Dar and A. Shapira. A Note on Maximizing the Spread of Influence in Social Networks. In X. Deng and F. C. Graham, editors, *Internet and Network Economics*, Lecture Notes in Computer Science, pages 281–286, Berlin, Heidelberg, 2007. Springer. ISBN 978-3-540-77105-0. .

- W. Falcon, J. Borovec, A. Walchli, and C. Mocholi. PyTorchLightning/pytorch-lightning. Pytorch Lightning, August 2021. URL <https://github.com/PyTorchLightning/pytorch-lightning>.
- W. Fan, Y. Ma, Q. Li, Y. He, E. Zhao, J. Tang, and D. Yin. Graph Neural Networks for Social Recommendation. *arXiv:1902.07243 [cs]*, November 2019. URL <http://arxiv.org/abs/1902.07243>.
- G. Fanti and P. Viswanath. Deanonymization in the Bitcoin P2P Network. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://papers.nips.cc/paper_files/paper/2017/hash/6c3cf77d52820cd0fe646d38bc2145ca-Abstract.html.
- K. Farrahi, R. Emonet, and M. Cebrian. Epidemic Contact Tracing via Communication Traces. *PLoS ONE*, 9(5):e95133, May 2014. ISSN 1932-6203. .
- K. Farrahi, R. Emonet, and M. Cebrian. Predicting a Community’s Flu Dynamics with Mobile Phone Data. In *Computer-Supported Cooperative Work and Social Computing*, Vancouver, Canada, March 2015. .
- J. M. Ferdinands. Waning 2-Dose and 3-Dose Effectiveness of mRNA Vaccines Against COVID-19—Associated Emergency Department and Urgent Care Encounters and Hospitalizations Among Adults During Periods of Delta and Omicron Variant Predominance — VISION Network, 10 States, August 2021–January 2022. *MMWR. Morbidity and Mortality Weekly Report*, 71, 2022. ISSN 0149-21951545-861X. .
- N. Ferguson, D. Laydon, G. Nedjati Gilani, N. Imai, K. Ainslie, M. Baguelin, S. Bhatia, A. Boonyasiri, ZULMA. Cucunuba Perez, G. Cuomo-Dannenburg, A. Dighe, I. Dorigatti, H. Fu, K. Gaythorpe, W. Green, A. Hamlet, W. Hinsley, L. Okell, S. Van Elsland, H. Thompson, R. Verity, E. Volz, H. Wang, Y. Wang, P. Walker, P. Winskill, C. Whittaker, C. Donnelly, S. Riley, and A. Ghani. Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID19 mortality and healthcare demand. Technical report, Imperial College London, March 2020.
- L. Ferretti. Quantifying dynamics of SARS-CoV-2 transmission suggests that epidemic control is feasible through instantaneous digital contact tracing, 2020. URL https://perso.math.univ-toulouse.fr/cattiaux/files/2020/04/slides_COVID19AppBasedContactTracing_Ferretti_30min.pdf.
- L. Ferretti, C. Wymant, M. Kendall, L. Zhao, A. Nurtay, L. Abeler-Dörner, M. Parker, D. Bonsall, and C. Fraser. Quantifying SARS-CoV-2 transmission suggests epidemic control with digital contact tracing. *Science*, 368, 2020. .
- T. Fetzer and T. Graeber. Measuring the scientific effectiveness of contact tracing: Evidence from a natural experiment. *Proceedings of the National Academy of Sciences*, 118(33):e2100814118, August 2021. .

- M. Fey and J. E. Lenssen. Fast Graph Representation Learning with PyTorch Geometric. *arXiv:1903.02428 [cs, stat]*, April 2019. URL <http://arxiv.org/abs/1903.02428>.
- A. Fick. V. On liquid diffusion. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 10(63):30–39, July 1855. ISSN 1941-5982. .
- S. Flaxman, S. Mishra, A. Gandy, H. J. T. Unwin, T. A. Mellan, H. Coupland, C. Whittaker, H. Zhu, T. Berah, J. W. Eaton, M. Monod, A. C. Ghani, C. A. Donnelly, S. Riley, M. A. C. Vollmer, N. M. Ferguson, L. C. Okell, and S. Bhatt. Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. *Nature*, 584(7820):257–261, August 2020. ISSN 1476-4687. .
- J.-B.-J. Fourier. *Théorie analytique de la chaleur*. F. Didot, 1822.
- J. Fournet and A. Barrat. Contact Patterns among High School Students. *PLOS ONE*, 9(9):e107878, September 2014. ISSN 1932-6203. .
- F. Frasca, E. Rossi, D. Eynard, B. Chamberlain, M. Bronstein, and F. Monti. SIGN: Scalable Inception Graph Neural Networks, 2020.
- L. C. Freeman. Centrality in social networks conceptual clarification. *Social Networks*, 1(3):215–239, January 1978. ISSN 0378-8733. .
- C. Fritz, E. Dorigatti, and D. Rügamer. Combining graph neural networks and spatio-temporal disease models to improve the prediction of weekly COVID-19 cases in Germany. *Scientific Reports*, 12(1):3930, March 2022. ISSN 2045-2322. .
- H. Fröhlich, J. K. Wegner, F. Sieker, and A. Zell. Optimal assignment kernels for attributed molecular graphs. In *In Proceedings of the 22nd International Conference on Machine Learning*, pages 225–232. ACM Press, 2005.
- S. Fujimoto, H. van Hoof, and D. Meger. Addressing Function Approximation Error in Actor-Critic Methods, October 2018. URL <http://arxiv.org/abs/1802.09477>.
- S. Garg, N. Bhatnagar, and N. Gangadharan. A Case for Participatory Disease Surveillance of the COVID-19 Pandemic in India. *JMIR Public Health and Surveillance*, 6(2):e18795, 2020. .
- M. Garry, L. Hope, R. Zajac, A. J. Verrall, and J. M. Robertson. Contact Tracing: A Memory Task With Consequences for Public Health. *Perspectives on Psychological Science*, 16(1):175–187, January 2021. ISSN 1745-6916. .
- C. Giattino. How epidemiological models of COVID-19 help us estimate the true number of infections, August 2020. URL <https://ourworldindata.org/covid-models>.

- D. T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *The Journal of Physical Chemistry*, 81(25):2340–2361, December 1977. ISSN 0022-3654, 1541-5740. .
- J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl. Neural Message Passing for Quantum Chemistry, June 2017.
- G. Giordano, F. Blanchini, R. Bruno, P. Colaneri, A. Di Filippo, A. Di Matteo, M. Colaneri, and t. C. I. S. M. P. T. Force. A SIDARTHE Model of COVID-19 Epidemic in Italy. *Nature Medicine*, 26(6):855–860, June 2020. ISSN 1078-8956, 1546-170X. .
- J. N. C. Gonçalves, H. S. Rodrigues, and M. T. T. Monteiro. A Contribution of Dynamical Systems Theory and Epidemiological Modeling to a Viral Marketing Campaign. In A. M. Madureira, A. Abraham, D. Gamboa, and P. Novais, editors, *Intelligent Systems Design and Applications*, Advances in Intelligent Systems and Computing, pages 974–983, Cham, 2017. Springer International Publishing. ISBN 978-3-319-53480-0. .
- Google and Apple. Exposure Notifications API — Google API for Exposure Notifications, 2020. URL <https://developers.google.com/android/exposure-notifications/exposure-notifications-api>.
- M. Gori, G. Monfardini, and F. Scarselli. A new model for earning in graph domains. In *Proceedings of the International Joint Conference on Neural Networks*, volume 2, pages 729–734 vol. 2, January 2005. ISBN 978-0-7803-9048-5. .
- GPAW. Global Pandemic App Watch (GPAW): COVID-19 Exposure Notification & Contact Tracing – CRAiEDL, December 2020. URL <https://craiedl.ca/gpaw/>.
- M. Granovetter. Threshold Models of Collective Behavior. *American Journal of Sociology*, 83(6):1420–1443, 1978. ISSN 0002-9602. URL <https://www.jstor.org/stable/2778111>.
- N. C. Grassly, M. Pons-Salort, E. P. K. Parker, P. J. White, N. M. Ferguson, K. Ainslie, M. Baguelin, S. Bhatt, A. Boonyasiri, N. Brazeau, L. Cattarino, H. Coupland, Z. Cucunuba, G. Cuomo-Dannenburg, A. Dighe, C. Donnelly, S. L. van Elsland, R. FitzJohn, S. Flaxman, K. Fraser, K. Gaythorpe, W. Green, A. Hamlet, W. Hinsley, N. Imai, E. Knock, D. Laydon, T. Mellan, S. Mishra, G. Nedjati-Gilani, P. Nouvellet, L. Okell, M. Ragonnet-Cronin, H. A. Thompson, H. J. T. Unwin, M. Vollmer, E. Volz, C. Walters, Y. Wang, O. J. Watson, C. Whittaker, L. Whittles, and X. Xi. Comparison of molecular testing strategies for COVID-19 control: A mathematical modelling study. *The Lancet Infectious Diseases*, 0(0), August 2020. ISSN 1473-3099, 1474-4457. .
- A. Gretton. Introduction to RKHS, and some simple kernel algorithms. Technical report, UCL, The Gatsby Unit, October 2019.

- A. Gretton, O. Bousquet, A. Smola, and B. Schölkopf. Measuring Statistical Dependence with Hilbert-Schmidt Norms. In D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, M. Sudan, D. Terzopoulos, D. Tygar, M. Y. Vardi, G. Weikum, S. Jain, H. U. Simon, and E. Tomita, editors, *Algorithmic Learning Theory*, volume 3734, pages 63–77. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005. ISBN 978-3-540-29242-5 978-3-540-31696-1. .
- S. Gu and Y. Guo. Learning SVM Classifiers with Indefinite Kernels. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, page 7, April 2012.
- N. Gupta, A. Narain, A. Arora, and D. Sharma. Correlating centralities of social networks. In *2016 IEEE International Conference on Advanced Networks and Telecommunications Systems (ANTS)*, pages 1–6, Bangalore, India, November 2016. IEEE Press. .
- T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor, August 2018. URL <http://arxiv.org/abs/1801.01290>.
- A. A. Hagberg, D. A. Schult, and P. J. Swart. Exploring network structure, dynamics, and function using NetworkX. In G. Varoquaux, T. Vaught, and J. Millman, editors, *Proceedings of the 7th Python in Science Conference*, pages 11–15, Pasadena, CA USA, 2008.
- N. Haider, A. Y. Osman, A. Gadzekpo, G. O. Akipede, D. Asogun, R. Ansumana, R. J. Lessells, P. Khan, M. M. A. Hamid, D. Yeboah-Manu, L. Mboera, E. H. Shayo, B. T. Mmbaga, M. Urassa, D. Musoke, N. Kapata, R. A. Ferrand, P.-C. Kapata, F. Stigler, T. Czypionka, A. Zumla, R. Kock, and D. McCoy. Lockdown measures in response to COVID-19 in nine sub-Saharan African countries. *BMJ Global Health*, 5(10): e003319, October 2020. ISSN 2059-7908. .
- E. Hairer, S. P. Nørsett, and G. Wanner. *Solving Ordinary Differential Equations I: Nonstiff Problems*. Springer Series in Computational Mathematics, Springer Ser.Comp.Mathem. Hairer,E.:Solving Ordinary Diff. Springer-Verlag, Berlin Heidelberg, 2 edition, 1993. ISBN 978-3-540-56670-0. .
- W. L. Hamilton, R. Ying, and J. Leskovec. Inductive Representation Learning on Large Graphs. *arXiv:1706.02216 [cs, stat]*, September 2018. URL <http://arxiv.org/abs/1706.02216>.
- L. Harwin and S. P. Comparison of SARSA algorithm and Temporal Difference Learning Algorithm for Robotic Path Planning for Static Obstacles. In *2019 Third International Conference on Inventive Systems and Control (ICISC)*, pages 472–476, January 2019. .

- W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, April 1970. ISSN 0006-3444. .
- D. Haussler. Convolution kernels on discrete structures. *undefined*, 1999. URL <https://www.semanticscholar.org/paper/Convolution-kernels-on-discrete-structures-Haussler/5ee0d8aeb2cb01ef4d8a858d234e72a7400c03ac>.
- J. He, M. Ostendorf, X. He, J. Chen, J. Gao, L. Li, and L. Deng. Deep Reinforcement Learning with a Combinatorial Action Space for Predicting Popular Reddit Threads. In *EMNLP*, July 2019. URL <https://openreview.net/forum?id=Sy4swGGubS>.
- S. He, Y. Peng, and K. Sun. SEIR modeling of the COVID-19 and its dynamics. *Nonlinear Dynamics*, 101(3):1667–1680, August 2020. ISSN 1573-269X. .
- J. Hellewell, S. Abbott, A. Gimma, N. I. Bosse, C. I. Jarvis, T. W. Russell, J. D. Munday, A. J. Kucharski, W. J. Edmunds, S. Funk, R. M. Eggo, F. Sun, S. Flasche, B. J. Quilty, N. Davies, Y. Liu, S. Clifford, P. Klepac, M. Jit, C. Diamond, H. Gibbs, and K. van Zandvoort. Feasibility of controlling COVID-19 outbreaks by isolation of cases and contacts. *The Lancet Global Health*, 8(4):e488–e496, April 2020. ISSN 2214109X. .
- J. Henley. Covid surges across Europe as experts warn not let guard down. *The Guardian*, June 2022. ISSN 0261-3077. URL <https://www.theguardian.com/world/2022/jun/21/covid-surges-europe-ba4-ba5-cases>.
- J. Herby, L. Jonung, and S. H. Hanke. Did lockdowns work? The verdict on Covid restrictions. Technical report, Institute of Economic Affairs, London, 2023. URL <https://iea.org.uk/publications/did-lockdowns-work-the-verdict-on-covid-restrictions>.
- R. Hinch, W. Probert, A. Nurtay, M. Kendall, C. Wymant, M. Hall, K. Lythgoe, A. B. Cruz, L. Zhao, A. Stewart, L. Ferretti, M. Parker, A. Meroueh, B. Mathias, S. Stevenson, D. Montero, J. Warren, N. K. Mather, A. Finkelstein, L. Abeler-Dörner, D. Bonsall, and C. Fraser. Digital contact tracing can slow or even stop coronavirus transmission and ease us out of lockdown, April 2020a. URL <https://www.research.ox.ac.uk/Article/2020-04-16-digital-contact-tracing-can-slow-or-even-stop-coronavirus-transmission-and-ease-us-out-of-lockdown>.
- R. Hinch, W. Probert, A. Nurtay, M. Kendall, C. Wymant, M. Hall, K. Lythgoe, A. B. Cruz, L. Zhao, A. Stewart, L. Ferretti, M. Parker, A. Meroueh, B. Mathias, S. Stevenson, D. Montero, J. Warren, N. K. Mather, A. Finkelstein, D. Bonsall, and C. Fraser. Effective Configurations of a Digital Contact Tracing App: A report to NHSX. Technical report, October 2020b. URL https://cdn.theconversation.com/static_files/files/1009/Report_-_Effective_App_Configurations.pdf?1587531217.

- NT. Hoang and T. Maehara. Revisiting Graph Neural Networks: All We Have is Low-Pass Filters. *arXiv:1905.09550 [cs, math, stat]*, May 2019. URL <http://arxiv.org/abs/1905.09550>.
- S. Hochreiter and J. Schmidhuber. Long Short-term Memory. *Neural computation*, 9: 1735–80, December 1997. .
- P. W. Holland, K. B. Laskey, and S. Leinhardt. Stochastic blockmodels: First steps. *Social Networks*, 5(2):109–137, June 1983. ISSN 0378-8733. .
- P. Holme and B. J. Kim. Growing scale-free networks with tunable clustering. *Physical Review E*, 65(2):026107, January 2002. .
- M. D. Homan and A. Gelman. The No-U-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *The Journal of Machine Learning Research*, 15(1): 1593–1623, January 2014. ISSN 1532-4435.
- J. Hopman and S. Mehtar. Country level analysis of COVID-19 policies. *eClinicalMedicine*, 25, August 2020. ISSN 2589-5370. .
- S. Hsiang, D. Allen, S. Annan-Phan, K. Bell, I. Bolliger, T. Chong, H. Druckenmiller, L. Y. Huang, A. Hultgren, E. Krasovich, P. Lau, J. Lee, E. Rolf, J. Tseng, and T. Wu. The effect of large-scale anti-contagion policies on the COVID-19 pandemic. *Nature*, 584(7820):262–267, August 2020. ISSN 1476-4687. .
- W. Hu, M. Fey, M. Zitnik, Y. Dong, H. Ren, B. Liu, M. Catasta, and J. Leskovec. Open Graph Benchmark: Datasets for Machine Learning on Graphs. In *Advances in Neural Information Processing Systems*, volume 33, pages 22118–22133. Curran Associates, Inc., 2020. URL <https://papers.neurips.cc/paper/2020/hash/fb60d411a5c5b72b2e7d3527cfc84fd0-Abstract.html>.
- Q. Huang, M. Yamada, Y. Tian, D. Singh, and Y. Chang. GraphLIME: Local Interpretable Model Explanations for Graph Neural Networks. *IEEE Transactions on Knowledge and Data Engineering*, pages 1–6, 2022. ISSN 1558-2191. .
- R. Huerta and L. S. Tsimring. Contact tracing and epidemics control in social networks. *Physical Review E*, 66(5):056115, November 2002. .
- A. Hyafil and D. Morriña. Analysis of the impact of lockdown on the reproduction number of the SARS-Cov-2 in Spain. *Gaceta Sanitaria*, 35(5):453–458, 2021. ISSN 0213-9111. .
- C. Ilin, S. Annan-Phan, X. H. Tai, S. Mehra, S. Hsiang, and J. E. Blumenstock. Public mobility data enables COVID-19 forecasting and management at local and global scales. *Scientific Reports*, 11(1):13531, December 2021. ISSN 2045-2322. .

- IOM. IOM COVID 19 Response - Situation Report 21 (26 June 2020). Technical Report 21, International Organization for Migration, June 2020. URL <https://dtm.iom.int/reports/iom-covid-19-response-situation-report-21-26-june-2020>.
- L. Isella, J. Stehlé, A. Barrat, C. Cattuto, J.-F. Pinton, and W. V. den Broeck. What's in a crowd? Analysis of face-to-face behavioral networks. *Journal of Theoretical Biology*, 271(1):166–180, February 2011. ISSN 00225193. .
- N. Islam, S. J. Sharp, G. Chowell, S. Shabnam, I. Kawachi, B. Lacey, J. M. Massaro, R. B. D'Agostino, and M. White. Physical distancing interventions and incidence of coronavirus disease 2019: Natural experiment in 149 countries. *BMJ*, 370:m2743, July 2020. ISSN 1756-1833. .
- C. Jacob. Branching Processes: Their Role in Epidemiology. *International Journal of Environmental Research and Public Health*, 7(3):1186–1204, March 2010. ISSN 1661-7827. .
- A. Jain, N. Kosaka, K.-M. Kim, and J. J. Lim. Know Your Action Set: Learning Action Relations for Reinforcement Learning. In *International Conference on Learning Representations*, March 2022. URL <https://openreview.net/forum?id=MljXVdp4A3N>.
- S. Janson, A. Rucinski, and T. Luczak. *Random Graphs*. John Wiley & Sons, September 2011. ISBN 978-1-118-03096-7.
- T. Jefferson, L. Dooley, E. Ferroni, L. A. Al-Ansary, M. L. van Driel, G. A. Bawazeer, M. A. Jones, T. C. Hoffmann, J. Clark, E. M. Beller, P. P. Glasziou, and J. M. Conly. Physical interventions to interrupt or reduce the spread of respiratory viruses. *The Cochrane Database of Systematic Reviews*, 1(1):CD006207, January 2023. ISSN 1469-493X. .
- B. Jhun. Effective vaccination strategy using graph neural network ansatz, December 2021.
- F. Jin, E. Dougherty, P. Saraf, Y. Cao, and N. Ramakrishnan. Epidemiological Modeling of News and Rumors on Twitter. In *Proceedings of the 7th Workshop on Social Network Mining and Analysis*, August 2013. .
- M. Jing, K. Y. Ng, B. M. Namee, P. Biglarbeigi, R. Brisk, R. Bond, D. Finlay, and J. McLaughlin. COVID-19 modelling by time-varying transmission rate associated with mobility trend of driving via Apple Maps. *Journal of Biomedical Informatics*, 122:103905, October 2021. ISSN 1532-0464. .
- A. R. Joffe. COVID-19: Rethinking the Lockdown Groupthink. *Frontiers in Public Health*, 9, 2021. ISSN 2296-2565. URL <https://www.frontiersin.org/articles/10.3389/fpubh.2021.625778>.

- F. D. Johansson and D. Dubhashi. Learning with Similarity Functions on Graphs using Matchings of Geometric Embeddings. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 467–476, Sydney NSW Australia, August 2015. ACM. ISBN 978-1-4503-3664-2. .
- C. K. Joshi, T. Laurent, and X. Bresson. An Efficient Graph Convolutional Network Technique for the Travelling Salesman Problem, 2019.
- L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1-2):99–134, May 1998. ISSN 00043702. .
- Y. Kajitani and M. Hatayama. Explaining the effective reproduction number of COVID-19 through mobility and enterprise statistics: Evidence from the first wave in Japan. *PLOS ONE*, 16(3):e0247186, March 2021. ISSN 1932-6203. .
- A. Kapoor, X. Ben, L. Liu, B. Perozzi, M. Barnes, M. Blais, and S. O’Banion. Examining COVID-19 Forecasting using Spatio-Temporal Graph Neural Networks. *San Diego*, 2020.
- S. Kapturowski, G. Ostrovski, J. Quan, R. Munos, and W. Dabney. Recurrent Experience Replay in Distributed Reinforcement Learning. In *International Conference on Learning Representations*, February 2022. URL <https://openreview.net/forum?id=r1lyTjAqYX>.
- A. Karaivanov, S. E. Lu, H. Shigeoka, C. Chen, and S. Pamplona. Face masks, public policies and slowing the spread of COVID-19: Evidence from Canada. *Journal of Health Economics*, 78:102475, July 2021. ISSN 1879-1646. .
- L. Katz. A new status index derived from sociometric analysis. *Psychometrika*, 18(1): 39–43, March 1953. ISSN 1860-0980. .
- M. J. Keeling and K. T. Eames. Networks and epidemic models. *Journal of The Royal Society Interface*, 2(4):295–307, September 2005. .
- M. J. Keeling, S. Moore, B. S. Penman, and E. M. Hill. The impacts of SARS-CoV-2 vaccine dose separation and targeting on the COVID-19 epidemic in England. *Nature Communications*, 14(1):1–10, February 2023. ISSN 2041-1723. .
- J. Keller, S. E. Liyanage, M. Hingorani, and A. Hingorani. Probability of encountering Covid-19 patients based on prevalence and testing during resumption of ophthalmology services. *Eye*, July:1–2, July 2020. ISSN 1476-5454. .
- D. Kempe, J. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD ’03, pages 137–146, New York, NY, USA, August 2003. Association for Computing Machinery. ISBN 978-1-58113-737-8. .

- W. O. Kermack, A. G. McKendrick, and G. T. Walker. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 115(772):700–721, August 1927. .
- K. Kersting, N. M. Kriege, C. Morris, P. Mutzel, and M. Neumann. Benchmark data sets for graph kernels, 2016. URL <http://graphkernels.cs.tu-dortmund.de>.
- S. Khan, M. Naseer, M. Hayat, S. W. Zamir, F. S. Khan, and M. Shah. Transformers in Vision: A Survey. *ACM Computing Surveys*, page 3505244, January 2022. ISSN 0360-0300, 1557-7341. .
- M. Kimura, K. Saito, R. Nakano, and H. Motoda. Extracting influential nodes on a social network for information diffusion. *Data Mining and Knowledge Discovery*, 20(1):70, October 2009. ISSN 1573-756X. .
- D. P. Kingma and J. Ba. Adam: A Method for Stochastic Optimization, January 2017.
- T. N. Kipf and M. Welling. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=SJU4ayYgl>.
- B. R. Kiran, I. Sobh, V. Talpaert, P. Mannion, A. A. A. Sallab, S. Yogamani, and P. Pérez. Deep Reinforcement Learning for Autonomous Driving: A Survey, January 2021. URL <http://arxiv.org/abs/2002.00444>.
- A. G. Klein. How to fight Holocaust denial in social media – with the evidence of what really happened, December 2020. URL <http://theconversation.com/how-to-fight-holocaust-denial-in-social-media-with-the-evidence-of-what-really-happened-150719>.
- L. Knox, G. C. Karantzas, D. Romano, J. A. Feeney, and J. A. Simpson. One year on: What we have learned about the psychological effects of COVID-19 social restrictions: A meta-analysis. *Current Opinion in Psychology*, 46:101315, August 2022. ISSN 2352-250X. .
- T. Kobayashi. Adaptive and Multiple Time-scale Eligibility Traces for Online Deep Reinforcement Learning. *Robotics and Autonomous Systems*, 151:104019, May 2022. ISSN 09218890. .
- S. Kojaku, L. Hébert-Dufresne, E. Mones, S. Lehmann, and Y.-Y. Ahn. The effectiveness of backward contact tracing in networks. *Nature Physics*, 17(5):652–658, May 2021. ISSN 1745-2481. .
- W. R. Kolk and R. A. Lerman. Analytic Solutions to Nonlinear Differential Equations. In W. R. Kolk and R. A. Lerman, editors, *Nonlinear System Dynamics*, pages 23–60. Springer US, Boston, MA, 1992. ISBN 978-1-4684-6494-8. .

- V. Konda and J. Tsitsiklis. Actor-Critic Algorithms. In *Advances in Neural Information Processing Systems*, volume 12. MIT Press, 1999. URL <https://proceedings.neurips.cc/paper/1999/hash/6449f44a102fde848669bdd9eb6b76fa-Abstract.html>.
- R. Kondor and J. Lafferty. Diffusion Kernels on Graphs and Other Discrete Input Spaces. *ICML*, Vol. 2, May 2002.
- W. Kool, H. van Hoof, and M. Welling. Attention, Learn to Solve Routing Problems! In *International Conference on Learning Representations*, December 2018. URL <https://openreview.net/forum?id=ByxBFsRqYm>.
- N. M. Kriege, P.-L. Giscard, and R. Wilson. On Valid Optimal Assignment Kernels and Applications to Graph Classification. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. URL <https://papers.nips.cc/paper/2016/hash/0efe32849d230d7f53049ddc4a4b0c60-Abstract.html>.
- N. M. Kriege, M. Neumann, C. Morris, K. Kersting, and P. Mutzel. A unifying view of explicit and implicit feature maps of graph kernels. *Data Mining and Knowledge Discovery*, 33(6):1505–1547, November 2019. ISSN 1573-756X. .
- N. M. Kriege, F. D. Johansson, and C. Morris. A survey on graph kernels. *Applied Network Science*, 5(1):1–42, December 2020. ISSN 2364-8228. .
- M. Kröger and R. Schlickeiser. Analytical solution of the SIR-model for the temporal evolution of epidemics. Part A: Time-independent reproduction factor. *Journal of Physics A: Mathematical and Theoretical*, 53(50):505601, November 2020. ISSN 1751-8121. .
- J. B. Kruskal. On the Shortest Spanning Subtree of a Graph and the Traveling Salesman Problem. *Proceedings of the American Mathematical Society*, 7(1):48–50, 1956. ISSN 0002-9939. .
- O. Krylova and D. J. D. Earn. Effects of the infectious period distribution on predicted transitions in childhood disease dynamics. *Journal of the Royal Society Interface*, 10(84), July 2013. ISSN 1742-5689. .
- A. J. Kucharski, T. W. Russell, C. Diamond, Y. Liu, J. Edmunds, S. Funk, R. M. Eggo, and Centre for Mathematical Modelling of Infectious Diseases COVID-19 working group. Early dynamics of transmission and control of COVID-19: A mathematical modelling study. *The Lancet. Infectious Diseases*, 20(5):553–558, May 2020. ISSN 1474-4457. .
- P. Kumar and A. Sinha. Information diffusion modeling and analysis for socially interacting networks. *Social Network Analysis and Mining*, 11(1):11, 2021. ISSN 1869-5450. .

- R. Kumar, J. Novak, and A. Tomkins. Structure and evolution of online social networks. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '06*, pages 611–617, New York, NY, USA, August 2006. Association for Computing Machinery. ISBN 978-1-59593-339-3. .
- S. Kumar, M. Saini, M. Goel, and N. Aggarwal. Modeling Information Diffusion In Online Social Networks Using SEI Epidemic Model. *Procedia Computer Science*, 171: 672–678, January 2020. ISSN 1877-0509. .
- J. Kunegis. KONECT: The Koblenz network collection. In *Proceedings of the 22nd International Conference on World Wide Web, WWW '13 Companion*, pages 1343–1350, New York, NY, USA, May 2013. Association for Computing Machinery. ISBN 978-1-4503-2038-2. .
- O. Lange and L. Perez. Traffic prediction with advanced Graph Neural Networks, September 2020. URL <https://deepmind.com/blog/article/traffic-prediction-with-advanced-graph-neural-networks>. .
- A. A. Lashari and P. Trapman. Branching process approach for epidemics in dynamic partnership network. *Journal of Mathematical Biology*, 76(1):265–294, 2018. ISSN 0303-6812. .
- A. Lazaridis, A. Fachantidis, and I. Vlahavas. Deep Reinforcement Learning: A State-of-the-Art Walkthrough. *Journal of Artificial Intelligence Research*, 69:1421–1471, December 2020. ISSN 1076-9757. .
- J. Leskovec. Stanford CS224W: Graph Neural Networks, February 2021. URL <http://web.stanford.edu/class/cs224w>. .
- K. Leung and J. T. Wu. Managing waning vaccine protection against SARS-CoV-2 variants. *The Lancet*, 399(10319):2–3, January 2022. ISSN 0140-6736, 1474-547X. .
- D. Lewis. Why many countries failed at COVID contact-tracing — but some got it right. *Nature*, 588(7838):384–387, December 2020. .
- D. Lewis. Superspreading drives the COVID pandemic — and could help to tame it. *Nature*, 590(7847):544–546, February 2021. .
- Y. Li, M. Liang, L. Gao, M. Ayaz Ahmed, J. P. Uy, C. Cheng, Q. Zhou, and C. Sun. Face masks to prevent transmission of COVID-19: A systematic review and meta-analysis. *American Journal of Infection Control*, 49(7):900–906, July 2021. ISSN 0196-6553. .
- Y. Li, H. V. Zhao, and Y. Chen. An epidemic model for correlated information diffusion in crowd intelligence networks. *International Journal of Crowd Science*, 3(2): 168–183, January 2019. ISSN 2398-7294. .

- 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, July 2019. URL <http://arxiv.org/abs/1509.02971>.
- M. Lipsitch, T. Cohen, B. Cooper, J. M. Robins, S. Ma, L. James, G. Gopalakrishna, S. K. Chew, C. C. Tan, M. H. Samore, D. Fisman, and M. Murray. Transmission Dynamics and Control of Severe Acute Respiratory Syndrome. *Science (New York, N.Y.)*, 300(5627):1966–1970, June 2003. ISSN 0036-8075. .
- D. Liu, Y. Jing, J. Zhao, W. Wang, and G. Song. A Fast and Efficient Algorithm for Mining Top-k Nodes in Complex Networks. *Scientific Reports*, 7(1):43330, February 2017. ISSN 2045-2322. .
- L. Liu, S. Vikram, J. Lao, X. Ben, A. N. D’Amour, S. O’Banion, M. Sandler, R. A. Saurous, and M. D. Hoffman. Estimating the Changing Infection Rate of COVID-19 Using Bayesian Models of Mobility. *medRxiv*, <https://www.medrxiv.org/content/10.1101/2020.08.06.20169664v1.full>, 2020a. URL <https://www.medrxiv.org/content/10.1101/2020.08.06.20169664v1>.
- T. Liu, D. Gong, J. Xiao, J. Hu, G. He, Z. Rong, and W. Ma. Cluster infections play important roles in the rapid evolution of COVID-19 transmission: A systematic review. *International Journal of Infectious Diseases*, 99:374–380, October 2020b. ISSN 1201-9712. .
- Y. Liu, S.-M. Diao, Y.-X. Zhu, and Q. Liu. SHIR competitive information diffusion model for online social media. *Physica A: Statistical Mechanics and its Applications*, 461:543–553, November 2016. ISSN 0378-4371. .
- J. O. Lloyd-Smith, S. J. Schreiber, P. E. Kopp, and W. M. Getz. Superspreading and the effect of individual variation on disease emergence. *Nature*, 438(7066):355–359, November 2005. ISSN 1476-4687. .
- J. Lopez Bernal, N. Andrews, C. Gower, E. Gallagher, R. Simmons, S. Thelwall, J. Stowe, E. Tessier, N. Groves, G. Dabrera, R. Myers, C. N. Campbell, G. Amirthalingam, M. Edmunds, M. Zambon, K. E. Brown, S. Hopkins, M. Chand, and M. Ramsay. Effectiveness of Covid-19 Vaccines against the B.1.617.2 (Delta) Variant. *New England Journal of Medicine*, 385(7):585–594, August 2021. ISSN 0028-4793. .
- I. Loshchilov and F. Hutter. Decoupled Weight Decay Regularization. *arXiv:1711.05101 [cs, math]*, January 2019. URL <http://arxiv.org/abs/1711.05101>.
- M. A. Lozano, Ò. G. i. Orts, E. Piñol, M. Rebollo, K. Polotskaya, M. A. Garcia-March, J. A. Conejero, F. Escolano, and N. Oliver. Open Data Science to Fight COVID-19: Winning the 500k XPRIZE Pandemic Response Challenge. In Y. Dong, N. Kourtellis, B. Hammer, and J. A. Lozano, editors, *Machine Learning and Knowledge Discovery in*

- Databases. Applied Data Science Track*, Lecture Notes in Computer Science, pages 384–399, Cham, 2021. Springer International Publishing. ISBN 978-3-030-86514-6. .
- B. Lucas, B. Vahedi, and M. Karimzadeh. A spatiotemporal machine learning approach to forecasting COVID-19 incidence at the county level in the USA. *International Journal of Data Science and Analytics*, 15(3):247–266, April 2023. ISSN 2364-4168. .
- S. Lundberg and S.-I. Lee. A Unified Approach to Interpreting Model Predictions, November 2017.
- R. Luss and A. d’Aspremont. Support Vector Machine Classification with Indefinite Kernels. In *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc., 2008. URL <https://proceedings.neurips.cc/paper/2007/hash/c0c7c76d30bd3dcaefc96f40275bdc0a-Abstract.html>.
- J. Ma. Estimating epidemic exponential growth rate and basic reproduction number. *Infectious Disease Modelling*, 5:129–141, January 2020. ISSN 2468-0427. .
- A. Madan, M. Cebrian, S. Moturu, K. Farrahi, and A. S. Pentland. Sensing the ‘Health State’ of a Community. *IEEE Pervasive Computing*, 11(4):36–45, October 2012. ISSN 1558-2590. .
- M. Mancastropa, C. Castellano, A. Vezzani, and R. Burioni. Stochastic sampling effects favor manual over digital contact tracing. *Nature Communications*, 12(1):1919, March 2021. ISSN 2041-1723. .
- H. Maron, H. Ben-Hamu, N. Shamir, and Y. Lipman. Invariant and Equivariant Graph Networks. *arXiv:1812.09902 [cs, stat]*, April 2019. URL <http://arxiv.org/abs/1812.09902>.
- M. Martinez-Garcia, E. Sansano-Sansano, A. Castillo-Hornero, R. Femenia, K. Roomp, and N. Oliver. Social isolation during the COVID-19 pandemic in Spain: A population study, January 2022.
- R. Mason, A. Allegretti, H. Devlin, and I. Sample. UK Treasury pushes to end most free Covid testing despite experts’ warnings. *The Guardian*, February 2022. ISSN 0261-3077. URL <https://www.theguardian.com/world/2022/feb/11/uk-treasury-pushes-to-end-most-free-covid-testing-despite-experts-warnings>.
- R. Mastrandrea, J. Fournet, and A. Barrat. Contact Patterns in a High School: A Comparison between Data Collected Using Wearable Sensors, Contact Diaries and Friendship Surveys. *PLOS ONE*, 10(9):e0136497, September 2015. ISSN 1932-6203. .
- N. Masuda. Immunization of networks with community structure. *New Journal of Physics*, 11(12):123018, December 2009. ISSN 1367-2630. .

- L. Matrajt and T. Leung. Evaluating the Effectiveness of Social Distancing Interventions to Delay or Flatten the Epidemic Curve of Coronavirus Disease. *Emerging Infectious Diseases*, 26(8):1740–1748, August 2020. ISSN 1080-6040. .
- Y. Matsubara, Y. Sakurai, B. A. Prakash, L. Li, and C. Faloutsos. Rise and fall patterns of information diffusion: Model and implications. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '12*, pages 6–14, New York, NY, USA, August 2012. Association for Computing Machinery. ISBN 978-1-4503-1462-6. .
- J. McDonald. What the Cochrane Review Says About Masks For COVID-19 - and What It Doesn't, March 2023. URL <https://www.factcheck.org/2023/03/scicheck-what-the-cochrane-review-says-about-masks-for-covid-19-and-what-it-doesnt/>.
- D. McEvoy, C. McAloon, A. Collins, K. Hunt, F. Butler, A. Byrne, M. Casey-Bryars, A. Barber, J. Griffin, E. A. Lane, P. Wall, and S. J. More. Relative infectiousness of asymptomatic SARS-CoV-2 infected persons compared with symptomatic individuals: A rapid scoping review. *BMJ open*, 11(5):e042354, May 2021. ISSN 2044-6055. .
- J. Mei, C. Xiao, B. Dai, L. Li, C. Szepesvari, and D. Schuurmans. Escaping the Gravitational Pull of Softmax. In *Advances in Neural Information Processing Systems*, volume 33, pages 21130–21140. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/f1cf2a082126bf02de0b307778ce73a7-Abstract.html>.
- E. Meiriom, H. Maron, S. Mannor, and G. Chechik. Controlling Graph Dynamics with Reinforcement Learning and Graph Neural Networks. In *Proceedings of the 38th International Conference on Machine Learning*, pages 7565–7577. PMLR, July 2021. .
- E. A. Meiriom, C. Milling, C. Caramanis, S. Mannor, S. Shakkottai, and A. Orda. Localized Epidemic Detection in Networks with Overwhelming Noise. *ACM SIGMETRICS Performance Evaluation Review*, 43(1):441–442, June 2015. ISSN 0163-5999. .
- E. A. Meiriom, C. Caramanis, S. Mannor, A. Orda, and S. Shakkottai. Detecting Cascades from Weak Signatures. *IEEE Transactions on Network Science and Engineering*, 5(4):313–325, October 2018. ISSN 2327-4697. .
- T. R. Mercer and M. Salit. Testing at scale during the COVID-19 pandemic. *Nature Reviews Genetics*, 22(7):415–426, July 2021. ISSN 1471-0064. .
- J. C. Miller and J. M. Hyman. Effective vaccination strategies for realistic social networks. *Physica A: Statistical Mechanics and its Applications*, 386(2):780–785, December 2007. ISSN 0378-4371. .

- S. Mishra, J. A. Scott, D. J. Laydon, S. Flaxman, A. Gandy, T. A. Mellan, H. J. T. Unwin, M. Vollmer, H. Coupland, O. Ratmann, M. Monod, H. H. Zhu, A. Cori, K. A. M. Gaythorpe, L. K. Whittles, C. Whittaker, C. A. Donnelly, N. M. Ferguson, and S. Bhatt. Comparing the responses of the UK, Sweden and Denmark to COVID-19 using counterfactual modelling. *Scientific Reports*, 11(1):16342, August 2021. ISSN 2045-2322. .
- T. Mitze, R. Kosfeld, J. Rode, and K. Wälde. Face masks considerably reduce COVID-19 cases in Germany. *Proceedings of the National Academy of Sciences*, 117(51):32293–32301, December 2020. .
- M. Mitzenmacher and E. Upfal. *Probability and Computing: An Introduction to Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, New York, 2005. ISBN 978-0-521-83540-4.
- K. Mizumoto, K. Kagaya, A. Zarebski, and G. Chowell. Estimating the asymptomatic proportion of coronavirus disease 2019 (COVID-19) cases on board the Diamond Princess cruise ship, Yokohama, Japan, 2020. *Eurosurveillance*, 25(10), March 2020. ISSN 1025-496X. .
- V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller. Playing Atari with Deep Reinforcement Learning, December 2013.
- V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, February 2015. ISSN 1476-4687. .
- S. Moein, N. Nickaeen, A. Roointan, N. Borhani, Z. Heidary, S. H. Javanmard, J. Ghaisari, and Y. Gheisari. Inefficiency of SIR models in forecasting COVID-19 epidemic: A case study of Isfahan. *Scientific Reports*, 11(1):4725, February 2021. ISSN 2045-2322. .
- M. Mohri, A. Rostamizadeh, and A. Talwalkar. *Foundations of Machine Learning*. Adaptive Computation and Machine Learning Series. MIT Press, Cambridge, MA, USA, 2 edition, December 2018. ISBN 978-0-262-03940-6.
- J. S. More and C. Lingam. A SI model for social media influencer maximization. *Applied Computing and Informatics*, 15(2):102–108, July 2019. ISSN 2210-8327. .
- C. Morris, M. Ritzert, M. Fey, W. L. Hamilton, J. E. Lenssen, G. Rattan, and M. Grohe. Weisfeiler and Leman Go Neural: Higher-order Graph Neural Networks. *arXiv:1810.02244 [cs, stat]*, February 2020. URL <http://arxiv.org/abs/1810.02244>.
- N. Moshiri. The dual-Barabási-Albert model, October 2018. URL <http://arxiv.org/abs/1810.10538>.

- T. Murata and H. Koga. Extended methods for influence maximization in dynamic networks. *Computational Social Networks*, 5(1):8, October 2018. ISSN 2197-4314. .
- C. J. Murray. Forecasting COVID-19 impact on hospital bed-days, ICU-days, ventilator-days and deaths by US state in the next 4 months, March 2020.
- E. Ç. Mutlu, A. Rajabi, and I. Garibay. CD-SEIZ: Cognition-Driven SEIZ Compartmental Model for the Prediction of Information Cascades on Twitter. In Z. Yang and E. von Briesen, editors, *Proceedings of the 2020 Conference of The Computational Social Science Society of the Americas*, Springer Proceedings in Complexity, pages 143–154, Cham, 2021. Springer International Publishing. ISBN 978-3-030-83418-0. .
- M. Nadini, L. Zino, A. Rizzo, and M. Porfiri. A multi-agent model to study epidemic spreading and vaccination strategies in an urban-like environment. *Applied Network Science*, 5(1):1–30, December 2020. ISSN 2364-8228. .
- V. Nair and G. E. Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on International Conference on Machine Learning, ICML'10*, pages 807–814, Madison, WI, USA, June 2010. Omnipress. ISBN 978-1-60558-907-7.
- J. T. Nardini, R. E. Baker, M. J. Simpson, and K. B. Flores. Learning differential equation models from stochastic agent-based model simulations. *Journal of The Royal Society Interface*, 18(176):20200987, March 2021. .
- neptune.ai. Neptune: Experiment management and collaboration tool, 2020. URL <https://neptune.ai>.
- M. Neumann, R. Garnett, C. Bauckhage, and K. Kersting. Propagation kernels: Efficient graph kernels from propagated information. *Machine Learning*, 102(2): 209–245, February 2016. ISSN 1573-0565. .
- M. E. J. Newman. The Structure and Function of Complex Networks. *SIAM Review*, 45(2):167–256, January 2003. ISSN 0036-1445, 1095-7200. .
- M. E. J. Newman and J. Park. Why social networks are different from other types of networks. *Physical Review E*, 68(3):036122, September 2003. .
- M. E. J. Newman, S. H. Strogatz, and D. J. Watts. Random graphs with arbitrary degree distributions and their applications. *Physical Review E*, 64(2):026118, July 2001. .
- M. E. J. Newman, D. J. Watts, and S. H. Strogatz. Random graph models of social networks. *Proceedings of the National Academy of Sciences*, 99(suppl 1):2566–2572, February 2002. ISSN 0027-8424, 1091-6490. .

- M. Newman. *Networks: An Introduction*. Oxford University Press, March 2010. ISBN 978-0-19-159417-5. URL <https://oxford.universitypressscholarship.com/view/10.1093/acprof:oso/9780199206650.001.0001/acprof-9780199206650>.
- L. Oettershagen, N. M. Kriege, C. Morris, and P. Mutzel. Temporal Graph Kernels for Classifying Dissemination Processes. *arXiv:1911.05496 [cs, stat]*, August 2021. URL <http://arxiv.org/abs/1911.05496>.
- Y. Okabe and A. Shudo. A Mathematical Model of Epidemics—A Tutorial for Students. *Mathematics*, 8(7):1174, July 2020. ISSN 2227-7390. .
- K. Oono and T. Suzuki. Graph Neural Networks Exponentially Lose Expressive Power for Node Classification. In *International Conference on Learning Representations*, December 2019. URL <https://openreview.net/forum?id=S1ld02EFPr>.
- D. P. Oran and E. J. Topol. Prevalence of Asymptomatic SARS-CoV-2 Infection. *Annals of Internal Medicine*, 173(5):362–367, June 2020. ISSN 0003-4819. .
- J. F. Padgett and C. K. Ansell. Robust Action and the Rise of the Medici, 1400-1434. *American Journal of Sociology*, 98(6):1259–1319, 1993. ISSN 0002-9602. URL <https://www.jstor.org/stable/2781822>.
- G. Panagopoulos, G. Nikolentzos, and M. Vazirgiannis. Transfer Graph Neural Networks for Pandemic Forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(6):4838–4845, May 2021. ISSN 2374-3468, 2159-5399. .
- J. A. Pandit, J. M. Radin, G. Quer, and E. J. Topol. Smartphone apps in the COVID-19 pandemic. *Nature Biotechnology*, 40(7):1013–1022, July 2022. ISSN 1546-1696. .
- R. Pastor-Satorras and A. Vespignani. Epidemic Spreading in Scale-Free Networks. *Physical Review Letters*, 86(14):3200–3203, April 2001. .
- R. Pastor-Satorras, C. Castellano, P. Van Mieghem, and A. Vespignani. Epidemic processes in complex networks. *Reviews of Modern Physics*, 87(3):925–979, August 2015. ISSN 0034-6861, 1539-0756. .
- A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimeshin, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://papers.nips.cc/paper/2019/hash/bdbca288fee7f92f2bfa9f7012727740-Abstract.html>.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau,

- M. Brucher, M. Perrot, and É. Duchesnay. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12(85):2825–2830, 2011. URL <http://jmlr.org/papers/v12/pedregosa11a.html>.
- S. Pei, S. Kandula, J. Cascante Vega, W. Yang, S. Foerster, C. Thompson, J. Baumgartner, S. D. Ahuja, K. Blaney, J. K. Varma, T. Long, and J. Shaman. Contact tracing reveals community transmission of COVID-19 in New York City. *Nature Communications*, 13(1):6307, October 2022. ISSN 2041-1723. .
- H. Peng, F. Long, and C. Ding. Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(8):1226–1238, August 2005. ISSN 1939-3539. .
- R. Picheta. Why the world is still arguing over face masks, 20 months into the pandemic, November 2021. URL <https://www.cnn.com/2021/11/02/health/face-mask-debate-covid-19-pandemic-cmd-intl/index.html>.
- M. J. Plank, A. James, A. Lustig, N. Steyn, R. N. Binny, and S. C. Hendy. Potential reduction in transmission of COVID-19 by digital contact tracing systems. Preprint doi:10.1101/2020.08.27.20068346, *Infectious Diseases (except HIV/AIDS)*, September 2020.
- F. P. Polack, S. J. Thomas, N. Kitchin, J. Absalon, A. Gurtman, S. Lockhart, J. L. Perez, G. Pérez Marc, E. D. Moreira, C. Zerbini, R. Bailey, K. A. Swanson, S. Roychoudhury, K. Koury, P. Li, W. V. Kalina, D. Cooper, R. W. Frenck, L. L. Hammitt, Ö. Türeci, H. Nell, A. Schaefer, S. Ünal, D. B. Tresnan, S. Mather, P. R. Dormitzer, U. Şahin, K. U. Jansen, and W. C. Gruber. Safety and Efficacy of the BNT162b2 mRNA Covid-19 Vaccine. *New England Journal of Medicine*, 383(27):2603–2615, 2020. ISSN 0028-4793. .
- C. Poletto, S. Meloni, V. Colizza, Y. Moreno, and A. Vespignani. Host Mobility Drives Pathogen Competition in Spatially Structured Populations. *PLOS Computational Biology*, 9(8):e1003169, August 2013. ISSN 1553-7358. .
- R. Prabakaran, S. Jemimah, P. Rawat, D. Sharma, and M. M. Gromiha. A novel hybrid SEIQR model incorporating the effect of quarantine and lockdown regulations for COVID-19. *Scientific Reports*, 11(1):24073, December 2021. ISSN 2045-2322. .
- B. A. Prakash, J. Vreeken, and C. Faloutsos. Spotting Culprits in Epidemics: How Many and Which Ones? In *2012 IEEE 12th International Conference on Data Mining*, pages 11–20, December 2012. .
- V. M. Preciado, M. Zargham, C. Enyioha, A. Jadbabaie, and G. J. Pappas. Optimal Resource Allocation for Network Protection Against Spreading Processes. *IEEE Transactions on Control of Network Systems*, 1(1):99–108, March 2014. ISSN 2325-5870. .

- W. H. Press and S. A. Teukolsky. Adaptive Stepsize Runge-Kutta Integration. *Computers in Physics*, 6(2):188, 1992. ISSN 08941866. .
- R. C. Prim. Shortest Connection Networks And Some Generalizations. *Bell System Technical Journal*, 36(6):1389–1401, 1957. ISSN 1538-7305. .
- N. Przulj, D. G. Corneil, and I. Jurisica. Modeling interactome: Scale-free or geometric? *Bioinformatics*, 20(18):3508–3515, December 2004. ISSN 1367-4803, 1460-2059. .
- M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., USA, 1st edition, 1994. ISBN 978-0-471-61977-2.
- F. Rahmadani and H. Lee. Hybrid Deep Learning-Based Epidemic Prediction Framework of COVID-19: South Korea Case. *Applied Sciences*, 10(23):8539, January 2020. .
- S. F. Railsback and V. Grimm. *Agent-Based and Individual-Based Modeling: A Practical Introduction*. Princeton University Press, 2011. ISBN 978-0-691-13674-5. URL <https://www.jstor.org/stable/j.ctt7sns7>.
- C. E. Rasmussen. Gaussian Processes in Machine Learning. In O. Bousquet, U. von Luxburg, and G. Rätsch, editors, *Advanced Lectures on Machine Learning: ML Summer Schools 2003, Canberra, Australia, February 2 - 14, 2003, Tübingen, Germany, August 4 - 16, 2003, Revised Lectures*, Lecture Notes in Computer Science, pages 63–71. Springer, Berlin, Heidelberg, 2004. ISBN 978-3-540-28650-9. .
- D. C. Rayner, N. R. Sturtevant, and M. Bowling. Subset Selection of Search Heuristics. In *IJCAI*, July 2019. URL <https://openreview.net/forum?id=HJZVf4f0bB>.
- A. Reinhart, L. Brooks, M. Jahja, A. Rumack, J. Tang, S. Agrawal, W. Al Saeed, T. Arnold, A. Basu, J. Bien, Á. A. Cabrera, A. Chin, E. J. Chua, B. Clark, S. Colquhoun, N. DeFries, D. C. Farrow, J. Forlizzi, J. Grabman, S. Gratzl, A. Green, G. Haff, R. Han, K. Harwood, A. J. Hu, R. Hyde, S. Hyun, A. Joshi, J. Kim, A. Kuznetsov, W. La Motte-Kerr, Y. J. Lee, K. Lee, Z. C. Lipton, M. X. Liu, L. Mackey, K. Mazaitis, D. J. McDonald, P. McGuinness, B. Narasimhan, M. P. O'Brien, N. L. Oliveira, P. Patil, A. Perer, C. A. Politsch, S. Rajanala, D. Rucker, C. Scott, N. H. Shah, V. Shankar, J. Sharpnack, D. Shemetov, N. Simon, B. Y. Smith, V. Srivastava, S. Tan, R. Tibshirani, E. Tuzhilina, A. K. Van Nortwick, V. Ventura, L. Wasserman, B. Weaver, J. C. Weiss, S. Whitman, K. Williams, R. Rosenfeld, and R. J. Tibshirani. An open repository of real-time COVID-19 indicators. *Proceedings of the National Academy of Sciences*, 118(51):e2111452118, December 2021. .
- M. T. Ribeiro, S. Singh, and C. Guestrin. ‘Why Should I Trust You?’: Explaining the Predictions of Any Classifier, August 2016.

- S. Riley and N. M. Ferguson. Smallpox transmission and control: Spatial dynamics in Great Britain. *Proceedings of the National Academy of Sciences*, 103(33):12637–12642, August 2006. .
- A. Rimmer. Sixty seconds on . . . the pingdemic. *BMJ*, 374:n1822, July 2021. ISSN 1756-1833. .
- N. Rolander. Face-Mask Skepticism Fades in Sweden as Infection Rates Rise. *Bloomberg.com*, February 2021. URL <https://www.bloomberg.com/news/articles/2021-02-23/swedish-face-mask-skepticism-fades-as-infection-rates-rise>.
- M. Roskams. Population and household estimates, England and Wales - Office for National Statistics, June 2022. URL <https://www.ons.gov.uk/peoplepopulationandcommunity/populationandmigration/populationestimates/bulletins/populationandhouseholdestimatesenglandandwales/census2021>.
- R. Ross. An application of the theory of probabilities to the study of a priori pathometry.—Part I. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 92(638):204–230, February 1916. .
- Royal Society. Face masks and coverings for the general public. 2020. URL <https://royalsociety.org/-/media/policy/projects/set-c/set-c-facemasks.pdf>.
- B. Rozemberczki, P. Scherer, Y. He, G. Panagopoulos, A. Riedel, M. Astefanoaei, O. Kiss, F. Beres, G. López, N. Collignon, and R. Sarkar. PyTorch Geometric Temporal: Spatiotemporal Signal Processing with Neural Machine Learning Models. *arXiv:2104.07788 [cs]*, June 2021. URL <http://arxiv.org/abs/2104.07788>.
- G. Rozhnova and A. Nunes. SIRS Dynamics on Random Networks: Simulations and Analytical Models. In J. Zhou, editor, *Complex Sciences*, Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, pages 792–797, Berlin, Heidelberg, 2009. Springer. ISBN 978-3-642-02466-5. .
- X. Rui, F. Meng, Z. Wang, G. Yuan, and C. Du. SPIR: The potential spreaders involved SIR model for information diffusion in social networks. *Physica A: Statistical Mechanics and its Applications*, 506(C):254–269, 2018. URL <https://ideas.repec.org//a/eee/phsmap/v506y2018icp254-269.html>.
- G. Rummery and M. Niranjana. On-Line Q-Learning Using Connectionist Systems. *Technical Report CUED/F-INFENG/TR 166*, November 1994.
- A. C. Rusu, R. Emonet, and K. Farrahi. Modelling digital and manual contact tracing for COVID-19. Are low uptakes and missed contacts deal-breakers? *PLOS ONE*, 16(11):e0259969, November 2021. ISSN 1932-6203. .

- A. C. Rusu, K. Farrahi, and M. Niranjani. Flattening the curve through reinforcement learning driven test and trace policies. In *16th EAI International Conference on Pervasive Computing Technologies for Healthcare (12/12/22 - 14/12/22)*, December 2022. URL <https://eprints.soton.ac.uk/473363/>.
- A. C. Rusu, K. Farrahi, and M. Niranjani. EpiCURB: Learning to derive epidemic control policies. *IEEE Pervasive Computing*, November 2023. ISSN 1536-1268. .
- M. Salathé and J. H. Jones. Dynamics and Control of Diseases in Networks with Community Structure. *PLOS Computational Biology*, 6(4):e1000736, April 2010. ISSN 1553-7358. .
- I. Sample and L. Geddes. Revised report on impact of Covid lockdowns ‘adds little insight’. *The Guardian*, June 2023. ISSN 0261-3077. URL <https://www.theguardian.com/society/2023/jun/05/revised-report-on-impact-of-covid-lockdowns-leaves-unanswered-questions>.
- R. Sato, M. Yamada, and H. Kashima. Random Features Strengthen Graph Neural Networks, January 2021.
- F. Scarselli, M. Gori, Ah Chung Tsoi, M. Hagenbuchner, and G. Monfardini. The Graph Neural Network Model. *IEEE Transactions on Neural Networks*, 20(1):61–80, January 2009. ISSN 1045-9227, 1941-0093. .
- B. Schölkopf, A. Smola, and K.-R. Müller. Kernel principal component analysis. In W. Gerstner, A. Germond, M. Hasler, and J.-D. Nicoud, editors, *Artificial Neural Networks — ICANN’97*, Lecture Notes in Computer Science, pages 583–588, Berlin, Heidelberg, 1997. Springer. ISBN 978-3-540-69620-9. .
- U. Schöning. Graph isomorphism is in the low hierarchy. *Journal of Computer and System Sciences*, 37(3):312–323, December 1988. ISSN 0022-0000. .
- J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal Policy Optimization Algorithms, August 2017.
- J. Schulman, P. Moritz, S. Levine, M. Jordan, and P. Abbeel. High-Dimensional Continuous Control Using Generalized Advantage Estimation, October 2018.
- M. Serafino, H. S. Monteiro, S. Luo, S. D. S. Reis, C. Igual, A. S. L. Neto, M. Travizano, J. S. A. Jr, and H. A. Makse. Digital contact tracing and network theory to stop the spread of COVID-19 using big-data on human mobility geolocalization. *PLOS Computational Biology*, 18(4):e1009865, April 2022. ISSN 1553-7358. .
- C. Shah, N. Dehmamy, N. Perra, M. Chinazzi, A.-L. Barabási, A. Vespignani, and R. Yu. Finding Patient Zero: Learning Contagion Source with Graph Neural Networks. June 2020. URL <http://arxiv.org/abs/2006.11913>.

- R. Sharan and T. Ideker. Modeling cellular machinery through biological network comparison. *Nature Biotechnology*, 24(4):427–433, April 2006. ISSN 1546-1696. .
- S. Shelke and V. Attar. Source detection of rumor in social network – A review. *Online Social Networks and Media*, 9:30–42, January 2019. ISSN 2468-6964. .
- N. Shervashidze, S. V. N. Vishwanathan, T. Petri, K. Mehlhorn, and K. Borgwardt. Efficient graphlet kernels for large graph comparison. In *Artificial Intelligence and Statistics*, pages 488–495. PMLR, April 2009. URL <http://proceedings.mlr.press/v5/shervashidze09a.html>.
- N. Shervashidze, P. Schweitzer, E. J. van Leeuwen, K. Mehlhorn, and K. M. Borgwardt. Weisfeiler-Lehman Graph Kernels. *The Journal of Machine Learning Research*, 12(null): 2539–2561, November 2011. ISSN 1532-4435.
- A. Sigal. Milder disease with Omicron: Is it the virus or the pre-existing immunity? *Nature Reviews Immunology*, 22(2):69–71, February 2022. ISSN 1474-1741. .
- D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, January 2016. ISSN 1476-4687. .
- D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, L. Sifre, D. Kumaran, T. Graepel, T. Lillicrap, K. Simonyan, and D. Hassabis. Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm, December 2017.
- B. W. Silverman. *Density Estimation for Statistics and Data Analysis*. Number 26 in Monographs on Statistics and Applied Probability. Chapman and Hall, London ; New York, 1986. ISBN 978-0-412-24620-3.
- D. E. Simmons, J. P. Coon, and A. Datta. The von Neumann Theil index: Characterizing graph centralization using the von Neumann index. *Journal of Complex Networks*, 6(6):859–876, December 2018. ISSN 2051-1329. .
- T. Smieszek, L. Fiebig, and R. W. Scholz. Models of epidemics: When contact repetition and clustering should be included. *Theoretical Biology and Medical Modelling*, 6(1):11, June 2009. ISSN 1742-4682. .
- G. D. Smith, F. Ng, and R. Watson. “Masking the evidence”: Perspectives of the COVID-19 pandemic. *Journal of Clinical Nursing*, 29(19-20):3580–3583, October 2020. ISSN 0962-1067. .

- J. Smith. Demand for Covid vaccines falls amid waning appetite for booster shots. *Financial Times*, April 2022.
- R. L. Smith, L. L. Gibson, P. P. Martinez, R. Ke, A. Mirza, M. Conte, N. Gallagher, A. Conte, L. Wang, R. Fredrickson, D. C. Edmonson, M. E. Baughman, K. K. Chiu, H. Choi, T. W. Jensen, K. R. Scardina, S. Bradley, S. L. Gloss, C. Reinhart, J. Yedetore, A. N. Owens, J. Broach, B. Barton, P. Lazar, D. Henness, T. Young, A. Dunnett, M. L. Robinson, H. H. Mostafa, A. Pekosz, Y. C. Manabe, W. J. Heetderks, D. D. McManus, and C. B. Brooke. Longitudinal Assessment of Diagnostic Test Performance Over the Course of Acute SARS-CoV-2 Infection. *The Journal of Infectious Diseases*, 224(6): 976–982, September 2021. ISSN 0022-1899. .
- H. Song, H. Jang, H. H. Tran, S.-e. Yoon, K. Son, D. Yun, H. Chung, and Y. Yi. Solving Continual Combinatorial Selection via Deep Reinforcement Learning. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, pages 3467–3474, Macao, China, August 2019. International Joint Conferences on Artificial Intelligence Organization. ISBN 978-0-9992411-4-1. .
- M. Starnini, A. Baronchelli, and R. Pastor-Satorras. Modeling human dynamics of face-to-face interaction networks. *Physical Review Letters*, 110(16):168701, April 2013. ISSN 0031-9007, 1079-7114. .
- J. Stehlé, N. Voirin, A. Barrat, C. Cattuto, V. Colizza, L. Isella, C. Régis, J.-F. Pinton, N. Khanafer, W. Van den Broeck, and P. Vanhems. Simulation of an SEIR infectious disease model on the dynamic contact network of conference attendees. *BMC Medicine*, 9(1):87, July 2011. ISSN 1741-7015. .
- G. Streftaris and G. J. Gibson. Non-exponential tolerance to infection in epidemic systems—modeling, inference, and assessment. *Biostatistics*, 13(4):580–593, September 2012. ISSN 1465-4644. .
- Z. Su, A. Cheshmehzangi, D. McDonnell, C. P. da Veiga, and Y.-T. Xiang. Mind the “Vaccine Fatigue”. *Frontiers in Immunology*, 13, 2022. ISSN 1664-3224. URL <https://www.frontiersin.org/articles/10.3389/fimmu.2022.839433>.
- A. Sudbury. The proportion of the population never hearing a rumour. *Journal of Applied Probability*, 22(2):443–446, June 1985. ISSN 0021-9002, 1475-6072. .
- S. R. Sukumar and J. Nutaro. Agent-Based vs. Equation-Based Epidemiological Models: A Model Selection Case Study. In *Proceedings of the 2012 ASE International Conference on BioMedical Computing, BioMedCom 2012*, pages 74–79, December 2012. .
- W. Sun, Q. Li, P. Wang, and J. Hou. Heuristic Network Security Risk Assessment Based on Attack Graph. In M. R. Khosravi, Q. He, and H. Dai, editors, *Cloud Computing*, Lecture Notes of the Institute for Computer Sciences, Social Informatics

- and Telecommunications Engineering, pages 181–194, Cham, 2022. Springer International Publishing. ISBN 978-3-030-99191-3. .
- R. S. Sutton. Learning to predict by the methods of temporal differences. *Machine Learning*, 3(1):9–44, August 1988. ISSN 1573-0565. .
- R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. Adaptive Computation and Machine Learning Series. The MIT Press, Cambridge, Massachusetts, second edition edition, 2018. ISBN 978-0-262-03924-6.
- M. Szummer and T. Jaakkola. Partially labeled classification with Markov random walks. In *Advances in Neural Information Processing Systems*, volume 14. MIT Press, 2002. URL <https://papers.nips.cc/paper/2001/hash/a82d922b133be19c1171534e6594f754-Abstract.html>.
- S. Tian, S. Mo, L. Wang, and Z. Peng. Deep Reinforcement Learning-Based Approach to Tackle Topic-Aware Influence Maximization. *Data Science and Engineering*, 5(1): 1–11, March 2020. ISSN 2364-1541. .
- J. Toenshoff, M. Ritzert, H. Wolf, and M. Grohe. Graph Neural Networks for Maximum Constraint Satisfaction. *arXiv:1909.08387 [cs]*, February 2020. URL <http://arxiv.org/abs/1909.08387>.
- J. Tolles and T. Luong. Modeling Epidemics With Compartmental Models. *JAMA*, 323(24):2515–2516, June 2020. ISSN 0098-7484. .
- A. Tomy, M. Razzanelli, F. Di Lauro, D. Rus, and C. Della Santina. Estimating the state of epidemics spreading with graph neural networks. *Nonlinear Dynamics*, 109(1): 249–263, July 2022. ISSN 1573-269X. .
- M. Tounsi. A dynamic heuristic for WSNs routing. *Cogent Engineering*, 8(1):1919040, January 2021. ISSN null. .
- L. S. Tsimring and R. Huerta. Modeling of contact tracing in social networks. *Physica A: Statistical Mechanics and its Applications*, 325(1):33–39, July 2003. ISSN 0378-4371. .
- Z. Tufekci. Opinion — Here’s Why the Science Is Clear That Masks Work. *The New York Times*, March 2023. ISSN 0362-4331. URL <https://www.nytimes.com/2023/03/10/opinion/masks-work-cochrane-study.html>.
- G. Tzortzis and A. Likas. The global kernel k-means clustering algorithm. In *2008 IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence)*, pages 1977–1984, June 2008. .
- H. Ueki, Y. Furusawa, K. Iwatsuki-Horimoto, M. Imai, H. Kabata, H. Nishimura, and Y. Kawaoka. Effectiveness of Face Masks in Preventing Airborne Transmission of SARS-CoV-2. *mSphere*, 5(5):e00637–20, October 2020. .

- L. van der Maaten and G. Hinton. Visualizing High-Dimensional Data Using t-SNE. *Journal of Machine Learning Research*, 9(nov):2579–2605, 2008. ISSN 1532-4435.
- H. van Hasselt, S. Madjiheurem, M. Hessel, D. Silver, A. Barreto, and D. Borsa. Expected Eligibility Traces. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(11):9997–10005, May 2021. ISSN 2374-3468, 2159-5399. .
- P. Van Mieghem and Q. Liu. Explicit non-Markovian susceptible-infected-susceptible mean-field epidemic threshold for Weibull and Gamma infections but Poisson curings. *Physical Review. E*, 100(2):022317, August 2019. ISSN 2470-0045. .
- A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is All you Need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL <https://papers.nips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>.
- P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio. Graph Attention Networks. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=rJXmpikCZ>.
- E. Vergu, H. Busson, and P. Ezanno. Impact of the Infection Period Distribution on the Epidemic Spread in a Metapopulation Model. *PLoS ONE*, 5(2), February 2010. ISSN 1932-6203. .
- C. L. Vestergaard and M. Génois. Temporal Gillespie Algorithm: Fast Simulation of Contagion Processes on Time-Varying Networks. *PLOS Computational Biology*, 11(10):e1004579, October 2015. ISSN 1553-7358. .
- J. Wallinga and M. Lipsitch. How generation intervals shape the relationship between growth rates and reproductive numbers. *Proceedings of the Royal Society B: Biological Sciences*, 274(1609):599–604, February 2007. .
- H. Wang, K. R. Paulson, S. A. Pease, S. Watson, H. Comfort, P. Zheng, A. Y. Aravkin, C. Bisignano, R. M. Barber, T. Alam, J. E. Fuller, E. A. May, D. P. Jones, M. E. Frisch, C. Abbafati, C. Adolph, A. Allorant, J. O. Amlag, B. Bang-Jensen, G. J. Bertolacci, S. S. Bloom, A. Carter, E. Castro, S. Chakrabarti, J. Chattopadhyay, R. M. Cogen, J. K. Collins, K. Cooperrider, X. Dai, W. J. Dangel, F. Daoud, C. Dapper, A. Deen, B. B. Duncan, M. Erickson, S. B. Ewald, T. Fedosseeva, A. J. Ferrari, J. J. Frostad, N. Fullman, J. Gallagher, A. Gamkrelidze, G. Guo, J. He, M. Helak, N. J. Henry, E. N. Hulland, B. M. Huntley, M. Kereselidze, A. Lazzar-Atwood, K. E. LeGrand, A. Lindstrom, E. Linebarger, P. A. Lotufo, R. Lozano, B. Magistro, D. C. Malta, J. Månsson, A. M. M. Herrera, F. Marinho, A. H. Mirkuzie, A. T. Misganaw, L. Monasta, P. Naik, S. Nomura, E. G. O’Brien, J. K. O’Halloran, L. T. Olana, S. M. Ostroff, L. Penberthy, R. C. R. Jr, G. Reinke, A. L. P. Ribeiro, D. F. Santomauro, M. I. Schmidt, D. H. Shaw, B. S. Sheena, A. Sholokhov, N. Skhvitaridze, R. J. D. Sorensen,

- E. E. Spurlock, R. Syailendrawati, R. Topor-Madry, C. E. Troeger, R. Walcott, A. Walker, C. S. Wiysonge, N. A. Worku, B. Zigler, D. M. Pigott, M. Naghavi, A. H. Mokdad, S. S. Lim, S. I. Hay, E. Gakidou, and C. J. L. Murray. Estimating excess mortality due to the COVID-19 pandemic: A systematic analysis of COVID-19-related mortality, 2020–21. *The Lancet*, 399(10334):1513–1536, April 2022. ISSN 0140-6736, 1474-547X. .
- K. Wang, Z. Shen, C. Huang, C.-H. Wu, Y. Dong, and A. Kanakia. Microsoft Academic Graph: When experts are not enough. *Quantitative Science Studies*, 1(1):396–413, February 2020a. ISSN 2641-3337. .
- L. Wang, X. Ben, A. Adiga, A. Sadilek, A. Tendulkar, S. Venkatramanan, A. Vullikanti, G. Aggarwal, A. Talekar, J. Chen, B. Lewis, S. Swarup, A. Kapoor, M. Tambe, and M. Marathe. Using Mobility Data to Understand and Forecast COVID19 Dynamics. *medRxiv*, page 2020.12.13.20248129, December 2020b. .
- L. Wang and B. C. Wood. An epidemiological approach to model the viral propagation of memes. *Applied Mathematical Modelling*, 35(11):5442–5447, November 2011. ISSN 0307-904X. .
- X. F. Wang and G. R. Chen. Complex networks: Small-world, scale-free and beyond. *IEEE Circuits and Systems Magazine*, 3(1):6–20, 2003. ISSN 1531-636X. .
- Y. Wang, W. Huang, L. Zong, T. Wang, and D. Yang. Influence maximization with limit cost in social network. *Science China Information Sciences*, 56(7):1–14, July 2013. ISSN 1869-1919. .
- M. Waniek, P. Holme, K. Farrahi, R. Emonet, M. Cebrian, and T. Rahwan. Trading contact tracing efficiency for finding patient zero. *Scientific Reports*, 12(1):22582, December 2022. ISSN 2045-2322. .
- C. Watkins. Learning From Delayed Rewards. January 1989.
- D. J. Watts and S. H. Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393(6684):440–442, June 1998. ISSN 1476-4687. .
- D. J. Watts, P. S. Dodds, and J. D. s. a. e. a. T. E. served as associate editor for this article. Influentials, Networks, and Public Opinion Formation. *Journal of Consumer Research*, 34(4):441–458, 2007. ISSN 0093-5301. .
- B. Y. Weisfeiler and A. A. Lehman. The reduction of a graph to canonical form and the algebra which appears therein. *Nauchno-Technicheskaya Informatsia*, 2(9):12–16, 1968. URL https://www.iti.zcu.cz/wl2018/pdf/wl_paper_translation.pdf.
- WHO. WHO Coronavirus (COVID-19) Dashboard, August 2023. URL <https://covid19.who.int>.

- F. Wu, T. Zhang, A. H. de Souza Jr., C. Fifty, T. Yu, and K. Q. Weinberger. Simplifying Graph Convolutional Networks. *arXiv:1902.07153 [cs, stat]*, June 2019. URL <http://arxiv.org/abs/1902.07153>.
- Y. Wu, M. Gao, M. Zeng, F. Chen, M. Li, and J. Zhang. BridgeDPI: A Novel Graph Neural Network for Predicting Drug-Protein Interactions. *arXiv:2101.12547 [cs, q-bio]*, January 2021. URL <http://arxiv.org/abs/2101.12547>.
- C. Wymant, L. Ferretti, D. Tsallis, M. Charalambides, L. Abeler-Dörner, D. Bonsall, R. Hinch, M. Kendall, L. Milsom, M. Ayres, C. Holmes, M. Briers, and C. Fraser. The epidemiological impact of the NHS COVID-19 app. *Nature*, 594(7863):408–412, June 2021. ISSN 1476-4687. .
- K. Xu, W. Hu, J. Leskovec, and S. Jegelka. How Powerful are Graph Neural Networks? In *arXiv:1810.00826 [Cs, Stat]*, February 2019. URL <http://arxiv.org/abs/1810.00826>.
- Z. Yan, X. Zhou, and R. Du. An enhanced SIR dynamic model: The timing and changes in public opinion in the process of information diffusion. *Electronic Commerce Research*, September 2022. ISSN 1572-9362. .
- Z. Yang, W. W. Cohen, and R. Salakhutdinov. Revisiting Semi-Supervised Learning with Graph Embeddings. *arXiv:1603.08861 [cs]*, May 2016. URL <http://arxiv.org/abs/1603.08861>.
- X. Yu, S. Lu, L. Guo, S.-H. Wang, and Y.-D. Zhang. ResGNet-C: A graph convolutional neural network for detection of COVID-19. *Neurocomputing*, 452:592–605, September 2021. ISSN 0925-2312. .
- W. Yuan, K. He, D. Guan, L. Zhou, and C. Li. Graph kernel based link prediction for signed social networks. *Information Fusion*, 46:1–10, March 2019. ISSN 1566-2535. .
- L. Zenk, G. Steiner, M. Pina e Cunha, M. D. Laubichler, M. Bertau, M. J. Kainz, C. Jäger, and E. S. Schernhammer. Fast Response to Superspreading: Uncertainty and Complexity in the Context of COVID-19. *International Journal of Environmental Research and Public Health*, 17(21):7884, October 2020. ISSN 1660-4601. .
- X.-X. Zhan, A. Hanjalic, and H. Wang. Information diffusion backbones in temporal networks. *Scientific Reports*, 9(1):6798, May 2019. ISSN 2045-2322. .
- S. Zhao and H. Chen. Modeling the epidemic dynamics and control of COVID-19 outbreak in China. *Quantitative Biology (Beijing, China)*, pages 1–9, March 2020. ISSN 2095-4689. .
- D. Zhou, O. Bousquet, T. Lal, J. Weston, and B. Schölkopf. Learning with Local and Global Consistency. In *Advances in Neural Information Processing Systems*, volume 16.

MIT Press, 2003. URL <https://papers.nips.cc/paper/2003/hash/87682805257e619d49b8e0dfdc14affa-Abstract.html>.

J. Zhou, G. Cui, S. Hu, Z. Zhang, C. Yang, Z. Liu, L. Wang, C. Li, and M. Sun. Graph neural networks: A review of methods and applications. *AI Open*, 1:57–81, 2020. ISSN 26666510. .