Optimizing Harmful Information Containment in Social Networks via Shortest-Path Interdiction

Muon Ha[1(✉)], Yulia Shichkina[2], Tran Duc Manh[2], Xuan -Hien Nguyen[3], and Thanh-Son Phan[4]

[1] Telecommunications University, Khanh Hoa, Vietnam
muon.ha@mail.ru
[2] St. Petersburg State Electrotechnical University, St. Petersburg, Russia
[3] Hanoi University of Industry, Ha Noi, Vietnam
[4] Thai Binh Duong University, Nha Trang, Khanh Hoa, Vietnam

Abstract. The spread of malicious information on social networks is often explosive and difficult to control, especially in small-world structures and echo chambers. Traditional intervention strategies based on degree centrality are effective in scale-free networks; however, they often perform poorly in relatively homogeneous community networks, where bridge nodes may be more important than hubs. In this paper, we propose SPID (Shortest-Path Interdiction for Diffusion), a corridor-based mitigation framework that targets fast-spreading "diffusion corridors", defined as shortest paths connecting high-risk source regions to target communities. The experimental implementation uses SPID++, a residual-updated version of SPID that adapts to rerouted paths after each intervention step. Experiments on the Watts–Strogatz network show that SPID++ reduces the total number of infected nodes by approximately 2.2 times compared with Degree Centrality. Budget sensitivity analysis further shows that SPID++ can protect up to 444 more nodes than Degree as the intervention budget increases. Runtime results indicate that batched shortest-path computation reduces the intervention-selection time of SPID++ while preserving the selected intervention set.

Keywords: shortest path · shortest-path interdiction · harmful information diffusion · small-world networks · intervention optimization · SEIR model

1 Introduction

The development of social media has created an ideal environment for the rapid spread of information. Alongside its positive benefits, social media has also become fertile ground for harmful content, including fake news, online scams, and the manipulation of public opinion. The core danger of this type of information lies in its "chain reaction" spread mechanism: from a small group of initiating nodes (seeds), information can quickly spread through social media links, creating infection peaks and exceeding the control capabilities of manual censorship systems.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2027
O. Gervasi et al. (Eds.): ICCSA 2026 Workshops, LNCS 16763, pp. 403–414, 2027.
https://doi.org/10.1007/978-3-032-30530-5_25

In cybersecurity management, the core objective is not only simply to remove content, but also to optimize prevention strategies to simultaneously reduce infection peaks and minimize the total number of affected individuals with the lowest intervention costs. Current common approaches often rely on centrality metrics to identify critical nodes and prioritize interventions. In this context, a Degree Centrality-based strategy (intervening with KOLs/hubs) is often considered effective for scale-free networks (like Twitter), where a few high-level nodes can strongly influence the flow of information. However, reality shows that social network structures are not always hub-centric. In closed groups, professional forums, or local communities, the network structure is often small-world, where dissemination relies more on "shortcuts" connecting clusters than on a few prominent individuals. In such cases, Degree Centrality becomes less effective because it is difficult to correctly identify the nodes acting as "bridges" and the crucial "corridors" for dissemination (Fig. 1).

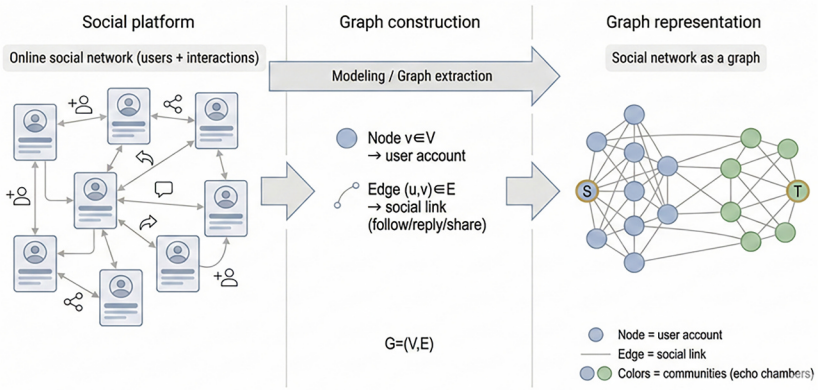

Fig. 1. Graph-based modeling of social media diffusion

A theoretically suitable approach to identify bridge nodes is Betweenness Centrality. However, this method faces significant computational obstacles when applied to large-scale graphs, making it impractical for networks with millions of nodes.

To fill this gap, the paper proposes a new method called SPID (Shortest-Path Interdiction for Diffusion). Instead of focusing on high-order nodes, SPID models the problem from a "path-blocking" perspective: finding and eliminating nodes on the shortest paths connecting the area of misinformation dissemination with vulnerable communities. This approach aims to cut off the fastest propagation routes, thereby reducing the spread of malicious information within the network.

The main contributions of the paper include: (i) proposing an algorithmic framework for node selection based on risk-aware path sampling and residual update mechanisms to adapt to changes in the information flow after intervention; (ii) using parallel shortest-path algorithms on large graphs as the core computation to address the performance issues encountered by Betweenness; and (iii) evaluating SPID on clustered network structures (Watts–Strogatz) and analyze budget-based sensitivity, thereby clarifying its potential application in scenarios where traditional methods perform poorly.

The paper is organized as follows. A review of related work and existing methods is presented in Sect. 2. Section 3 details the proposed parallel algorithm and its technical refinements. Experimental results obtained on large-scale datasets are reported and analyzed in Sect. 4. Finally, Sect. 5 concludes the paper and discusses possible avenues for future research.

2 Related Works

Recent research indicates that the rapid spread of misinformation is largely determined by the topology of the underlying network, such as community structures, echo chambers, and bridging nodes. These findings suggest that network-based models provide more accurate representations than traditional well-mixed population assumptions. Empirical data further demonstrate that fake news often diffuses faster and reaches a wider audience than factual information [1, 2]. Additionally, social bots act as significant amplifiers of untrustworthy content, particularly when such information aligns with existing community biases or current trends [3]. Furthermore, the "complex contagion" theory emphasizes the role of repeated exposure in clustered neighborhoods, increasing the probability of acceptance/sharing and contributing to widely spread dynamics [4].

To describe behavioral latency, SEIR-type models are often used because the Exposed state represents the "received but not shared" situation. This choice is suitable for modern studies that combine propagation with social context and temporal evolution. In terms of data, FakeNewsNet integrates content, social context, and spatiotemporal signals, supporting simultaneous analysis of propagation, detection, and intervention [7]. In terms of detection, surveys have systematized feature groups (content, users, propagation) and method families (classical machine learning, deep learning, multimodal learning, and explainable models) [5, 6]. For rumors, specialized reviews highlight the lack of early-stage data, label uncertainty, and the importance of conversational signals and community responses [8]. More recently, detection systems are able to interpret and exploit user content and feedback to provide "evidence" for decision-making have also been proposed [5, 6].

Mitigation methods can be broadly divided into "soft" and "hard" interventions. Soft interventions – for example, "accuracy nudges" – improve the quality of sharing without directly blocking content [9], while prebunking/inoculation increases resistance to manipulation techniques, even in the context of COVID-19 misinformation [10]. Community-based strategies also apply crowdsourcing to integrate detection and mitigation through collective labeling and feedback [11]. In contrast, hard interventions often model the intervention site selection problem as a network optimization problem under budget constraints, aiming to reduce exposure or spread intensity. Typical approaches involve point process-based intervention models [13], propagation control under diffusion processes [12], and recent extensions incorporating fairness/social constraints [14], time-based penalties under heterogeneous propagation rates [15], or adaptive selection of correctors through learning-to-intervene methods [16].

From a network optimization perspective, the network interdiction problem provides a bilevel framework: the upper layer selects nodes/edges to break or degrade under budget constraints, and the lower layer optimizes a network goal on the post-intervention graph

(e.g., shortest path or maximum flow). A recent survey has systematized interdiction models and algorithms, in which the shortest-path interdiction is a classic and widely studied case [17]. During the period 2016–2025, shortest-path interdiction has been expanded in directions such as arc improvements and structural analysis [18], settings with uncertainty and information asymmetry [19], incomplete information regimes [20], and policy learning approaches to increase scalability [21]. Notably, misinformation control has been modeled as a two-tiered stochastic interdiction problem to minimize expected reach to vulnerable groups [22], thereby reinforcing the link between breaking structural "corridors" and limiting propagation.

Finally, shortest-path/interdiction-based strategies often require iterative shortest-path computations across multiple post-intervention graph instances that make computational efficiency a key factor. Recent acceleration efforts leverage graph-algebra frameworks (e.g., GraphBLAS) to implement efficient SSSP on parallel hardware (e.g., delta-stepping) [23]. Complementary to this direction, an SSSP method tailored to graphs with cycles is introduced in [24] and subsequently extended to parallel execution on large graphs in [25], providing a practical implementation basis for shortest-path-based interventions.

3 Methodology

3.1 Problem Statement

Consider a directed graph $G = (V, E)$, where V is the set of nodes (e.g., user accounts) and E is the set of edges (e.g., interactions). Let $S_0 \subseteq V$ denote the initial set of sources that disseminate harmful information. Diffusion follows an SEIR (Susceptible–Exposed–Infected–Recovered) process with parameters $\Theta = \{\beta, \kappa, \gamma\}$, representing the exposure, infection, and recovery rates, respectively.

An intervention strategy selects a set I of elements to disable subject to a budget constraint: the total intervention cost over I does not exceed budget B, where each element x (node or edge) has cost $c(x)$. We consider two intervention modes:

- **Node interdiction**: $I \subseteq V$ (e.g., suspending accounts to reduce dissemination capacity).
- **Edge interdiction:** $I \subseteq E$ (e.g., weakening or removing transmission interactions).

After intervention, the network becomes $G' = (V', E')$. Motivated by shortest-path interdiction [17, 18, 22], effective interventions disrupt fast "diffusion corridors" by increasing distances from sources to target communities or, equivalently, reducing reachability.

We formulate the problem with one of two primary objectives:

Structural objective: maximize the aggregate shortest-path distance from sources S_0 to a target set T_0 (optionally via a monotone distance-based utility, with weights if needed).

Dynamical objective: minimize the number of infected nodes at equilibrium under SEIR simulation on G'.

Recent studies connect these views through interdiction-based diffusion control, modeling reachability or dynamics under budget constraints [17, 18, 22]. Section 3 details how shortest-path structure can be exploited to reduce SEIR spread.

3.2 The Proposed SPID Method

SPID (Shortest-Path Interdiction for Diffusion) selects an intervention set I (nodes or edges) under a finite budget B to reduce diffusion in SEIR simulations. Unlike approaches that prioritize globally central nodes via static centrality measures, SPID targets "diffusion corridors" approximated by shortest paths from sources S_0 to targets T_0. Figure 2 illustrates the simulated social network and SPID's core intuition: interdiction of shortest-path corridors.

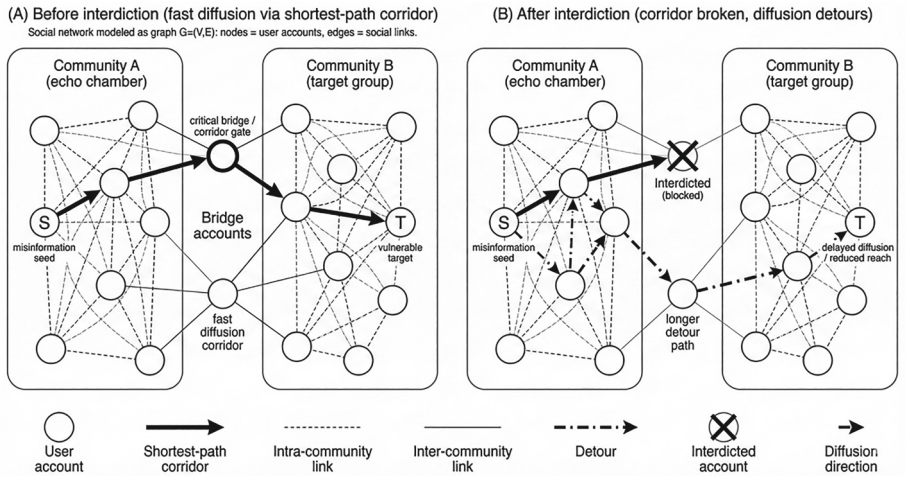

Fig. 2. Social network graph simulation and SPID intuition: shortest-path corridor interdiction

From the formulation above, the ultimate goal is to reduce SEIR propagation on the post-intervention graph. Directly optimizing SEIR outcomes is expensive because it typically requires many running stochastic simulations. SPID therefore optimizes a structural surrogate: by breaking short source-to-target routes, diffusion is forced onto longer detours, reducing both speed and reach.

Given G, S_0, T_0, and budget B, SPID samples K representative source–target pairs $P = \{(s_i, t_i)\}_{i=1}^{K}$. Sampling can be performed in two modes:

Uniform: sample uniformly at random from $S_0 \times T_0$.
Risk-aware: sample according to weights $r(s)$ (e.g., source out-degree) and $v(t)$ (e.g., target community size). When such metadata are unavailable, uniform sampling is used. For each pair $(s, t) \in P$, SPID computes the shortest-path set $\mathrm{SP}(s, t)$ (or a single representative shortest path). Each candidate node v (excluding sources and targets) receives a score based on how frequently it appears on these paths, multiplied by a priority weight that favors earlier interdiction near the source (e.g., a decreasing function of distance from s).
To deal with rerouting after interdiction, we use SPID $++$, a greedy version of SPID with selective recalculation: at each step, it recomputes shortest paths only for the sampled pairs affected by the newly removed element, reducing unnecessary work.

Algorithm 1: SPID++

Input: $G = (V,E)$, S_0, T_0, budget B, sample size K.
Output: intervention node set U.

1. Sample K pairs $P = \{(s,t)\}$ from $S_0 \times T_0$ (uniform or risk-aware).
2. For each $(s,t) \in P$, compute $\mathrm{SP}(s,t)$ and initialize node scores.
3. $U \leftarrow \varnothing$.
4. For $b = 1,\ldots,B$:
   a. Select $u \in V \backslash U$ with the highest score; add u to U; remove u from G.
   b. Identify affected pairs $P_u = \{(s,t) \in P : u \text{ lies on the current path for } (s,t)\}$.
   c. Recompute $\mathrm{SP}(s,t)$ only for $(s,t) \in P_u$ on the updated graph.
   d. Update scores by subtracting the previous contributions and adding the new contributions induced by P_u.
5. Return U.

The computational bottleneck of SPID is repeated shortest-path computation and updates across many pairs. SPID addresses this by integrating an efficient SP-Kernel based on a blocked adjacency-matrix formulation for parallelization [24], further extended to multi-core execution with workload partitioning and synchronization for large graphs [25].

SP-Kernel supports batched initial shortest-path computation and selective updates for affected pairs, balancing interdiction effectiveness with practical runtime on large-scale networks. Initialization scales with K shortest-path queries. Each greedy iteration depends on the number of affected pairs (typically much smaller than K). In experiments, we use $K = 5000$, a distance-decay weight with $\alpha = 1$, and budgets $B = 10$–60 for sensitivity analysis over budget levels.

4 Experiments and Evaluation

This section describes the datasets, baselines, SEIR simulation settings, and evaluation protocol used to validate SPID under two different network regimes: hub-dominated and clustered small-world structures.

4.1 Datasets

The two networks considered here have different topological features for the analysis of the efficacy of the intervention strategies proposed. They are a real-world network and a synthetic network.

- **PolitiFact (FakeNewsNet):** The network used comes from the FakeNewsNet repository [7]. The network is characterized by scale-free properties, which means that the network follows a power-law degree distribution and has high-degree nodes. After

preprocessing, there are 2,500 nodes and 7,442 edges. Since scale-free networks rely heavily on their hubs in terms of diffusion of information, this network provides useful data to evaluate interventions that target hubs and analyze the behavior of SPID when the network is dominated by hubs.

- **Watts-Strogatz (Synthetic):** This is a synthetic network that is created using the Watts-Strogatz graph (N = 2,000). The graph represents a small-world network where the clustering coefficient is high, which makes the transmission process depend heavily on bridging edges.

4.2 Baselines and Parameter Settings

We evaluate SPID + + as the main implementation of the proposed SPID framework against four baselines: No Intervention, Random selection, Degree Centrality, and Betweenness Centrality. Unless otherwise noted, all methods are tested under the common configuration reported in Table 1.

Table 1. Experimental parameters.

Parameter group	Symbol	Value
SEIR	β	0.05
SEIR	δ (E $\rightarrow$ I)	0.25
SEIR	γ (I $\rightarrow$ R)	0.10
SEIR	T (steps)	30
SPID	K (sample pairs)	5000
SPID	α (distance weight)	1
Intervention	B (budget)	{10,20,30,40,50,60}
Sources/targets	$\lvert S_0 \rvert$	10 ($\approx$0.5% of node)
Sources/targets	T_0	$V \setminus S_0$
Monte Carlo	R	20

We chose $K = 5000$ to balance node-score stability and running time while keeping the SEIR parameters fixed across all methods to isolate the effect of the intervention strategy. Preliminary tests showed that smaller values of K led to unstable node rankings, while larger values only slightly changed the top-ranked nodes but significantly increased the shortest-path computation. Therefore, K functions primarily as a sampling-cost parameter rather than a diffusion parameter.

4.3 Evaluation Metrics

We evaluate mitigation performance using two measures that characterize outbreak severity and diffusion extent.

- **Peak Infected (Peak):** the maximum number of nodes in state I at any time during the simulation. This measure indicates the highest outbreak intensity and is closely related to short-term control pressure.
- **Total Infected:** the number of nodes that ever become infected during the simulation, corresponding to the cumulative affected population. This measure captures the final reach of the diffusion process.

Results for both metrics are presented across intervention budgets and network topologies.

4.4 Experimental Protocol

For each experimental setting – defined by the dataset, SEIR parameters, intervention budget B, and strategy – we run the stochastic SEIR simulation 20 times and report the average results. Repeating the simulation in this way helps reduce stochastic variation and keeps the comparison fair across strategies. In addition to comparing methods at a fixed budget, we examine budget sensitivity by varying B over a range of values. This analysis allows us to study how each strategy benefits from additional intervention resources and whether SPID scales more effectively than Degree Centrality or Betweenness Centrality across different network topologies.

5 Results and Discussion

This section reports the performance of SPID ++, the residual-updated implementation of SPID, in comparison with the baseline methods. We focus first on mitigation effectiveness in small-world networks represented by the Watts-Strogatz model, and then examine how performance changes as the intervention budget increases. The results are subsequently compared with those obtained on the real PolitiFact network in order to highlight how different network structures influence strategy behavior.

5.1 Performance on Small-World Networks (Watts–Strogatz)

In Watts-Strogatz networks, node degrees are relatively homogeneous, whereas clustering remains high. Under these conditions, cross-community diffusion is often driven by shortcut links and bridge nodes or edges. As a result, strategies based solely on degree may not identify the most influential transmission routes. By contrast, SPID is designed to target diffusion corridors approximated through shortest paths, which may provide an advantage in this type of topology.

Table 2 compares Peak and Total Infected at the same intervention budget (Table 3 and Fig. 3).

Random yields little improvement and, in some runs, slightly increases Total Infected, likely because the selected nodes vary substantially in importance. Degree Centrality reduces Total Infected by 4.3%, while SPID++ achieves a larger reduction of 9.5%, corresponding to about 2.2 times the improvement obtained by Degree in the small-world setting. Betweenness Centrality also performs reasonably well, with a

Table 2. Comparison of mitigation performance on the Watts-Strogatz network.

Strategy	Peak	Total Infected	% Reduction (Total)
No Intervention	397.6	1741.0	0.0%
Random	395.6	1764.0	-1.3%
Degree	375.4	1666.8	4.3%
Betweenness	360.9	1608.5	7.6%
SPID++ (Ours)	**351.9**	**1575.2**	**9.5%**

Table 3. Indicative CPU running time of intervention selection at $B = 60$.

Strategy	PolitiFact	Watts–Strogatz
Degree Centrality	0.01 s	0.01 s
Betweenness Centrality	1.42 s	1.18 s
SPID++ without batching	0.76 s	0.84 s
SPID++ with batching	0.41 s	0.46 s

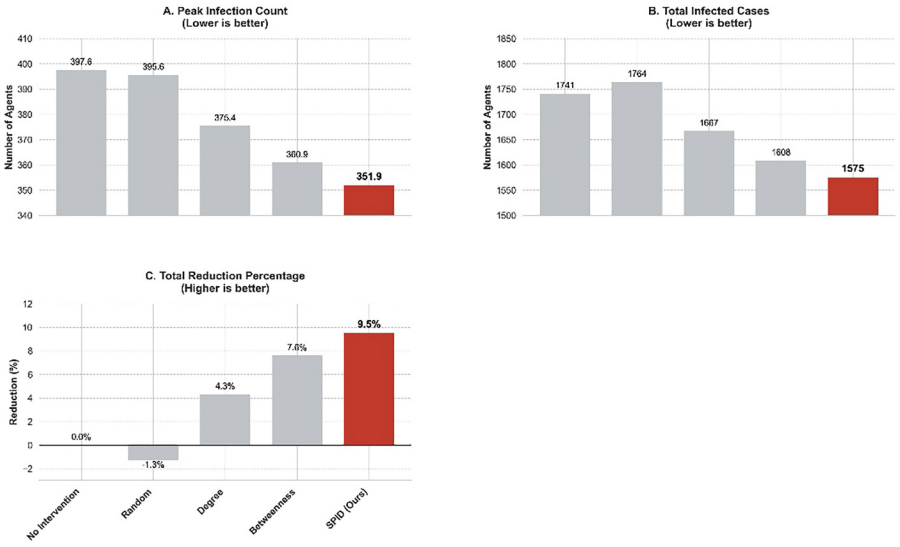

Fig. 3. Comparison of Mitigation Strategies on Watts-Strogatz Networks

reduction of 7.6%, which is consistent with its sensitivity to bridging structure. Even so, SPID++ performs best, likely because it targets diffusion corridors through source–target shortest-path sampling and updates the intervention set after each step.

5.2 Budget Sensitivity and Scaling Behavior

We further study budget sensitivity by varying the intervention budget B and measuring Total Infected for each method. Two consistent trends are observed. Degree Centrality saturates early: as B increases, the incremental reduction in Total Infected becomes smaller, suggesting that the most influential high-degree nodes are selected early and that further removals contribute relatively little to reducing inter-community transmission. SPID++, by contrast, continues to benefit from larger budgets. Its residual updating allows the method to adapt to rerouted diffusion paths after each intervention step, so additional budget can disrupt newly emerging transmission corridors.

At $B = 60$, SPID++ protects 444 more nodes than Degree Centrality, meaning that Total Infected is lower by 444. This result indicates that, in clustered networks, larger budgets are used more effectively when the intervention strategy targets diffusion corridors rather than relying only on node degree. We also report indicative CPU running times for the intervention selection step at the same budget level, where shortest-path computation is the main cost of SPID. Batching reduces the running time of SPID++ while keeping the selected intervention set unchanged.

These results show that SPID++ is more expensive than Degree Centrality, but remains faster than Betweenness Centrality; the batched version gives about a 1.8 $\times$ speedup on both tested graphs.

5.3 Comparison on the Real PolitiFact Network (Hub-Dominated)

On the PolitiFact network, which exhibits the hub-dominated characteristics typical of a scale-free topology, Degree Centrality is particularly effective at reducing Peak Infected, achieving a reduction of about 99%, whereas SPID++ achieves about 88%. This is consistent with scale-free intuition: disabling hubs can sharply suppress early-stage dissemination capacity.

Nevertheless, SPID remains highly effective and practically relevant in settings where direct hub intervention is infeasible (e.g., policy constraints, accounts that cannot be removed, or requirements for less intrusive actions). In such cases, SPID's corridor-interdiction mechanism serves as a complementary alternative: it reduces spread by cutting short routes from the source region to other clusters, rather than depending entirely on hub removal.

6 Conclusion and Limitations

We propose SPID, a corridor-based framework for mitigating harmful-information diffusion in social networks. SPID selects intervention nodes according to their coverage of shortest paths from source regions to target communities. In the experiments, we use SPID++, a residual-updated implementation of SPID that adapts to alternative routes after each intervention step.

Experiments show that SPID++ is particularly effective on high-clustering small-world networks. On Watts–Strogatz graphs, SPID++ reduces Total Infected more than Degree Centrality and maintains a clear advantage as the budget increases, with the

performance gap widening at higher budgets. On the real PolitiFact network, Degree Centrality is stronger in Peak reduction due to hub dominance, yet SPID++ still provides high mitigation and remains useful as a non-hub-dependent intervention strategy. For deployment, SPID++ can be paired with an efficient shortest-path kernel that supports batched shortest-path queries and improves feasibility on larger graphs.

This study assumes a static network per simulation episode, known initial sources, medium-sized graphs, and node-based interventions under a fixed budget. Future work will extend SPID to larger and dynamic graphs, partially observed source sets, and soft interventions, where transmission is attenuated rather than nodes or edges being fully removed.

Acknowledgment. This work was supported by the Ministry of Science and Higher Education (assignment No. FSEE-2025-0015).

References

1. Vosoughi, S., Roy, D., Aral, S.: The spread of true and false news online. Science **359**(6380), 1146–1151 (2018). https://doi.org/10.1126/science.aap9559
2. Del Vicario, M., Bessi, A., Zollo, F., et al.: The spreading of misinformation online. Proc. Natl. Acad. Sci. U.S.A. **113**(3), 554–559 (2016). https://doi.org/10.1073/pnas.1517441113
3. Shao, C., Ciampaglia, G.L., Varol, O., Yang, K.-C., Flammini, A., Menczer, F.: The spread of low-credibility content by social bots. Nat. Commun. **9**, 4787 (2018). https://doi.org/10.1038/s41467-018-06930-7
4. Törnberg, P.: Echo chambers and viral misinformation: modeling fake news as complex contagion. PLoS ONE **13**(9), e0203958 (2018). https://doi.org/10.1371/journal.pone.0203958
5. Shu, K., Sliva, A., Wang, S., Tang, J., Liu, H.: Fake news detection on social media: a data mining perspective. SIGKDD Explor. Newsl. **19**(1), 22–36 (2017). https://doi.org/10.1145/3137597.3137600
6. Zhou, X., Zafarani, R.: A survey of fake news: fundamental theories, detection methods, and opportunities. ACM Comput. Surv. **53**(5), Article 109 (2020). https://doi.org/10.1145/3395046
7. Shu, K., Mahudeswaran, D., Wang, S., Lee, D., Liu, H.: FakeNewsNet: a data repository with news content, social context, and spatiotemporal information for studying fake news on social media. Big Data **8**(3), 171–188 (2020). https://doi.org/10.1089/big.2020.0062
8. Zubiaga, A., Aker, A., Bontcheva, K., Liakata, M., Procter, R.: Detection and resolution of rumours in social media: a survey. ACM Comput. Surv. **51**(2), Article 32 (2018). https://doi.org/10.1145/3161603
9. Pennycook, G., McPhetres, J., Zhang, Y., Lu, J.G., Rand, D.G.: Fighting COVID-19 misinformation on social media: experimental evidence for a scalable accuracy-nudge intervention. Psychol. Sci. **31**(7), 770–780 (2020). https://doi.org/10.1177/0956797620939054
10. Basol, M., Roozenbeek, J., van der Linden, S.: Towards psychological herd immunity: cross-cultural evidence for two prebunking interventions against COVID-19 misinformation. Big Data Soc. **8**(1) (2021). https://doi.org/10.1177/20539517211013868
11. Kim, J., Tabibian, B., Oh, A., Schölkopf, B., Gomez-Rodriguez, M.: Leveraging the crowd to detect and reduce the spread of fake news and misinformation. In: Proceedings of the 11th ACM International Conference on Web Search and Data Mining (WSDM 2018), pp. 324–332. ACM (2018). https://doi.org/10.1145/3159652.3159734

12. Tong, G., Wu, W., Du, D.-Z.: On misinformation containment in online social networks. In: Advances in Neural Information Processing Systems (NeurIPS 2018). arXiv:1809.06486 (2018)
13. Farajtabar, M., et al.: Fake news mitigation via point process based intervention. In: Proceedings of the 34th International Conference on Machine Learning (ICML 2017), Proceedings of Machine Learning Research, vol. 70, pp. 1097–1106. PMLR (2017)
14. Abouzeid, A., et al.: Socially fair mitigation of misinformation on social networks via constraint stochastic optimization. In: *Proc. AAAI Conf. on Artificial Intelligence* 36(11), 11801–11809 (2022). https://doi.org/10.1609/aaai.v36i11.21436
15. Simpson, M., Hashemi, F., Lakshmanan, L.V.S.: Misinformation mitigation under differential propagation rates and temporal penalties. Proc. VLDB Endow. **15**(10), 2216–2229 (2022). https://doi.org/10.14778/3547305.3547324
16. Xu, L., et al.: Selecting debunkers via self-imitation learning. In: Proceedings of the AAAI Conference on Artificial Intelligence **38**(20), 22447–22456 (2024). https://doi.org/10.1609/aaai.v38i20.30252
17. Smith, J.C., Song, Y.: A survey of network interdiction models and algorithms. Eur. J. Oper. Res. **283**(3), 797–811 (2020). https://doi.org/10.1016/j.ejor.2019.10.030
18. Holzmann, T., Smith, J.C.: Shortest-path interdiction problem with arc improvement recourse: a multiobjective approach. Nav. Res. Logist. **66**(3), 230–252 (2019). https://doi.org/10.1002/nav.21839
19. Punla-Green, S.Z., et al.: Shortest path network interdiction with asymmetric uncertainty. Networks **85**(5), 886–905 (2024). https://doi.org/10.1002/net.22208
20. Azizi, E., Seifi, A.: Shortest path network interdiction with incomplete information: a robust optimization approach. Ann. Oper. Res. **335**(2), 727–759 (2024). https://doi.org/10.1007/s10479-023-05350-1
21. Zhang, L., Chen, Z., Lu, C.-T., Zhao, L.: Network interdiction goes neural. arXiv:2405.16409 (2024)
22. Moug, K., Shen, S.: Stochastic bilevel interdiction for fake news control in online social networks. Comput. Oper. Res. **173**, 106872 (2025). https://doi.org/10.1016/j.cor.2024.106872
23. Sridhar, U., Blanco, M., Mayuranath, R., Spampinato, D.G., Low, T.M., McMillan, S.: Delta-stepping SSSP: from vertices and edges to GraphBLAS implementations. In: Proceedings of the 2019 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW), pp. 241–250. IEEE (2019). https://doi.org/10.1109/IPDPSW.2019.00047
24. Shichkina, Y., Nguyen, XH., Ha, M., Tran, D.M.: Shortest path search method on a graph with cycles. In: Gervasi, O., Murgante, B., Garau, C., Taniar, D., C. Rocha, A.M.A., Faginas Lago, M.N. (eds.) Computational Science and Its Applications – ICCSA 2024 Workshops. ICCSA 2024. Lecture Notes in Computer Science, vol. 14825, pp. 344–356. Springer, Cham (2024). https://doi.org/10.1007/978-3-031-65343-8_26
25. Ha, M., Shichkina, Y., Nguyen, X.H.: Extending the shortest path algorithm for large graphs with cycles and parallel computing capabilities. In: Gervasi, O., et al. Computational Science and Its Applications – ICCSA 2025 Workshops. ICCSA 2025. Lecture Notes in Computer Science, vol. 15894, pp. 191–204. Springer, Cham (2026). https://doi.org/10.1007/978-3-031-97648-3_13