Efficient Epidemic Intervention Generation: A Graph Adversarial Attack Perspective

Yunxiao Li, Mingyang Wei, Zewen Liu, Max S.Y. Lau, Wei Jin
Emory University
{yunxiao.li2,mingyang.wei,zewen.liu,msy.lau,wei.jin}@emory.edu

Abstract

Epidemic outbreaks spread through networked interactions, making graph-based modeling essential for effective intervention strategies. In this study, we propose EpiAdv, a scalable approach that integrates a modified Susceptible-Infected-Recovered (SIR) model with graph adversarial attack techniques to control infection spread on time-varying graphs. EpiAdv dynamically adjusts node states or edge weights using gradient descent, optimizing an objective function to balance infection reduction and intervention. Experimental results show that EpiAdv significantly outperforms baseline methods, such as expert-based and degree-based heuristics, in reducing infection rates while minimizing isolation costs. Furthermore, adaptability tests demonstrate the robustness of EpiAdv across different intervention timing scenarios, providing an effective and efficient framework for managing epidemics in dynamic networks.

ACM Reference Format:
Yunxiao Li, Mingyang Wei, Zewen Liu, Max S.Y. Lau, Wei Jin. 2018. Efficient Epidemic Intervention Generation: A Graph Adversarial Attack Perspective. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation emai (Conference acronym 'XX)*. ACM, New York, NY, USA, 5 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

Pandemics have consistently posed significant threats to global public health, evidenced by historical outbreaks such as SARS, the COVID-19 pandemic, and seasonal influenza. While these events vary in severity, they share common characteristics: rapid transmission, substantial morbidity and mortality, and the capacity to overwhelm healthcare systems. Similar to the COVID-19 crisis, pandemics frequently disrupt social and economic structures [16, 19], calling for timely and effective interventions. However, controlling pandemics presents multifaceted challenges, underscoring the urgent need for innovative approaches to intervention strategies.

Interventions such as city lockdowns, quarantines, hospitalizations, vaccinations, and mask mandates [30] have been critical in controlling the spread of diseases. However, the effectiveness of these measures is often hindered by prolonged implementation timelines and significant economic costs. To optimize outcomes, intervention strategies must involve a meticulous selection of targets, such as vulnerable populations and high-risk geographic regions. This process heavily relies on expert judgment, which can be labor-intensive and may introduce delays, affecting the timeliness and scalability of responses in dynamic epidemic scenarios. Consequently, there is a growing need for data-driven strategies to enhance adaptability and scalability.

Given that epidemic transmission inherently aligns with graph-based modeling and involves a sequential decision-making problem, graph neural networks (GNNs) [11, 29, 31] and reinforcement learning [27] have demonstrated significant potential for improving the scalability and precision of epidemic response strategies [4, 8, 17]. Individuals are usually represented as nodes, and their interactions as edges, naturally capturing the networked nature of disease spread. Targeted interventions, such as reducing social interactions or implementing quarantines, can be modeled as modifications to edge weights or network structure. GNNs excel in capturing disease propagation patterns within these graphs, while reinforcement learning provides a powerful mechanism for optimizing sequential intervention decisions over time. However, these methods often require a large amount of high-quality data for accurate predictions, which can be difficult to obtain during rapidly evolving outbreaks. Computational complexity also poses scalability concerns, especially in large, densely connected graphs.

Our Contributions: To address these limitations, we draw on the concept of graph adversarial attacks—techniques that manipulate graph structures or node features to effectively and efficiently disrupt the performance of graph-based algorithms [9, 33]. We view interventions such as quarantines and social distancing as forms of adversarial manipulation, strategically deployed to control disease spread. More specifically, our approach integrates a time-varying SIR model with adversarial, gradient-driven optimization to design an efficient intervention strategy, which balances infection mitigation and associated social costs. In this manner, our method not only reduces reliance on extensive labeled data but also scales efficiently to large, dynamically evolving networks. In contrast to existing methods, our framework offers the following key advantages:

(a) **Robust Epidemic Modeling:** By capturing potential changes in real-world graph structures, our SIR model more accurately represents the complex interplay between human contact behavior and disease spread.
(b) **An Adversarial Lens:** Our method uses minimal node isolation to disrupt transmission pathways, thereby maximizing the impact of limited interventions.
(c) **Balanced Intervention Costs:** Our method combines infection rates and isolation overhead into a unified objective, ensuring an optimal trade-off between public health benefits and social/economic costs.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference acronym 'XX, June 03–05, 2018, Woodstock, NY

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/18/06
https://doi.org/XXXXXXX.XXXXXXX

2 Related Work

Network-Based Epidemic Modeling. Early approaches, including classical SIR and SEIR models [1, 6], used differential equations for smaller or homogeneous contexts. However, with the widespread availability of social network and mobile data, researchers began incorporating network structures into epidemic modeling to account for the complex and heterogeneous contact patterns among individuals [20]. In this context, many static network models assume that node connections and transmission probabilities remain unchanged throughout the process [7]. In reality, however, social interactions evolve over time, thus giving rise to the need for time-varying network models [15, 26].

Machine Learning for Epidemic Interventions. Recent advances in big data and machine learning have driven the exploration of supervised and reinforcement learning for epidemic intervention design. Graph Neural Networks (GNNs): Liu et al.[14] review GNN applications in epidemic tasks, emphasizing the ability to capture complex relational dynamics but noting significant computational and data demands. Jhun et al.[8] propose a GNN-based vaccination strategy that surpasses centrality approaches, albeit with high resource costs for large networks. Reinforcement Learning (RL): Moving beyond pure GNN methods, Song et al.[24] introduce DURLECA, a dual-objective RL framework integrating Flow-GNN for interregional mobility management, which relies on robust aggregated data. Feng et al.[4] present IDRLECA, merging GNN-based models with individual-level infections to balance epidemic control and costs, yet it hinges on precise data. Meirom et al. [17] developed a GNN-RL system for intervention prioritization under partial observability, though scalability and rapid adaptation remain challenges.

Graph Adversarial Attack. Graph adversarial attack strategically modifies a network's structure or node features, often by editing edges or altering node attributes to degrade the performance of models [2, 10, 18, 25]. This adversarial perspective aligns with resource-constrained real-world scenarios where testing, quarantine, or vaccination capacities are limited. To the best of our knowledge, our work presents the first study for applying graph adversarial attacks to *epidemic interventions*.

3 Problem Definition

We consider a sequence of time-varying graphs $\{G_t\}$ with a fixed node set V and dynamic edges E_t. Each edge $e \in E_t$ is associated with a transmission probability $p_e(t) \in [0, 1]$, representing the likelihood of infection transmission at time t. Each node $v \in V$ can be in one of three states at any time t: susceptible (S), infected (I), or recovered (R). The state $S_v(t)$ of node v at time t depends on its own state at time $t-1$, the states of its neighboring nodes at time $t-1$, and the interactions through edges in E_t. Our objective is to design an intervention strategy that curbs infection spread by selectively isolating nodes while minimizing the total cost.

4 Methodology

In this section, we present EpiAdv, a novel framework for optimizing intervention strategies on dynamic graphs. EpiAdv combines a modified SIR model with a graph adversarial attack formulation to disrupt transmission pathways while minimizing interventions.

4.1 Infection Forecasting

We propose a *modified SIR* model [15] on time-varying graphs, addressing limitations of static contact structures and uniform infection periods. Classical SIR models assume $S \longrightarrow I \longrightarrow R$ transitions with fixed connectivity and homogeneous infection/recovery durations, which can oversimplify real-world dynamics. By contrast, our model tracks evolving edges and assigns node-specific infectious durations, enabling more realistic forecasts. Formally, we simulate disease spread on a *time-varying graph* $\{G_t\}$, where nodes progress from susceptible (S) to infected (I) and eventually to recovered (R). Upon recovery, a node remains in R indefinitely, reflecting immunity. Such time-dependent edges and heterogeneous infectious periods are crucial for adaptive, accurate intervention strategies.

Transmission Probability: Each edge $e = (v, u) \in E_t$ in the graph at time t is assigned a *transmission probability* $p_e(t) \in [0, 1]$, representing the likelihood that infection is passed along that edge within the interval $(t-1, t]$. As real-world contact patterns vary (e.g., due to behavior or policy changes), the set of edges E_t and the probabilities $p_e(t)$ can change over time. This temporal flexibility is key to accurately reflecting shifting social structures and intervention effects. *Infection Mechanism:* For each node v at time t, we define $E_v(t) = \{e \in E_t \mid e = (v, u)$ and $S_u(t-1) = I\}$ as the set of edges through which v can contract the infection from neighbors that were infected at time $(t-1)$. Assuming independence among potential transmission events, the probability that v remains in state S at time t is: $P_v^S(t) = \prod_{e \in E_v(t)} \big[1 - p_e(t)\big]$. Thus, the probability that v transitions to I at time t is: $1 - P_v^S(t) = 1 - \prod_{e \in E_v(t)} \big[1 - p_e(t)\big]$. Once infected, v may transmit the pathogen to other susceptible neighbors in subsequent time steps. *Recovery Process:* Each infected node v remains infectious for a random duration D_v, sampled from a distribution aligned with the disease's natural history (e.g., geometric, exponential, or gamma). This feature addresses individual variability in how long people remain contagious, a key limitation in classical models. Let T_v denote the infection onset for node v. At $(T_v + D_v)$, v transitions from I to R and gains immunity permanently, reflecting the well-established assumption of no reinfection.

Furthermore, we have three modeling assumptions: (a) *No Reinfection*: Once a node transitions to R, it stays recovered indefinitely. (b) *Full Observability*: We assume node states are fully known. In reality, data-driven methods may be required to infer these states. (c) *Dynamic Parameters*: The values of $p_e(t)$ and D_v maybe informed by empirical data, and can be adjusted to reflect interventions (e.g., isolation, distancing), offering flexible "what-if" analyses.

4.2 Implementing Intervention Strategies from the Perspective of Graph Adversarial Attack

To effectively control the spread of infection, we frame our intervention strategy within the context of graph adversarial attack. Graph attack refers to the deliberate modification of a graph's structure—through the removal or alteration of nodes and edges—to influence the outcome of dynamic processes on the graph. Our approach treats the epidemic control problem as a specific instance of a graph attack, aiming to minimally perturb the graph structure to achieve maximal disruption of the infection spread.

In graph adversarial attack, subtle perturbations (e.g., adding or removing edges) are introduced to degrade the performance of graph-based machine learning models. Here, we leverage a similar idea but with a distinctly different goal: *protecting* the system by limiting the disease's ability to propagate, rather than disrupting it covertly. Specifically, we propose an intervention strategy represented by an intervention state vector $\boldsymbol{\theta}(t) \in \{0,1\}^{|V|}$. Each element $\theta_v(t)$ corresponds to the intervention state of node v at time t: (a) $\theta_v(t) = 0$: Node v is in the no intervention state, and the transmission probabilities along its edges remain unchanged. (b) $\theta_v(t) = 1$: Node v is in the isolated state, and all edges connected to v have their infection probabilities set to zero, effectively preventing any transmission to or from this node. This is analogous to removing the node in a graph attack. Formally, the transmission probability $p_{uv}(t)$ for an edge (u, v) at time t is given by:

$$p_{uv}(t) = \begin{cases} p_{uv}, & \text{if } \theta_u(t) = 0 \text{ and } \theta_v(t) = 0, \\ 0, & \text{otherwise.} \end{cases}$$

where p_{uv} is the original infection probability of the edge (u, v). By controlling $\theta_v(t)$ for each node over time, the strategy aims to disrupt the underlying transmission pathways, akin to a strategic adversarial perturbation, but in service of epidemic mitigation rather than model degradation.

Having established the connection to graph adversarial attack, we now detail the concrete implementation of our proposed intervention strategy. Specifically, we formulate epidemic control as a constrained optimization problem designed to minimize infection spread while limiting the extent of interventions. We first define two competing objectives: (1) reducing overall infection, and (2) restricting the cost associated with frequent or excessive interventions. Formally, we seek to minimize the total number of infections subject to a cap on allowable interventions: $\theta_{\text{individuals}} = \arg\min_\theta J_{\text{infection}}(\theta)$, subject to $J_{\text{intervention}}(\theta) \leq \gamma$, where γ is a predefined upper limit on the total allowable intervention effort. To make the above constraint more tractable in practice, we combine infection suppression and intervention limitation into a single cost function: $J(\boldsymbol{\theta}) = J_{\text{infection}}(\boldsymbol{\theta}) + \lambda J_{\text{intervention}}(\boldsymbol{\theta})$, where $\lambda > 0$ is a weighting parameter that balances the trade-off between minimizing the epidemic and avoiding overly disruptive interventions. The infection cost is expressed as $J_{\text{infection}}(\boldsymbol{\theta}) = \sum_{t=0}^{T} \sum_{v \in V} I_v(t)$, where $I_v(t)$ is the infection state of node v at time t, and T denotes the final time step of the simulation; the intervention cost is expressed as $J_{\text{intervention}}(\boldsymbol{\theta}) = \sum_{t=0}^{T} \sum_{v \in V} \left|\theta_v(t) - \theta_v(t-1)\right|$, representing the total number of times nodes switch from no intervention ($\theta_v = 0$) to isolated ($\theta_v = 1$) or vice versa. This penalizes strategies that rapidly toggle node states.

Gradient-Based Optimization. We minimize $J(\boldsymbol{\theta})$ via a gradient descent approach. Let $\eta > 0$ be the learning rate, and let $\nabla_{\boldsymbol{\theta}} J(\boldsymbol{\theta}(t))$ denote the gradient of J with respect to $\boldsymbol{\theta}$ at time t. We iteratively update: $\boldsymbol{\theta}(t+1) = \boldsymbol{\theta}(t) - \eta \nabla_{\boldsymbol{\theta}} J(\boldsymbol{\theta}(t))$. This procedure systematically adjusts each node's intervention state, seeking to reduce infections while incurring minimal intervention overhead. In contrast to classical adversarial attacks designed to degrade predictive models, our framework applies the adversarial perspective *constructively*, using minimal "perturbations" (node isolations) to achieve epidemic suppression instead of disruption.

By casting epidemic mitigation as a constrained optimization and applying a gradient-based solution, our intervention strategy selectively isolates high-risk nodes with minimal toggling, effectively balancing epidemiological impact and resource constraints.

5 Experiments

5.1 Experimental Setup

Table 1: Dataset Information.

Dataset	Nodes	Edges	Avg. Degree
E-Network	1,005	25,571	50.8
CA-GrQc	5,242	14,496	5.53
GEMSEC-RO	41,773	125,826	6.02

Real-World Datasets. We evaluated our proposed algorithm and baselines on three real-world networks that capture diverse facets of human social interaction. Specifically, we used **CA-GrQc** [21], **E-Network** [12], and **GEMSEC-RO** [22], each showcasing distinct topological and structural characteristics while reflecting various modes of social connectivity. Ranging from around 4,000 to over 40,000 nodes, these networks have been extensively utilized in prior research on network science and epidemiological modeling [23, 32] because they span different domains and exhibit diverse structural properties (e.g., varying average degrees, levels of clustering, and modularity). Such diversity allows us to evaluate how well our method generalizes across different contact patterns and scales. Table 1 summarizes key statistics for each dataset.

Infection Simulation. We employ a modified SIR model [15], as outlined in the Methodology section, to simulate the spread of infection within the networks. The transmission probability $p_e(t)$ for each edge $e \in E$ at time t is initialized based on empirical studies relevant to each dataset. This probability governs the likelihood of disease transmission between connected nodes and may be influenced by the intervention levels $\boldsymbol{\theta}(t)$. The model iteratively updates the states of the nodes over discrete time steps, allowing us to simulate the progression of the epidemic and assess the impact of various intervention strategies.

Baseline Methods. Although some reinforcement learning methods [24] have also been developed, their code and datasets were not provided. Therefore, in this study, we compare our method with several baselines commonly used in epidemic research: (a) **No Intervention**: Serves as the control scenario with no modifications to transmission probabilities. (b) **Lockdown** [5]: Simulates widespread social distancing measures by uniformly reducing transmission probabilities across all edges by a fixed percentage. (c) **Expert Random**: Selects a subset of nodes for intervention randomly, guided by expert-defined heuristics. (d) **Degree-Sample** [28]: Randomly samples nodes for intervention with probabilities proportional to their degrees, thereby targeting more connected nodes with higher likelihood. (e) **Degree-Order** [3]: Prioritizes interventions on nodes with larger degrees.

Evaluation Metrics. The performance of each intervention strategy is assessed using the following metrics: (a) **Infection Proportion**: The proportion of infected individuals relative to the total number of nodes, calculated as: $\frac{N_I}{N}$, where N_I is the number of infected nodes and N is the total number of nodes in the network. (b) **Isolation Proportion**: The proportion of isolated nodes relative to the total number of nodes, calculated as: $\frac{N_{\text{isolated}}}{N}$, where N_{isolated} is the number of nodes isolated due to interventions.

All the baselines and datasets are implemented with the comprehensive toolkit *EpiLearn* [13].

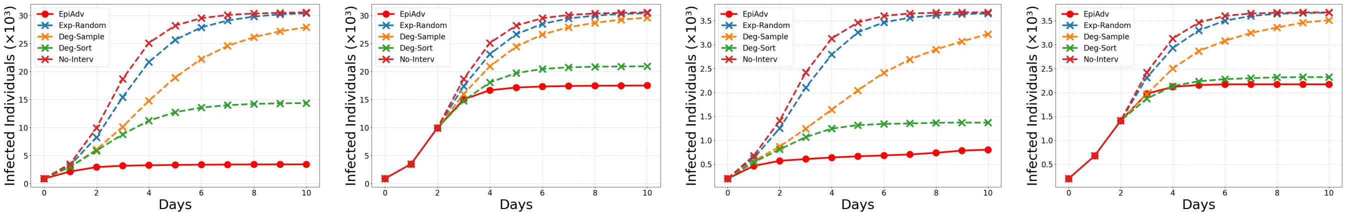

(a) Scenario Default for *GEMSEC*. **(b) Scenario Late for *GEMSEC-RO*** **(c) Scenario Default for *CA-GrQc*** **(d) Scenario Late for *CA-GrQc***

Figure 1: Comparison of intervention performance under Default and Late scenarios.

Table 2: Daily infection (%) | isolation proportion (%).

Method	E-Network	CA-GrQc	GEMSEC-RO
No Intervention	89.35 | 0.00	96.42 | 0.00	56.66 | 0.00
Exp-Random	78.51 | 67.00	46.48 | 64.95	39.70 | 68.23
Degree Sample	85.57 | 50.65	44.07 | 49.51	51.94 | 35.62
Degree Sorting	71.59 | 9.95	13.91 | 9.99	11.92 | 14.98
EpiAdv	**3.38 | 2.68**	**12.73 | 6.29**	**10.61 | 10.79**

5.2 Intervention Performance

Comparison Scenarios. We define t_{start} as the number of days that elapse before initiating epidemic interventions following the identification of the first patient. In this study, we consider the two scenarios: (a) **Scenario-Default** ($t_{\text{start}} = 1$): Interventions are initiated immediately upon detection of the outbreak; (b) **Scenario-Late** ($t_{\text{start}} = 3$): Interventions start three days after detection, introducing a two-day delay compared to the default scenario.

Effectiveness of EpiAdv. Under Scenario-Default on three datasets (*E-Network*, *CA-GrQc*, *GEMSEC-RO*), EpiAdv significantly outperforms standard methods in reducing infection and isolation, as shown in Table 2. Specifically, on the *E-Network* dataset, EpiAdv reduces the infection proportion to 3.38% and maintains a low isolation proportion of 2.68%, outperforming Degree Sorting, which results in an infection proportion of 71.59% and an isolation proportion of 9.95%. Similar superior performance is observed on the other two datasets. These results demonstrate that EpiAdv effectively curtails infection spread while minimizing social disruption, achieving better outcomes than commonly used methods.

Adaptability of EpiAdv to Different Scenarios. To assess the adaptability of EpiAdv, we apply two intervention cases—Default ($t_{\text{start}} = 1$) and Late ($t_{\text{start}} = 3$)—across two datasets that differ in scale: *GEMSEC-RO* (∼ 40k nodes) and *CA-GrQc* (∼ 5k nodes). From Figure 1, EpiAdv consistently maintains the lowest infection curves in both Default ($t_{\text{start}} = 1$) and Late ($t_{\text{start}} = 3$) scenarios. Even when delaying the interventions by two days, EpiAdv curbs the outbreak and plateaus at a lower level than competing methods, demonstrating its resilience to suboptimal response times. From Figures 1b and 1d, EpiAdv consistently performs well on both large (*GEMSEC-RO* ∼ 40,000 nodes) and small dataset (*CA-GrQc* ∼ 5,000 nodes). EpiAdv adapts effectively to networks of different scales and demonstrates low sensitivity to variations in network topology, underscoring its scalability and generality.

5.3 Practical Real-World Considerations

Flexible Adjustment of the Infection-Isolation Tradeoff: In real-world scenarios, the economic and logistical costs of isolation vary significantly across different regions. Therefore, our model must adjust the infection-isolation tradeoff according to diverse real-world conditions. This is achieved by tuning the parameter λ, which reflects the relative penalty on isolation. As illustrated in Figures 2(a) and 2(b), a smaller λ prioritizes reducing infection numbers and therefore results in more widespread isolation, whereas a larger λ lowers isolation counts at the expense of slightly higher infections. This flexibility allows decision-makers to tailor interventions to diverse economic, social, and epidemiological settings. **Real-World Constraints:** Real deployments often operate at coarser timescales (e.g., weekly updates), facing logistical hurdles, administrative delays, and infrastructural challenges. EpiAdv's modular framework adapts to daily, weekly, or bi-weekly re-optimizations, flexibly identifying which nodes or groups to isolate. **Social and Ethical Aspects:** Minimizing isolation is vital for public trust, economic stability, and psychological well-being. EpiAdv targets the most influential transmission nodes while limiting the size of isolation groups, balancing effective disease control with social costs.

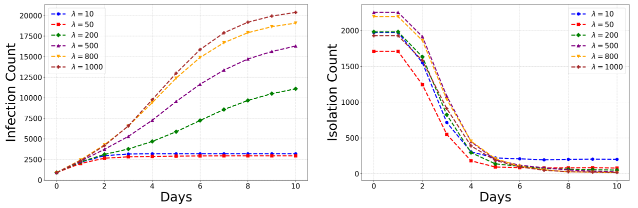

(a) Daily infections over time. **(b) Daily isolations over time.**

Figure 2: Daily infection (a) and isolation (b) counts under different penalty coefficients for isolation.

6 Conclusion

In this study, we introduce EpiAdv, an intervention strategy for epidemic control on time-varying graphs. EpiAdv integrates a modified SIR model with graph-based adversarial attack methods, striking a balance between infection control and intervention costs via gradient descent. Experiments on three real-world networks show it outperforms expert- and degree-based heuristics, with robustness tests confirming adaptability to varied intervention timings, while flexible tuning of the infection-isolation tradeoff accommodates diverse conditions and prioritizes minimal isolation to address social and ethical considerations.Emphasizing scalability and data-driven methods, EpiAdv addresses key epidemic management challenges, with potential extensions to heterogeneous node attributes and multi-layer networks.

References

[1] 1927. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 115, 772 (August 1927), 700–721. https://doi.org/10.1098/rspa.1927.0118

[2] Divya Anand Sinha, Yezi Liu, Ruijie Du, and Yanning Shen. 2024. Gradient Inversion Attack on Graph Neural Networks. *arXiv e-prints* (2024), arXiv–2411.

[3] Stephen Eubank, Hasan Guclu, V. S. Anil Kumar, Madhav V. Marathe, Aravind Srinivasan, Zoltán Toroczkai, and Nan Wang. 2004. Modelling disease outbreaks in realistic urban social networks. *Nature* 429, 6988 (May 2004), 180–184. https://doi.org/10.1038/nature02541

[4] Tao Feng, Sirui Song, Tong Xia, and Yong Li. 2023. Contact Tracing and Epidemic Intervention via Deep Reinforcement Learning. *ACM Transactions on Knowledge Discovery from Data* (2023).

[5] Thomas Hale, Noam Angrist, Beatriz Kira, Anna Petherick, Toby Phillips, and Samuel Webster. 2020. Variation in government responses to COVID-19. (2020).

[6] Tiberiu Harko, Francisco S.N. Lobo, and M.K. Mak. 2014. Exact analytical solutions of the Susceptible-Infected-Recovered (SIR) epidemic model and of the SIR model with equal death and birth rates. *Appl. Math. Comput.* 236 (2014), 184–194. https://doi.org/10.1016/j.amc.2014.03.030

[7] Petter Holme and Jari Saramäki. 2012. Temporal networks. *Physics Reports* 519, 3 (Oct. 2012), 97–125. https://doi.org/10.1016/j.physrep.2012.03.001

[8] Bukyoung Jhun. 2021. Effective vaccination strategy using graph neural network ansatz. *arXiv preprint arXiv:2111.00920* (2021).

[9] Wei Jin, Yaxing Li, Han Xu, Yiqi Wang, Shuiwang Ji, Charu Aggarwal, and Jiliang Tang. 2021. Adversarial attacks and defenses on graphs. *ACM SIGKDD Explorations Newsletter* 2 (2021).

[10] Wei Jin, Yao Ma, Xiaorui Liu, Xianfeng Tang, Suhang Wang, and Jiliang Tang. 2020. Graph structure learning for robust graph neural networks. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*. 66–74.

[11] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations*. https://openreview.net/forum?id=SJU4ayYgl

[12] Jure Leskovec, Jon Kleinberg, and Christos Faloutsos. 2007. Graph evolution: Densification and shrinking diameters. *ACM Transactions on Knowledge Discovery from Data* 1, 1 (March 2007), 2. https://doi.org/10.1145/1217299.1217301

[13] Zewen Liu, Yunxiao Li, Mingyang Wei, Guancheng Wan, Max SY Lau, and Wei Jin. 2024. EpiLearn: A Python Library for Machine Learning in Epidemic Modeling. *arXiv preprint arXiv:2406.06016* (2024).

[14] Zewen Liu, Guancheng Wan, B. Aditya Prakash, Max S. Y. Lau, and Wei Jin. 2024. A Review of Graph Neural Networks in Epidemic Modeling. *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (2024).

[15] Leonardo López and Xavier Rodo. 2021. A modified SEIR model to predict the COVID-19 outbreak in Spain and Italy: simulating control scenarios and multi-scale epidemics. *Results in Physics* 21 (2021), 103746.

[16] Wyatt G Madden, Wei Jin, Benjamin Lopman, Andreas Zufle, Benjamin Dalziel, C Jessica E. Metcalf, Bryan T Grenfell, and Max SY Lau. 2024. Deep neural networks for endemic measles dynamics: Comparative analysis and integration with mechanistic models. *PLOS Computational Biology* 20, 11 (2024), e1012616.

[17] Eli A. Meirom, Haggai Maron, Shie Mannor, and Gal Chechik. 2021. Controlling Graph Dynamics with Reinforcement Learning and Graph Neural Networks. *Proceedings of the 38th International Conference on Machine Learning (ICML)* (2021).

[18] Eli A. Meirom, Haggai Maron, Shie Mannor, and Gal Chechik. 2021. Controlling Graph Dynamics with Reinforcement Learning and Graph Neural Networks. arXiv:2010.05313 [cs.LG] https://arxiv.org/abs/2010.05313

[19] Maria Nicola, Zaid Alsafi, Catrin Sohrabi, Ahmed Kerwan, Ahmed Al-Jabir, Christos Iosifidis, Maliha Agha, and Riaz Agha. 2020. The socio-economic implications of the coronavirus pandemic (COVID-19): A review. *International journal of surgery* 78 (2020), 185–193.

[20] Romualdo Pastor-Satorras and Alessandro Vespignani. 2001. Epidemic Spreading in Scale-Free Networks. *Phys. Rev. Lett.* 86 (Apr 2001), 3200–3203. Issue 14. https://doi.org/10.1103/PhysRevLett.86.3200

[21] Ryan A. Rossi and Nesreen K. Ahmed. 2014. NetworkRepository: An Interactive Data Repository with Multi-scale Visual Analytics. https://doi.org/10.48550/ARXIV.1410.3560

[22] Benedek Rozemberczki, Ryan Davies, Rik Sarkar, and Charles Sutton. 2018. GEMSEC: Graph Embedding with Self Clustering. (2018). https://doi.org/10.48550/ARXIV.1802.03997

[23] Prathyush Sambaturu, Bijaya Adhikari, B Aditya Prakash, Srinivasan Venkatramanan, and Anil Vullikanti. 2020. Designing effective and practical interventions to contain epidemics. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*.

[24] Sirui Song, Zefang Zong, Yong Li, Xue Liu, and Yang Yu. 2020. Reinforced Epidemic Control: Saving Both Lives and Economy. *Proceedings of KDD 2020* (2020).

[25] Shuchang Tao, Huawei Shen, Qi Cao, Liang Hou, and Xueqi Cheng. 2020. Adversarial Immunization for Certifiable Robustness on Graphs. *arXiv* (2020). https://doi.org/10.48550/ARXIV.2007.09647

[26] Guancheng Wan, Zewen Liu, Max SY Lau, B Aditya Prakash, and Wei Jin. 2024. Epidemiology-Aware Neural ODE with Continuous Disease Transmission Graph. *arXiv preprint arXiv:2410.00049* (2024).

[27] Xu Wang, Sen Wang, Xingxing Liang, Dawei Zhao, Jincai Huang, Xin Xu, Bin Dai, and Qiguang Miao. 2022. Deep reinforcement learning: A survey. *IEEE Transactions on Neural Networks and Learning Systems* 35, 4 (2022), 5064–5078.

[28] Duncan J. Watts and Steven H. Strogatz. 1998. Collective dynamics of 'small-world' networks. *Nature* 393, 6684 (June 1998), 440–442. https://doi.org/10.1038/30918

[29] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. 2020. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems* 32, 1 (2020), 4–24.

[30] Yue Xiang, Yonghong Jia, Linlin Chen, Lei Guo, Bizhen Shu, and Enshen Long. 2021. COVID-19 epidemic prediction and the impact of public health interventions: A review of COVID-19 epidemic models. *Infectious Disease Modelling* 6 (2021), 324–342.

[31] Jiaxuan Xie, Yezi Liu, and Yanning Shen. 2022. Explaining dynamic graph neural networks via relevance back-propagation. *arXiv preprint arXiv:2207.11175* (2022).

[32] Shenghao Yang, Priyabrata Senapati, Di Wang, Chris T Bauch, and Kimon Fountoulakis. 2021. Targeted pandemic containment through identifying local contact network bottlenecks. *PLoS computational biology* 17, 8 (2021), e1009351.

[33] Daniel Zügner, Amir Akbarnejad, and Stephan Günnemann. 2018. Adversarial attacks on neural networks for graph data. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. 2847–2856.

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009