# Community-Aware Influence Maximization for Suppressing Cryptocurrency Scam Misinformation

Naglaa Mostafa, Hatem Abdelkader, Asmaa H.Ali
Information System Department-Faculty of Computers and Information, Menoufia University, Egypt

***Abstract*—Cryptocurrency fraud campaigns often rely on large-scale social-media diffusion to recruit victims, normalize false claims, and coordinate multi-level marketing behavior. This study examines the dynamics of the One Coin scam. It proposes an influence-maximization (IM)-driven workflow for identifying high-impact accounts whose intervention can reduce future misinformation diffusion. A directed Twitter engagement network from retweet/reply interactions is constructed and studied, and the accounts that should be prioritized for intervention to reduce the reach of future scam-promoting misinformation are identified. We evaluate six seed selection strategies: Degree, Betweenness, PageRank, k-core, CELF (lazy greedy), and Reverse Influence Sampling (RIS) under the classical Independent Cascade (IC) and Linear Threshold (LT) diffusion models using a weighted-cascade parameterization when ground-truth transmission probabilities are unavailable. Across the tested seed budgets, CELF achieves the highest expected spread, but with the highest computational cost. At the largest seed budget, Degree is effectively tied with CELF (within 0.09% under LT and 1.4% under IC), indicating a hub-dominated engagement structure in which simple reach-based heuristics can be highly competitive. RIS provides a strong quality–efficiency trade-off, remaining within approximately 9.7% (LT) and 9.5% (IC) of CELF while requiring substantially less computation. We further introduce a community-aware variant using Leiden partitions and proportional seed allocation to improve cross-community coverage; at larger budgets, this improves methods sensitive to seed over-concentration, increasing LT spread by about 9.8% for k-core and 8.6% for RIS. Overall, the results quantify practical trade-offs between spread and runtime for deployable suppression workflows and show when community-aware planning better aligns with the heterogeneous structure of scam recruitment ecosystems.**

***Keywords*—*Influence maximization; misinformation suppression; cryptocurrency scams; OneCoin; diffusion models; Leiden; community-aware seeding***

## I. Introduction

Online social networks have become a primary channel for disseminating both legitimate information and harmful content. Fraudulent investment narratives are especially sensitive to viral diffusion: repeated exposure can create perceived legitimacy, recruit new participants, and sustain the financial incentives of a scam ecosystem. OneCoin, a cryptocurrency launched in 2014, has been described by U.S. prosecutors as a multibillion-dollar global fraud scheme, with prosecutions and public warnings continuing over several years (U.S. Department of Justice [1][2][3]. From a network-science perspective, such campaigns can be studied as diffusion processes on an interaction graph, where certain accounts play outsized roles in initiating and amplifying narratives.

A key practical question is: Which accounts should be prioritized for intervention to most effectively curb the spread of scam-promoting misinformation? Intervention may include platform trust-and-safety actions, targeted fact-checking prompts, or counter-messaging allocation. This problem relates to influence maximization (IM): given a graph and a diffusion model, IM selects a small seed set of nodes (accounts to which an intervention or corrective exposure is applied) to maximize the expected diffusion reach [4]. In misinformation suppression settings, these seeds represent the most leverageable points for disruption or for seeding corrective information, aligning conceptually with competing-cascade work on limiting misinformation [5]. We focus on OneCoin because it is a well-documented, large-scale fraud campaign with extensive public enforcement records and warnings, making it a realistic case for studying diffusion-driven recruitment dynamics. This work provides the following contributions:

- A reproducible workflow to build a directed OneCoin interaction network from Twitter engagement signals (retweets and replies), and to model propagation under IC and LT diffusion [4].
- A comparative experimental study of classic heuristics (Degree, Betweenness, PageRank, k-core), CELF-accelerated greedy selection [6], and Reverse Influence Sampling (RIS) [7] on the same dataset and seed budgets.
- A community-aware variant using Leiden partitions [8], with proportional seed allocation across communities, demonstrates how community constraints affect diffusion and runtime.

The remainder of this study is organized as follows: Section II will cover the related work and background. In Section III, the conceptual pipeline of the proposed OneCoin misinformation diffusion study is detailed. Section IV presents the experiments and evaluations, and we conclude in Section V.

## II. Related Work

### A. Influence Maximization and Scalable Algorithms

Kempe et al. formalized IM under the IC and LT models and showed that expected spread is monotone and submodular, enabling a $(1-1/e)$ approximation via greedy selection [4]. Because naive greedy requires many spread estimations, Leskovec et al. proposed CELF (Cost-Effective Lazy Forward), which reduces recomputation of marginal gains while preserving greedy quality [6]. To scale IM to large graphs, RIS-type methods estimate influence by sampling reverse-reachable (RR) sets, yielding near-linear-time algorithms with provable approximation guarantees [7], [22]. Recent surveys summarize these algorithm families and extensions to dynamic and competitive settings [10]–[12].

### B. Misinformation Diffusion and Influence Blocking

A central line of work models misinformation containment as a competitive diffusion problem, where a 'good' campaign (e.g., corrective information) competes with a 'bad' campaign (misinformation) [5]. Subsequent studies have proposed variations that incorporate awareness raising, time constraints, refutation, and vertex blocking, yielding influence-blocking maximization objectives and algorithms [13]–[15]. While our experiments evaluate classical IM under IC/LT to characterize which accounts can most amplify diffusion, the intervention interpretation is consistent with this containment literature: identifying a small set of high-leverage accounts where friction, monitoring, or corrective seeding can meaningfully reduce cascade potential.

### C. Community Structure and Community-Aware Influence

Maximization. Diffusion often concentrates within dense communities and crosses communities via a smaller set of bridging accounts, motivating community-aware IM that allocates seeds across communities to improve coverage and diversity [16]–[18]. Reliable community detection is therefore important when community boundaries constrain seeding decisions. Leiden improves upon Louvain by refining partitions to ensure well-connected communities and often achieves higher-quality modularity solutions efficiently [8]. In this study, we use Leiden as a principled preprocessing step and evaluate how proportional community allocation changes both spread and runtime in a real scam-engagement network.

Positioning of this study. Prior work has extensively developed IM algorithms and misinformation containment models, but fewer studies provide end-to-end, reproducible evaluations of how classic IM baselines, scalable RIS methods, and community-aware constraints trade off effectiveness and runtime in a documented cryptocurrency scam setting. Our contribution is to ground this comparison in a OneCoin engagement network and to quantify when community-aware allocation improves cross-community coverage and, at sufficient budgets, overall diffusion reach.

## III. Methodology

This study investigates how OneCoin scam-related misinformation can propagate through a Twitter interaction network. It demonstrates how influence maximization [9] can be used to identify a small set of high-influence accounts. To address misinformation propagation in social networks, this study combines information diffusion models, influence evaluation, and influence maximization (see Fig. 1). A directed, weighted engagement graph is constructed from retweet/reply links and analyzed under two widely used diffusion models, Independent Cascade (IC) and Linear Threshold (LT), with a parameter-free weighted-cascade setting based on each node's in-degree, enabling principled propagation simulation when accurate transmission probabilities are unavailable. Significantly, integrating Leiden community detection and enforcing community-proportional seed allocation changes the seed composition toward broader sub-community coverage. Six seed-selection strategies (Degree, Betweenness, PageRank, k-core, CELF, and RIS) were used across multiple seed budgets to achieve the most significant expected spread consistent with submodular maximization theory, but at substantially higher computational cost. Also, we show that, in influence maximization methods, it is essential to provide a strong spread-time compromise suitable for scalable analyses, and that simple heuristics remain competitive across several budgets due to the network's sparse, hub-dominated engagement structure.

### A. Data Collection and Case Context (OneCoin Scam Misinformation)

In this subsection, we represent the input evidence used to study diffusion: OneCoin-related Twitter content and engagement traces (e.g., tweets promoting the scam narrative and retweets/replies that amplify it). In misinformation-suppression framing, the goal is not only to describe content, but to capture who amplifies whom and how narratives propagate through user-to-user interactions. In practice, this stage yields: 1) a set of users, 2) interaction events (retweet/reply), and 3) timestamps and metadata that can later support time-sliced diffusion analyses. The OneCoin case is well documented as a large fraud scheme by official authorities, motivating a realistic misinformation setting for network-based interventions. Operationally, the data layer comprises OneCoin-related social-media traces (e.g., posts and engagements). In this work, the emphasis is on engagement events (retweets/replies) because they are direct, observable signals of information amplification and conversational diffusion, enabling network-based modeling.

### B. Interaction-Network Construction (Twitter Engagement Graph)

This subsection maps raw engagement events into a directed graph suitable for diffusion modeling. We define a directed weighted interaction network $G = (V, E, W)$, where each node $v \in V$ represents a user account and each directed edge $(u, v) \in E$ indicates that user $u$ interacted with content produced by user $v$ (retweet/reply). Edge weight $w_{uv}$ counts repeated engagements to capture the intensity of interaction. This representation is widely used because diffusion and influence processes can be simulated on graphs and optimized over node sets [4]. The resulting structure is typically sparse and exhibits heterogeneous degree distributions, which motivates comparing both optimization-based methods (e.g., greedy/CELF, RIS) and centrality-based heuristics that may perform well in hub-dominated graphs.

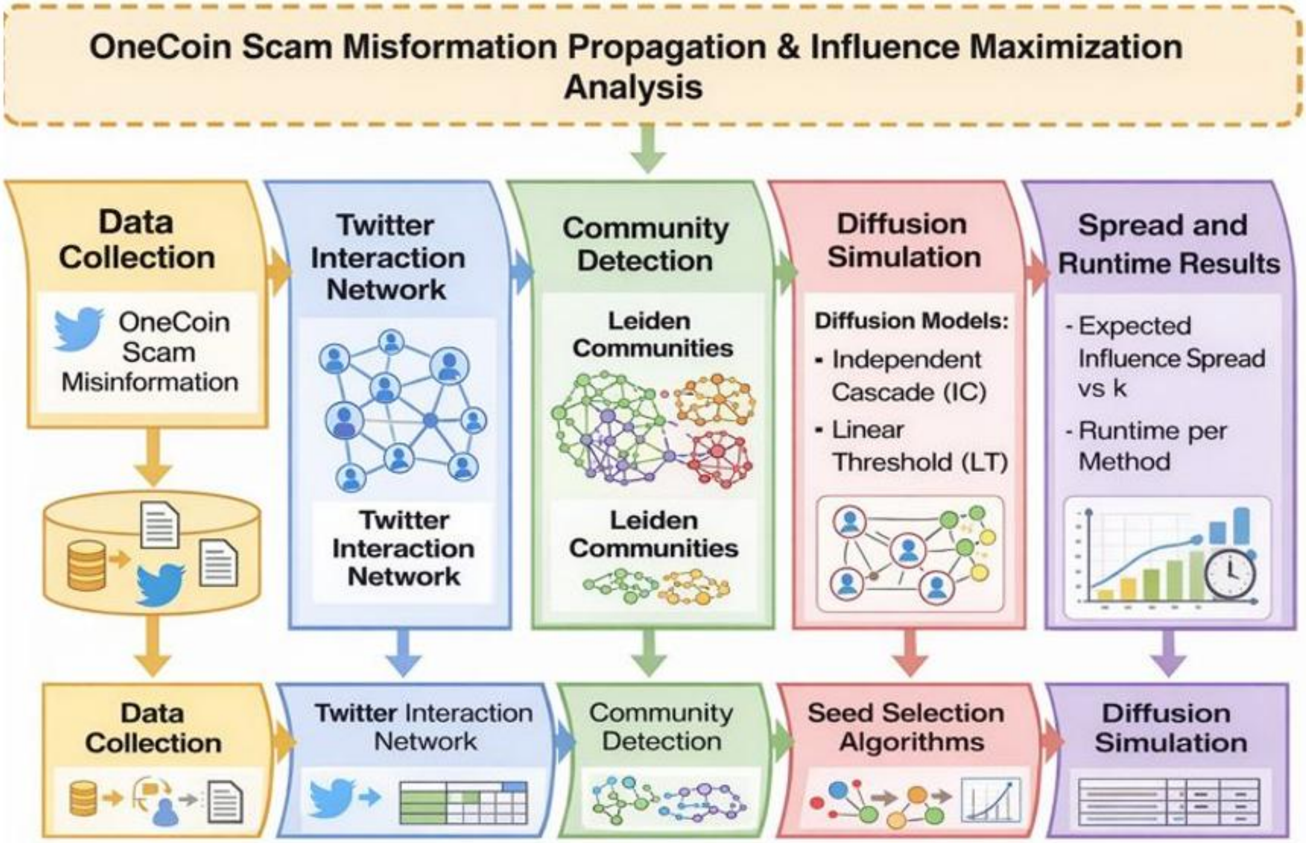


Fig. 1. Conceptual pipeline of the proposed OneCoin misinformation diffusion study.

### C. Community Detection (Leiden Partitioning)

Community structure is central to misinformation diffusion because amplification frequently occurs within dense clusters, while a smaller number of accounts facilitates cross-community transmission. To capture this structure, we apply the Leiden algorithm in Fig. 2 to obtain a partition $\mathcal{C} = \{C_1, \dots, C_r\}$. Leiden is specifically chosen because it addresses known weaknesses of Louvain by improving community connectivity and providing stronger guarantees about partition quality (well-connected communities), while remaining efficient [8]. In the proposed research problem, communities represent sub-audiences or recruitment circles; identifying them supports targeted suppression strategies that aim to reduce cascade potential across the ecosystem rather than within only one dense group.

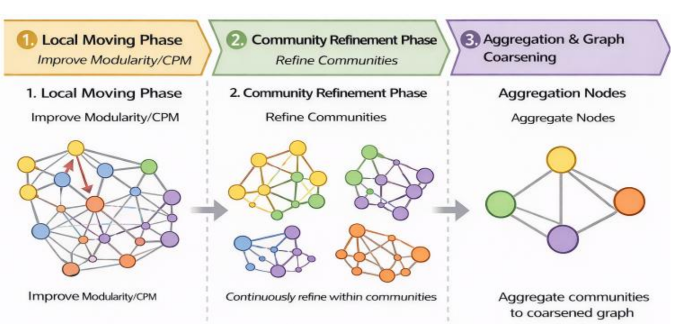


Fig. 2. Community detection of the Leiden algorithm.

### D. Diffusion Models for Misinformation Propagation

This subsection formalizes how misinformation may spread in the constructed graph. Independent Cascade (IC). Starting from a seed set $S$, diffusion proceeds in discrete steps. When node $u$becomes active, it gets a single opportunity to activate each inactive neighbor $v$with probability $p_{uv}$. Eq. (1) clearly states the influence (spread) function as:

$$\sigma(S) = \mathbb{E}[|\, A(S) \,|] \tag{1}$$

where, $A(S)$ is the final set of activated nodes[4].Linear Threshold (LT). Each node $v$samples a threshold. $\theta_v \in [0,1]$. Node $v$activates when the cumulative influence of its active in-neighbors reaches the threshold [see Eq. (2)]:

$$\textstyle\sum_{u \in N^{-}(v) \cap A} w_{uv} \geq \theta_v. \tag{2}$$

LT is frequently interpreted as "accumulated exposure" or social reinforcement, which is conceptually appropriate for the persuasion-like adoption of misinformation [4]. Weighted-cascade parameterization. When accurate transmission probabilities are unavailable, a standard practical setting, as mentioned in Eq. (3) and Eq. (4), is:

$$p_{uv} = \frac{1}{\deg^{-}(v)} (\text{IC}) \tag{3}$$

$$w_{uv} = \frac{1}{\deg^{-}(v)} (\text{LT}) \tag{4}$$

This equalizes influence contribution across in-neighbors of each node.

### E. Influence Maximization Objective (Identifying High-Impact Accounts for Suppression)

This subsection establishes the optimization target. Given a budget $k$, the classical influence maximization (IM) problem seeks: Under IC and LT, $\sigma(\cdot)$ is monotone and submodular, which supports near-optimal approximation via greedy selection with provable guarantees [4]. In the suppression setting, the same optimization identifies accounts whose intervention (e.g., friction, monitoring, counter-messaging exposure, or de-amplification) can most strongly reduce future diffusion potential across the misinformation network, conceptually aligned with misinformation limitation via competing campaign/mitigation strategies.

### F. Seed-Selection

This subsection corresponds to "Algorithms" in the diagram and details the families evaluated.

*1) Centrality/structural heuristics (fast baselines):*

*a) Out-degree* prioritizes nodes with many outgoing neighbors (high immediate reach).

*b) Betweenness centrality* captures brokerage/bridge positions that can connect communities; the classic formulation is based on Freeman's betweenness centrality [19].

*c) PageRank assigns global* importance to nodes in directed graphs, initially introduced in the context of web link analysis [20].

*d)* k-core (core number) identifies nodes embedded in cohesive substructures; foundational discussion is linked to minimum-degree core decomposition [21].

*2) Greedy influence maximization with CELF* acceleration. Greedy selection iteratively adds the node with maximum marginal gain in Eq. (5):

$$\Delta = \sigma(S \cup \{v\}) - \sigma(S) \tag{5}$$

But repeated spread estimation is expensive. CELF reduces computation via lazy marginal-gain updates while preserving greedy behavior and is widely used as a strong quality baseline [4].

*3) Reverse Influence Sampling (RIS):* RIS approximates $\sigma(S)$ using random reverse-reachable (RR) sets; selecting nodes that cover the most RR sets provides a scalable approximation with theoretical grounding, making RIS-type methods suitable for large graphs and repeated experiments [22].

### G. Community-Aware Influence Maximization

After computing communities $\{C_i\}$, we impose a coverage constraint by distributing the seed budget across communities (e.g., proportional allocation) as mentioned below in Eq. (6) and Eq. (7):

$$k_i \approx k \cdot \frac{|C_i|}{|V|} \tag{6}$$

$$\sum_i k_i = k \tag{7}$$

And then select $k_i$ seeds within each community using the same ranking/optimization method (Degree/CELF/RIS/etc), producing [Eq. (8)]:

$$S = \bigcup_i S_i, |S_i| = k_i \tag{8}$$

The scientific rationale is that misinformation ecosystems are heterogeneous; community-aware seeding mitigates the risk of concentrating interventions in a single dense cluster and supports broader, cross-community mitigation strategies. The choice of Leiden strengthens this component because community integrity (connectivity/quality) influences downstream allocation and interpretation.

## IV. Results and Discussion

### A. Network Summary and Evaluation Setting

The OneCoin interaction network is modeled as a directed engagement graph built from retweet/reply relations. The details of the datasets are shown in Table I. Across all experiments, diffusion is evaluated under the two standard influence models—Independent Cascade (IC) and Linear Threshold (LT)—using the weighted-cascade normalization (equal influence share across a node's in-neighbors) to parameterize edge influence when ground-truth probabilities are unavailable. This modeling choice is consistent with the classical formulation of diffusion-based influence maximization [4].

TABLE I. Dataset Statistics

| | |
|---|---|
| Number of nodes | 5415 |
| Number of edges | 6117 |
| Average Degree | 2.2327 |
| Maximum Degree | 1582 |

### B. Influence Spread Under the LT Model

Fig. 3 compares the expected influence spread achieved by Degree, CELF, $k$-core, Betweenness, PageRank, and RIS under LT across budgets $k = \{5,10,15,20,25,30\}$. Three patterns are salient. First, CELF and Degree dominate across budgets, with the highest spreads observed at larger $k$; for example, at $k = 30$, CELF reaches ≈1620.9 activated users while Degree reaches ≈1619.5, indicating that, on this dataset, a simple out-degree heuristic is highly competitive once the seed budget is moderate to large. Second, RIS consistently provides a strong quality–efficiency trade-off, tracking the leading methods closely as $k$ increases (e.g., ≈1463.2 at $k$=30), which aligns with the design objective of RIS-style algorithms: scalable approximation to near-greedy performance via reverse-reachable sampling [22]. Third, the structural centrality baselines (Betweenness, PageRank, $k$-core) improve gradually with increasing $k$ but typically remain below CELF/Degree; this suggests that, in the OneCoin engagement network, large cascades are more strongly driven by outward amplification capacity (high out-degree and greedy marginal gain) than by purely bridge-based shortest-path control.

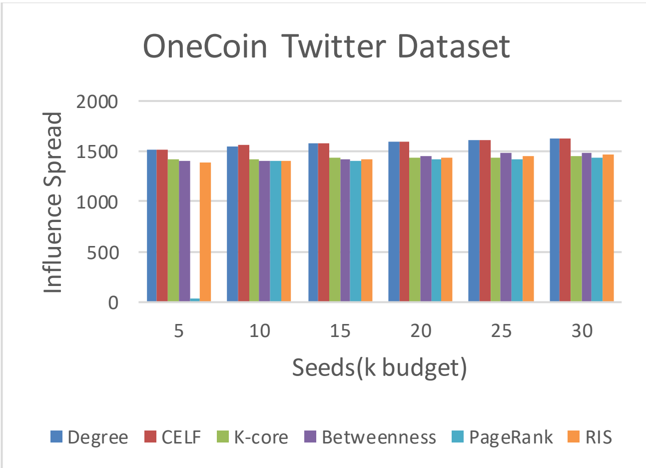


Fig. 3. Influence spread achieved by the evaluated seed-selection methods under the LT diffusion model on the OneCoin dataset.

### C. Influence Spread Under the IC Model

Under IC (see Fig. 4), the ranking trends are mainly consistent with LT. CELF achieves the most significant expected spread across budgets, reflecting the theoretical motivation for greedy selection under submodular influence functions [4]. At k = 30, CELF reaches ≈1597.9, followed by Degree at ≈1576.2. RIS remains competitive (≈1445.8 at k = 30 ), again reinforcing its role as a practical large-scale alternative to greedy [22]. An additional observation is the instability of PageRank at minimal budgets (e.g., at k = 5, PageRank produces very low spread). In directed engagement graphs, PageRank can rank nodes that are globally "prestigious" but not necessarily effective at outward activation (e.g., under the diffusion parameterization, limited outgoing influence pathways). As k increases, PageRank begins to include more structurally active accounts, and its spread becomes comparable to that of other centralities.

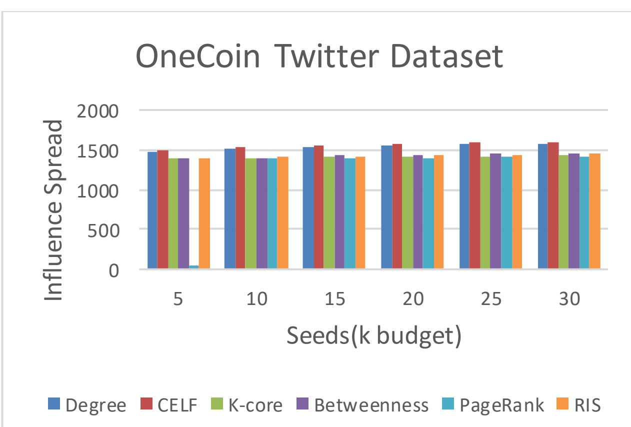


Fig. 4. Influence spread achieved by the evaluated seed-selection methods under the IC diffusion model on the OneCoin dataset.

### D. Runtime Comparison Under LT and IC

Fig. 5 (LT) and Fig. 6 (IC) report the total runtime per method (seed selection and spread evaluation). The computational cost is dominated by two factors: 1) the complexity of constructing the seed set and 2) the repeated spread estimation required by the evaluation protocol.

*1)* CELF is substantially more expensive than all alternatives because it relies on repeated marginal-gain evaluations (even with lazy updates) [4]. Concretely, the total runtime is ≈15.44 **s** under LT and ≈35.87 **s** under IC, both of which are far lower than those of the other methods.

*2)* Heuristic methods are operationally efficient**.** Degree and PageRank have very low seed-selection costs and total times, typically under a few seconds, in these runs.

*3)* Betweenness is moderately expensive relative to other heuristics due to shortest-path–based centrality computation, which explains its higher total time, even though it does not match the top-performing spreads.

Overall, the runtime plots support a practical conclusion: CELF maximizes spread but is least suitable for frequent recomputation. At the same time, RIS and Degree offer the best "deployable" trade-off between spread and runtime.

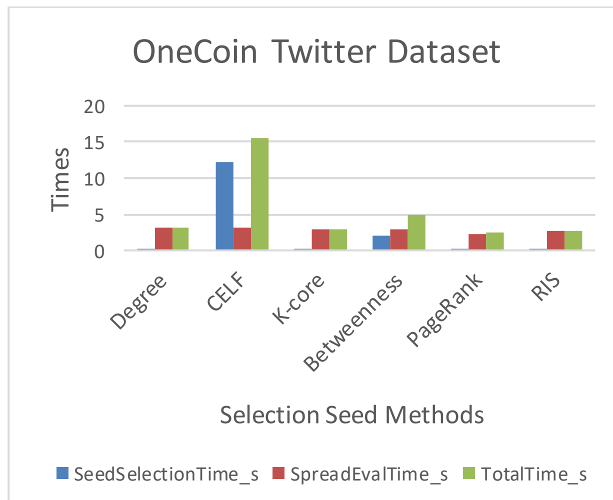


Fig. 5. Total computational time (in seconds) required by each seed-selection method under the LT diffusion model on the OneCoin dataset.

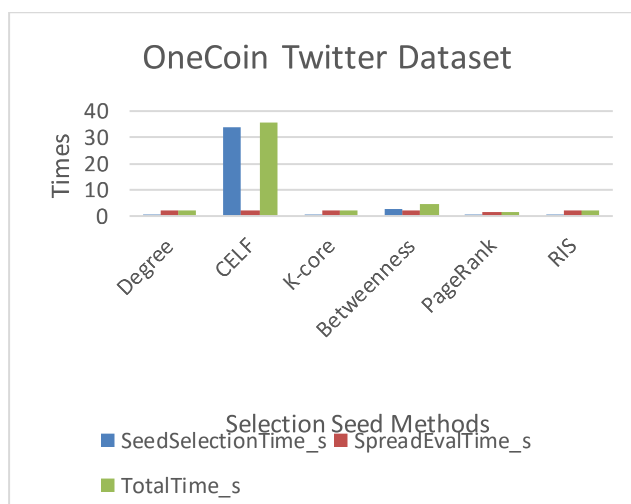


Fig. 6. Total computational time (in seconds) required by each seed-selection method under the IC diffusion model on the OneCoin dataset.

### E. Community-Aware Results with Leiden

Fig. 7 and Fig. 8 examine the effect of incorporating Leiden communities and enforcing community-aware seed allocation under LT. Leiden is chosen because it is designed to avoid poorly connected community partitions and to produce higher-quality communities efficiently, which is essential when community boundaries are used to constrain seeding decisions [8].

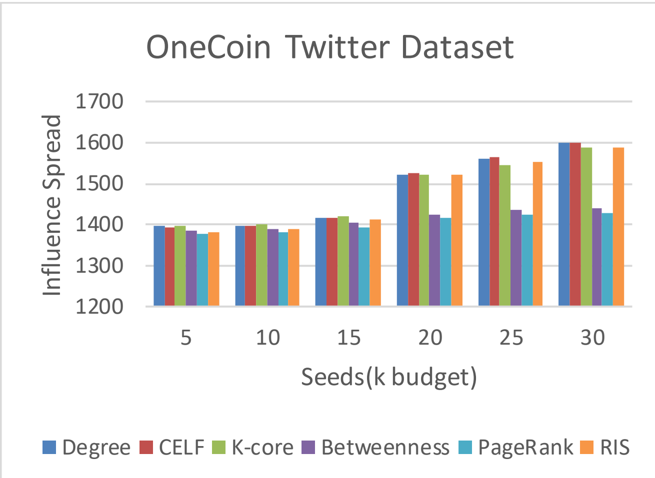


Fig. 7. Influence spread achieved by the evaluated seed-selection methods under the LT diffusion model on the OneCoin data communities using the Leiden algorithm.

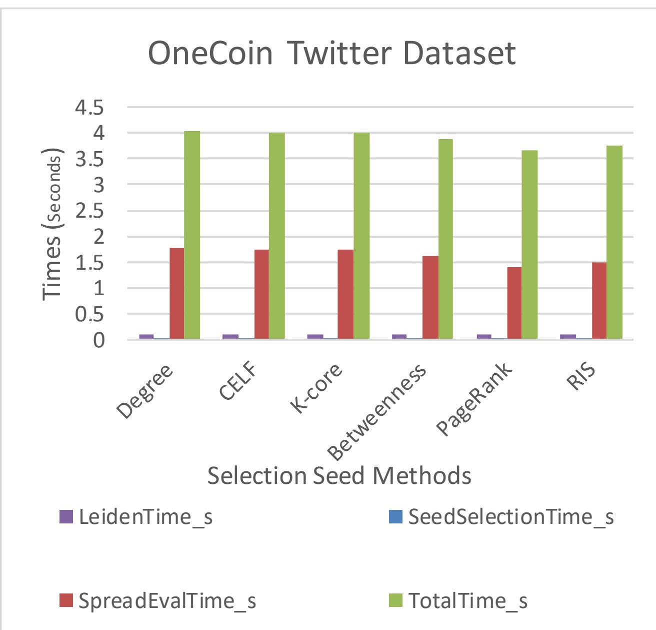


Fig. 8. Total computational time (in seconds) required by each seed-selection method under the LT diffusion model on the OneCoin data communities using the Leiden algorithm.

The results show a clear coverage–concentration trade-off:

- For Degree and CELF, community constraints typically reduce spread at small-to-mid budgets (e.g., around $k = 10-15$), because the unconstrained versions tend to concentrate seeds in the single densest, most cascade-prone region of the graph; community-aware allocation forces distributing seeds across multiple communities, which can reduce immediate cascade size. For methods that are sensitive to seed "over-concentration" effects, community awareness can be beneficial. Notably, $k$-core and RIS improve at larger budgets under Leiden constraints. For example, at $k = 30$, $k$-core increases from ≈1445.5 (global LT) to ≈1587.4 (Leiden-aware LT), and RIS increases from ≈1463.2 to ≈1589.1. This indicates that, once enough budget is available, spreading seeds across communities helps capture multiple diffusion pockets simultaneously, improving overall reach while also achieving better representativeness across the misinformation ecosystem.

- Community-aware allocation also stabilizes methods that can fail at minimal budgets (e.g., PageRank at $k = 5$), because the constraint prevents all seeds from being drawn from a narrow set of globally ranked nodes and instead enforces multi-community representation.

- Runtime under the Leiden-aware pipeline (see Fig. 8) remains within a few seconds for all methods in these runs, as the added community-detection step is computationally modest relative to greedy marginal-gain recomputation, consistent with Leiden's known efficiency properties [8].

### F. Discussion and Interpretation

Across both diffusion models, the results follow established influence-maximization theory: greedy selection is near-optimal under monotone submodular spread, which explains CELF's consistently highest spread [4], [6]. The runtime gap is also expected because CELF still relies on repeated marginal-gain estimation (even with lazy updates) [6]. The strong performance of the Degree heuristic (nearly matching CELF at larger k) indicates that the OneCoin engagement graph is sparse and hub-dominated; in such graphs, immediate outward reach can approximate marginal gains well. RIS provides a strong near-greedy alternative with substantially lower runtime, consistent with reverse-reachable (RR) set sampling methods designed for near-linear-time approximation [7], [22].

Community-aware seeding changes the intervention footprint. By allocating seeds proportionally across Leiden communities, the approach discourages concentrating interventions in a single dense cluster and instead targets multiple sub-audiences. This explains the observed coverage–concentration trade-off: at small-to-mid budgets, coverage constraints can reduce the single largest cascade, while at larger budgets they can improve overall reach by activating multiple diffusion pockets in parallel, consistent with community-based IM objectives [16]–[18]. Leiden is a suitable choice because it refines partitions to ensure well-connected communities, improving the reliability of downstream community-constrained seeding [8].

### G. Limitations and Threats to Validity

Our evaluation has limitations in modeling and data. IC and LT are simplified abstractions that omit content semantics, heterogeneous susceptibility, recommender effects, and time-varying exposure. Weighted-cascade normalization is a standard practical parameterization when true edge probabilities are unknown [4], but it may not match platform-specific amplification dynamics and can affect which nodes appear influential. The engagement network is built from observable retweets/replies and is sensitive to collection choices (keywords, time window, API limits), missing/deleted/private content, and automated accounts, which can introduce sampling bias and missing edges. These factors can affect both community detection and seed ranking. Results from OneCoin may also not generalize to other scam domains. Future work should incorporate temporal diffusion (time-sliced graphs), explicit competitive/containment objectives [5], robustness checks over alternative probability parameterizations, and hybrid models that combine network structure with content/account-level signals (e.g., credibility and bot likelihood).

## V. Conclusion

This study presented a diffusion-based, community-aware influence-maximization workflow to support the suppression of scam-related misinformation in the OneCoin Twitter engagement ecosystem. Using a directed interaction graph of 5,415 accounts and 6,117 engagement edges, we evaluated six seed-selection strategies under IC and LT diffusion, using the weighted-cascade parameterization [4]. CELF achieved the largest expected spread (e.g., ≈1620.9 under LT and ≈1597.9 under IC at k=30), but required the most computation (≈15.44 s under LT; ≈35.87 s under IC), reflecting the cost of repeated marginal-gain evaluation despite lazy acceleration [6]. RIS delivered a strong near-greedy alternative with substantially lower runtime, consistent with RR-set sampling guarantees and prior evidence of practical efficiency [7], [22].

Integrating Leiden community detection and proportional seed allocation shifted seed composition toward broader sub-community coverage, which is operationally meaningful for scam ecosystems sustained by multiple semi-separated recruitment clusters [8]. At larger budgets, community-aware allocation can also improve diffusion reach for methods that otherwise over-concentrate seeds (e.g., k-core and RIS under LT). Overall, the results provide actionable guidance for mitigation planning: greedy-quality methods maximize reach when offline computation is acceptable, while RIS and simple degree-based heuristics offer deployable trade-offs between spread and runtime for scalable, repeatedly updated intervention workflows.

Future research should move beyond static-graph IC/LT simulations by incorporating temporal dynamics, competitive influence-blocking objectives, and sensitivity analyses over propagation parameters, while enriching network signals with content and account-level features to better represent real platform diffusion and enforcement constraints.

## References

[1] 2019. [Online]. Available: https://www.justice.gov/usao-sdny/pr/manhattan-us-attorney-announces-charges-against-leaders-onecoin-multibillion-dollar. Accessed: 26-Dec-2025. U.S. Department of U.S. Attorney's Office SDNY "Manhattan U.S. Attorney announces charges against leaders of OneCoin multibillion dollar fraud scheme Justice.

[2] 2023. [Online]. Available: https://www.justice.gov/usao-sdny/pr/co-founder-multibillion-dollar-cryptocurrency-scheme-onecoin-sentenced-20-years-prison. Accessed: 26-Dec-2025. U.S. Department of U.S. Attorney's Office SDNY "Co-founder of multibillion-dollar cryptocurrency scheme OneCoin sentenced to 20 years in prison Justice.

[3] 2022. [Online]. Available: https://www.fbi.gov/news/stories/ruja-ignatova-added-to-fbis-ten-most-wanted-fugitives-list. Accessed: 26-Dec-2025 Federal Bureau of "Ruja Ignatova added to FBI's Ten Most Wanted Fugitives list Investigation.

[4] 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, 2003, pp. 137–146.

[5] C. Budak, D. Agrawal, and A. El Abbadi, "Limiting the spread of misinformation in social networks," in Proceedings of the 20th international conference on World wide web, 2011, pp. 665–674.

[6] J. Leskovec, A. Krause, C. Guestrin, C. Faloutsos, J. VanBriesen, and N. Glance, "Cost-effective outbreak detection in networks," in Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining, 2007, pp. 420–429.

[7] C. Borgs, M. Brautbar, J. Chayes, and B. Lucier, "Maximizing social influence in nearly optimal time," in Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms, 2014, pp. 946–957.

[8] V. A. Traag, L. Waltman, and N. J. Van Eck, "From Louvain to Leiden: guaranteeing well-connected communities," Scientific reports, vol. 9, no. 1, pp. 1–12, 2019.

[9] Y. Tang, X. Xiao, and Y. Shi, "Influence Maximization: Near-Optimal Time Complexity Meets Practical Efficiency," in Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data (SIGMOD), 2014.

[10] M. Jaouadi and others, "A Survey on Influence Maximization Models," Expert Systems with Applications, 2024.

[11] Y. Ye and others, "Influence Maximization in Social Networks," Software Impacts, 2022.

[12] S. Solanki and others, "A Survey on Information Diffusion and Competitive Influence Maximization," Social Network Analysis and Mining, 2025.

[13] H. Zhang, H. Zhang, X. Li, and M. T. Thai, "Limiting the Spread of Misinformation While Effectively Raising Awareness in Social Networks," in Computational Social Networks (CSoNet 2015), 2015.

[14] Q. Luo, D. Yu, D. Wang, Y. Zhang, Y. Zheng, and Z. Cai, "Influence Blocking Maximization Under Refutation," Social Network Analysis and Mining, vol. 13, p. 143, 2023.

[15] J. Xie, F. Zhang, K. Wang, X. Lin, and W. Zhang, "Minimizing the Influence of Misinformation via Vertex Blocking," in Proceedings of the IEEE International Conference on Data Engineering (ICDE), 2023.

[16] A. Bozorgi, H. Haghighi, M. S. Zahedi, and M. Rezvani, "INCIM: A Community-Based Algorithm for Influence Maximization Problem Under the Linear Threshold Model," Information Processing & Management, vol. 52, no. 6, pp. 1188–1199, 2016.

[17] S. Banerjee, M. Jenamani, and D. K. Pratihar, "ComBIM: A Community-Based Solution Approach for the Budgeted Influence Maximization Problem," Expert Systems with Applications, vol. 125, pp. 1–13, 2019.

[18] X. Qin, C. Zhong, and H. Lin, "Community-Based Influence Maximization Using Network Embedding in Dynamic Heterogeneous Social Networks," ACM Transactions on Intelligent Systems and Technology, 2023.

[19] L. C. Freeman, "A set of measures of centrality based on betweenness," Sociometry, pp. 35–41, 1977.

[20] S. Brin and L. Page, "The anatomy of a large-scale hypertextual web search engine," Computer networks and ISDN systems, vol. 30, no. 1–7, pp. 107–117, 1998.

[21] S. B. Seidman, "Network structure and minimum degree," Social networks, vol. 5, no. 3, pp. 269–287, 1983.

[22] C. Chekuri, Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms. SIAM, 2014.