Misinformation Mitigation under Differential Propagation Rates and Temporal Penalties

Michael Simpson
University of British Columbia
Vancouver, Canada
mesimp@cs.ubc.ca

Farnoosh Hashemi
University of British Columbia
Vancouver, Canada
farsh@cs.ubc.ca

Laks V.S. Lakshmanan
University of British Columbia
Vancouver, Canada
laks@cs.ubc.ca

ABSTRACT

We propose an information propagation model that captures important temporal aspects that have been well observed in the dynamics of fake news diffusion, in contrast with the diffusion of truth. The model accounts for differential propagation rates of truth and misinformation and for user reaction times. We study a time-sensitive variant of the *misinformation mitigation* problem, where k seeds are to be selected to activate a truth campaign so as to minimize the number of users that adopt misinformation propagating through a social network. We show that the resulting objective is non-submodular and employ a sandwiching technique by defining submodular upper and lower bounding functions, providing data-dependent guarantees. In order to enable the use of a reverse sampling framework, we introduce a weighted version of reverse reachability sets that captures the associated differential propagation rates and establish a key equivalence between weighted set coverage probabilities and mitigation with respect to the sandwiching functions. Further, we propose an offline reverse sampling framework that provides $(1 - 1/e - \epsilon)$-approximate solutions to our bounding functions and introduce an importance sampling technique to reduce the sample complexity of our solution. Finally, we show how our framework can provide an anytime solution to the problem. Experiments over five datasets show that our approach outperforms previous approaches and is robust to uncertainty in the model parameters.

PVLDB Reference Format:
Michael Simpson, Farnoosh Hashemi, and Laks V.S. Lakshmanan. Misinformation Mitigation under Differential Propagation Rates and Temporal Penalties. PVLDB, 15(10): 2216 - 2229, 2022.
doi:10.14778/3547305.3547324

PVLDB Artifact Availability:
The source code, data, and/or other artifacts have been made available at https://github.com/stamps/temporal_misinformation_mitigation.

1 INTRODUCTION

Social networks have rapidly transformed into a prominent hub for political campaigns, viral marketing, and the dissemination of news and health information. As an unfortunate side effect, there has been an increase in the number of "bad actors", such as spammers, hackers, and bots, exploiting these platforms to spread fake news and misinformation. A fundamental question is: *How can one limit the spread of misinformation in social networks?* Once misinformation is detected, one feasible approach is to introduce a *truth campaign* with a goal of reaching users before, or at least not much later than, they are reached by the misinformation. The *misinformation mitigation* (MM) problem [6, 14] aims to select effective seed nodes for the truth campaign such that the spread of misinformation can be limited as much as possible.

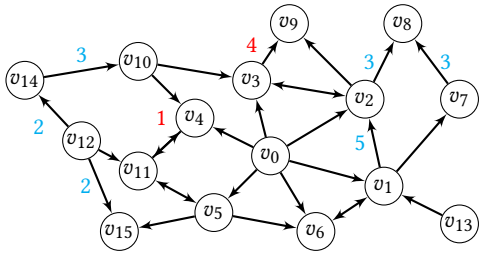

Figure 1: Sample instance: edge labels = meeting delays (ML); MLs not shown are 1; numbers besides nodes are reaction times = activation window (AW) lengths; AW lengths not shown are 0.

Notably, existing propagation models considered in the MM literature [6, 11, 31, 33, 34, 39, 40, 44–47] fail to incorporate critical temporal aspects that have been well observed in the dynamics of fake news diffusion [13, 28, 49]. We argue that it is important to distinguish between the relative propagation dynamics of fake news and truth as it has been observed that fake news often "spreads like wildfire" online [49] while the adoption of truth occurs much slower. For instance, Fig. 1 depicts a sample instance of misinformation/truth propagation, where all the edge propagation probabilities are 1 and the edge labels indicate the diffusion delay for truth. E.g., if node v_2 adopts misinformation at time t, it will propagate to v_8 at time $t + 1$. By contrast, if v_2 adopted truth at t, it will propagate to v_8 at $t + 3$. The differential propagation rates of truth and fake news may have considerable consequences for the selection of effective seed nodes. Secondly, important dynamics are at play in user decision making. E.g., a recent study [13] observed that approximately 59% of users forego reading articles linked in social media posts before acting on them, while the time spent reading linked articles [28] ranges from under a minute for short articles to several minutes for longer articles. Thus, the "reaction times" of users can vary considerably. To illustrate, in Fig. 1, all users save v_3, v_4 may react instantly (no reading), while v_3 and v_4 may read linked articles before reacting, with v_4 reacting quicker (1 time unit) than v_3 (4 time units). Finally, mitigation strategies for combating misinformation are only effective when the *truth* arrives at a user

Proceedings of the VLDB Endowment, Vol. 15, No. 10 ISSN 2150-8097.
doi:10.14778/3547305.3547324

either (a) before the arrival of the fake news or (b) shortly after the user becoming aware of the fake news. In particular, truth arriving *later* than misinformation may still have an effect, up to a reasonable delay. E.g., suppose v_0 is a seed for misinformation while v_{12} is a seed for truth, with both adopting the information at time 0. Then v_3, v_4, v_5 will become aware of the misinformation at time 1, while the truth reaches v_3, v_4, v_5 at times 6, 2, 2 respectively. Meanwhile, v_3, v_4, v_5 will react at times $1 + 4 = 5$, $1 + 1 = 2$, and $1 + 0 = 1$ respectively. Thus, truth reaches v_4 in time for its reaction while it arrives at v_3 and v_5 too late. Further, as illustrated above, as the delay between the arrival of misinformation and truth increases, the *effectiveness* of the mitigating campaign drops off significantly [10, 32]. *Models used in prior misinformation mitigation studies fail to account for these phenomena, supported by real-world observations.*

Novelty. In this paper, we study an interesting and realistic variant of the classic MM problem incorporating the twin time-critical aspects of misinformation propagation that have been observed and validated in earlier studies [13, 28, 49]: differential propagation rates and user reaction times. Our approach uses a node-level dynamic penalty function based on the delay between the arrival of the competing campaigns. We propose a new propagation model, the *Temporal Competitive Independent Cascade* (TCIC) model which, unlike existing propagation models, accounts for differential propagation rates and user reaction times, by employing two critical components which work jointly to properly model the dynamics of diffusion: edge-level campaign-specific *time-delayed* propagation and node-level *activation windows* for making adoption decisions.

We then define a novel optimization problem for misinformation mitigation under the TCIC model. Unlike prior MM models [6, 36, 39, 40, 44, 46, 47], which are based on competitive IC [7], whose objectives satisfy submodularity when the campaigns share propagation probabilities, we show that submodularity does not hold for our objective in the TCIC model even when the campaigns share probabilities. We also show that the recent guarantees shown for Greedy [4] based on curvature and submodularity ratio when applied to non-submodular objectives, lead to degenerate results for our objective function. To overcome this challenge, we employ the *Sandwich Approximation (SA)* [27] and develop non-trivial upper and lower bounding submodular functions to produce solutions with a data-dependent approximation guarantee. Further, existing state-of-the-art solutions for *influence maximization* (under a single campaign), such as *IMM* [42], *SSA* [15, 30], and *OPIM* [41], are based on *reverse sampling (RS)*. However, adapting the *RS* machinery (based on *reverse reachability (RR)* sets) to our propagation model comes with the challenge that the set of nodes reached by the misinformation is an *unknown* meaning our algorithm cannot inherit the sample complexity lower bound from prior work and thus requires a novel derivation, which we provide. We develop a *RS* framework by building on the state-of-the-art *OPIM* algorithm. Due to the complex interactions that occur during the propagation of the fake and mitigating campaigns as well as the new temporal model components, the construction of the analog to RR sets under our new model requires great care. Further, adapting RR sets to our setting requires pushing the idea of tie-breaking between campaigns into the notion of RR sets. To the best of our knowledge, we are the *first to incorporate a proportional tie-breaking rule* into the reverse sampling framework, which introduces additional challenges in constructing RR sets. Second, in order to further reduce the number of samples required by the *RS* framework, we define an unbiased estimator for our upper and lower bound objectives based on *importance sampling* leading to reduced variance and tighter concentration bounds. We tackle the challenge of the sample complexity depending on the expected influence of the misinformation, an *unknown*, by developing a novel normalization term.

Table 1: Frequently used notation.

Notation	Description
G, m, n	Social network graph with n nodes and m edges
F, M, S_F, S_M	The misinformation and mitigation campaigns and their seed sets, respectively
$m_F(u,v), m_M(u,v)$	Meeting event probabilities
X, R_F^X	Possible world of the TCIC model and set of nodes reachable from F in X
h_e^F, h_e^M	Sampled meeting event edge lengths along edge e in X
τ_v, π_v	Sampled AW length and in-neighbour permutation in X
t_v^F, t_v^M	First step v meets with F, M
$\rho_X(v,S)$	Reward achieved in X at v by S
$\mu_X(S_M), \mu(S_M)$	Total reward achieved by S i nX and in expectation, respectively
$\overline{\mu}, \underline{\mu}$	Submodular upper/lower bounding functions of μ
INF_F, INF_1	Expected influence of F and largest expected influence of any size-1 node set
Γ	Misinformation sampling error
EPT	Expected complexity of generating an RDR set

Our main contributions are: (1) we introduce a novel propagation model capturing important temporal aspects pertaining to the diffusion of and reaction to truth and misinformation and define a novel MM problem formulation with delay-dependent reward (§ 2). (2) We develop non-trivial upper & lower bounding submodular functions for our non-submodular objective to use in a sandwiching technique (§ 3). (3) We introduce an importance sampling technique to reduce the sample complexity of our solution (§ 4). (4) We develop a reverse sampling framework that provides a $(1-1/e-\epsilon)$-approximate solution to the upper & lower bounding objectives yielding an instance-dependent approximation to the MM objective (§ 5). (5) We show that our algorithm can provide an anytime solution to the MM problem (§ 5). (6) We present a thorough experimental validation (§ 6). For lack of space, some proofs are sketched. Complete details can be found in [37].

2 PRELIMINARIES

We formalize the notions of differential propagation rates and activation windows and then present our new propagation model referred to as TCIC (for Temporal Competitive Independent Cascade). Let $G = (V, E)$ be a social network with sets of nodes V and directed edges E, where $|V| = n$, $|E| = m$. Let F (for "*F*ake") and M (for "*M*itigating") denote two influence campaigns with seed sets S_F and S_M, respectively.[1] The seeds S_F (S_M) are active in campaign

[1] We assume $S_F \cap S_M = \emptyset$, w.l.o.g.

F (resp. M) at time $t = 0$. We assume each edge $e = (u, v)$ in E is associated with a propagation probability $p(u, v) \in (0, 1]$.

Meeting Events. We associate with each edge $e = (u, v)$ *meeting probabilities* $m_F(u, v), m_M(u, v) \in (0, 1]$. At *any* step, an active node u in campaign F (or M) meets any of its currently inactive neighbours v independently with probability $m_F(u, v)$ (or $m_M(u, v)$). A node activated at time T attempts to meet with its inactive out-neighbours at *every* step $t \geq T+1$ until there is a successful meeting. Then, if a meeting event occurs between u and v in F (or M) for the *first* time at step t, then u is given a single chance to activate v in F (or M) with independent probability $p(u, v)$. If the activation attempt is successful, v becomes active in F (or M) at time t and enters an activation window, described below. Once activated, subsequent meeting attempts from in-neighbours are ignored.

With the above formulation, we can capture the observation that *truth travels more slowly than fake news* by setting $m_F(u, v) \geq m_M(u, v), \forall$ edge $(u, v) \in E$. As a special case, we can set $m_F(u, v) = 1$, i.e., only apply the meeting events to nodes active in the mitigating campaign M. This results in F traversing every edge in a single hop, while M may be delayed by several meeting attempts. For convenience, under this special case, for an edge $e = (u, v)$ we will write $m_M(u, v)$ as $m(u, v)$ or $m(e)$. We shall henceforth assume this special case for ease of exposition, although our theory and techniques apply to the general case.

Activation Windows. Motivated by concepts well established in Sociology and Marketing [1, 2, 13, 18, 27, 28], we distinguish between *awareness* and *adoption* in our TCIC model by defining an *activation window (AW)* for each node. The AW augments the propagation behaviour defined by meeting events such that successful activation attempts now result in awareness, and not adoption. Specifically, we say node u becomes *aware* in the *first* step that an in-neighbour of u active in either campaign M or F succeeds in meeting with u. However, before deciding to commit to activation in either campaign, u enters its *activation window*: a period of time during which it may receive conflicting and/or reinforcing information, which it factors in making an adoption decision. The dynamics of the AW are governed by a node-level parameter $\gamma(u)$ that determines the length of the AW for u. The *closing* function $\gamma(\cdot) \rightarrow \mathbb{Z}^+$ allows the model to capture varying window sizes, typically related to the time spent reading articles linked in social media posts. Example choices for the closing function include: (i) a constant function where all nodes wait for some fixed number of steps before making a decision, (ii) a uniform function where window lengths are chosen uniformly at random between 0 and some closing time τ, or (iii) an attenuating function (e.g. exponential decay with an appropriate mapping to $\mathbb{Z}^+$) where some users may wait substantially longer (to gather additional information) before making an adoption decision. Furthermore, $\gamma(\cdot)$ can be made *node-specific* to capture the individual behaviour of users.

Tie-breaking Policy. Observe that (active users from) both campaigns may meet with a node u within the AW and at the end of the window, u must decide to adopt F or M. In such a scenario, both campaigns are attempting to activate u and so we require a tie-breaking policy. We employ a *weighted random choice* policy based on in-neighbour activation counts for each campaign over the duration of the AW, described as follows: when the AW closes, the probability that node u activates in campaign F is $|N_F^-(u)|/|N^-(u)|$ where $N_F^-(u)$ (resp. $N^-(u)$) is the set of in-neighbours of u that met with u from campaign F (resp. either campaign). Similarly, the probability that node u activates in campaign M is $|N_M^-(u)|/|N^-(u)|$.

TCIC Model. Our new propagation model is defined by incorporating the edge-level meeting events and node-level activation windows into a standard competitive independent cascade model [7]. The propagation process terminates when all active nodes in both campaigns have met with all their out-neighbours and no new nodes can be activated in either campaign. The model parameters allow for a great deal of expressiveness to accurately describe many observed real-world adoption behaviours.

Problem Statement. We argue that in gauging the extent of misinformation mitigation, merely counting the number of users prevented from adopting the misinformation after the propagation terminates is too restrictive. The reason is that the quality of the mitigation depends on the delay, if any, in the arrival of the truth. To help capture this, assume that there is a reward function $\rho(\cdot) : V \times 2^V \rightarrow R^{>=0}$, which given a node v and a seed set S_M, returns a real number indicating the effectiveness of the mitigation at v after the propagation terminates. We will provide more details on $\rho(\cdot)$ in the next section. Given such a reward function, we define the *expected mitigation* by $\mu(S_M) = \mathbb{E}[\sum_v \rho(v, S_M)]$ where the expectation is taken over the randomness in the propagation process. We are now ready to formally state the problem we study in this work.

Problem 1. *Given a misinformation seed set S_F, the* misinformation mitigation (MM) problem *under the TCIC model is to find a seed set S_M with at most k nodes that maximizes the expected mitigation. Formally, find a seed set S_M satisfying* $\text{argmax}_{S_M \subset V \setminus S_F, |S_M| \leq k} \mu(S_M)$.

Possible Worlds. We can view the stochastic propagation process under TCIC using an equivalent "possible worlds" interpretation. Suppose that before the propagation process starts, a set of outcomes for all meeting event attempts (i.e., the number of failed meeting attempts before two neighbours successfully meet), activation windows parameters, and edge liveness are pre-determined. Specifically, for each edge $e = (u, v) \in E$, we declare the edge "live" with probability $p(u, v)$, or "blocked" otherwise. Further, for each edge, we sample from a geometric distribution parameterized by success probability $m_F(e)$ ($m_M(e)$) the (random) number of meeting event attempts, denoted h_e^F (h_e^M), required by campaign F (or M) along e. Next, a set of outcomes for all activation windows are pre-determined. Specifically, for each node $v \in V$ we sample a window length τ_v from the closing function $\gamma(v)$ and a random permutation π_v of the active in-neighbours of v (i.e., those active in-neighbours connected to v by live edges) for the purpose of tie-breaking. All random events, including coin flips are independent. Thus, a certain set of outcomes of all coin flips and sampled parameter values corresponds to one possible world under the TCIC model. A possible world, denoted X, is a deterministic graph obtained by conditioning on a particular set of outcomes. We denote by R_F^X the set of nodes reachable by F in possible world X in the absence of M.

Next, we define the notion of distance in X for each campaign. Consider a live edge $e = (u, v)$ in X. Traditionally, without meeting events, v is reachable from u in a single hop. Now with pre-determined meeting event attempts, v is reachable from u in h_e^F (or

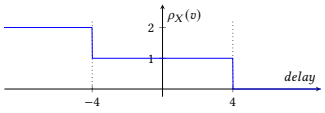

Figure 2: $\rho_X(v_3)$ **with** $\tau_{v_3} = 4$.

h_e^M) hops. Then, the *delayed-distance* of a path P from u to v is the total number of hops along the edges of P plus the sum of activation window lengths of nodes in $P \setminus \{u, v\}$. Finally, for campaign $A \in \{F, M\}$, the delayed-distance $dd_X(S_A, v)$ from a seed set S_A to v in X is the delayed-distance of the live-path P from $u \in S_A$ to v that minimizes $\sum_{e \in P} h_e^A + \sum_{x \in P \setminus \{u,v\}} \tau_x$.

Example 1. *We illustrate the propagation process on the network in Fig. 1 under the TCIC model. For simplicity, we assume each edge has propagation probability $p(e) = 1$. The sampled meeting delays h_e^M and AW lengths τ_v are indicated in cyan and red labels respectively. Let $\pi_{v_4} = (v_{11}, v_{10}, v_F)$ and $\pi_{v_{15}} = (v_5, v_{12})$ be the sampled permutations. Consider seeds v_0 for campaign F and v_{12} for campaign M and consider the resulting cascade through the network. At time 1, v_1, v_2, v_5, v_6 adopt F and v_{11} adopts M. Meanwhile, v_3, v_4 open their AW's. At time 2, both F and M reach v_{15}. Additionally, M reaches v_4 as its AW closes. Tie-breaks result in v_4 and v_{15} adopting M and F respectively. Meanwhile v_7, v_8, v_9 adopt F and v_{14} adopts M. Finally, at time 5 v_{10} adopts M and v_3 adopts F.*

Using the above sampled parameters, with v_0 as the F seed, let us compare M seeds v_{12} and v_{14}. It can be verified that v_{12} causes the nodes $v_{12}, v_{11}, v_4, v_{14}, v_{10}$ to either adopt or be informed of M. Of these only v_{11} and v_3 would have adopted F if there was no M campaign, so intuitively v_{12} "saves" 2 nodes. By contrast, v_{14} causes v_{14}, v_{10}, v_3 to adopt or be informed of M, of which only v_3 would have adopted F if there was no M campaign, so v_{14} "saves" v_3. □

2.1 Mitigation Reward

Motivated by the time-critical nature of the MM problem, we introduce a novel reward function that intuitively captures the penalty paid when mitigation arrives too late after the misinformation. In particular, the reward function is designed such that in case adoption of the truth cannot be secured, awareness is encouraged.

Delay-specific Reward Function. The reward function $\rho_X(\cdot)$ is defined w.r.t. activating a node in campaign M relative to the behaviour of campaign F. First, for nodes $v \notin R_F^X$ that would *not* have been activated in F in the absence of M we define $\rho_X(v, S_M) = 0$. Next, consider a node v that would have been activated in F in the absence of M, i.e. there exists a path from S_F to v in X. Let t_v^F be the first step in which v meets with a node in campaign F and t_v^M be the first step in which v meets with a node in campaign M. We use the convention that t_v^F (or t_v^M) is infinite if no meeting with a node from F (or M) occurs. We define the reward $\rho_X(v, S_M)$ for node v as a function of the amount of time that has passed between the mitigation and the fake news arriving at v. In particular, we consider the step function given by Eq. (1). There are three cases for the amount of reward achieved: (i) reward 2 if the truth arrives at v sufficiently early such that the misinformation arrives after the AW closes or if the presence of M stops F from ever reaching v (in the case of $t_v^F = \infty$), (ii) reward 1 if both the truth and misinformation arrive at v within the AW and (iii) no reward if the truth arrives after the AW closes or not at all (i.e., $t_v^M = \infty$). The reward function is illustrated in Fig. 2 for v_3 from Fig. 1 where $\tau_{v_3} = 4$.

$$\rho_X(v, S_M) = \begin{cases} 2 & \text{if } t_v^M < t_v^F - \tau_v \\ 1 & \text{if } |t_v^M - t_v^F| \leq \tau_v \\ 0 & \text{if } t_v^M > t_v^F + \tau_v \end{cases} \quad (1)$$

When the context is clear, we write $\rho_X(v, S_M)$ as $\rho(v)$. We refer to the reward achieved by set S_M, given S_F, after the propagation terminates as the *mitigation* and denote it by $\mu_X(S_M) = \sum_v \rho_X(v)$. Further, we denote the *expected mitigation* by $\mu(S_M) = \mathbb{E}[\mu_X(S_M)]$.[2]

Design Decisions. The advantage of our new reward function, and the purpose of the middle condition of Equation 1, is to promote campaign M reaching nodes that are reached by campaign F, despite not being able to *guarantee* adoption of M, owing to tie-breaking. Thus, even if a node does not end up activating in M, the user will be exposed to the true information, which we argue is a natural goal to strive for. Clearly, in the event that adoption of M cannot be guaranteed, promoting its awareness is preferable to inaction. If the truth reaches a user sufficiently early (compared to the misinformation) then their adoption decision will be uncontested by the misinformation. Thus, the subsequent propagation of truth by the user leads to the desired outcome. Hence, this favourable scenario contributes maximal reward towards the objective. Meanwhile, due to the tie-breaking policy, reaching users within their AW always provides an opportunity for the adoption of truth and its subsequent propagation. However, when both the misinformation and truth arrive within the AW, the adoption of truth is no longer uncontested and, as such, contributes less reward. Finally, when mitigation arrives too late after the misinformation, we penalize the mitigation campaign to capture its reduced effectiveness. Further, recall that the reward function only attributes non-zero values when campaign F reaches node v. Since we do not credit any reward to those nodes reached by the mitigation but not the misinformation, this encourages a solution to "focus" on mitigating the spread of misinformation, and discourages solutions that blindly maximize the spread of the truth. We note that in place of a step function, we could use an arbitrary non-increasing reward function. We revisit this point in § 3.

Problem Properties. The original MM problem under the CIC model [6] is **NP**-hard. It is a special case of MM under the TCIC model, with all $m(u, v) = 1$, a length zero activation window, and dominant tie-breaking.

Proposition 1. *The misinformation mitigation problem is **NP**-hard under the TCIC model.*

An important observation is that, while the expected mitigation $\mu(\cdot)$ is monotone under the TCIC model with reward function $\rho(\cdot)$, it is not submodular.

Theorem 1. *Given a seed set S_F, the mitigation function $\mu(\cdot)$ is not submodular in general under the TCIC model.*

Bian et al. [4] recently provided approximation guarantees for the standard greedy algorithm on non-submodular objectives based on the submodularity ratio and curvature of the objective. Unfortunately, the submodularity ratio of the MM objective can be as

[2] Mitigation depends on S_F, but we omit S_F as an argument of $\mu(.)$ since S_F is a fixed input to the problem.

small as 0 and thus their result does not provide any non-trivial guarantees. To overcome the challenges of a non-submodular objective function, we leverage the *Sandwich Approximation* of [27] by developing appropriate upper and lower bounding functions for our mitigation objective, yielding a solution with non-trivial data-dependent approximation guarantees.

3 SANDWICHING THE OBJECTIVE

The *Sandwich Approximation (SA)* of [27] leverages upper ($\overline{\mu}$) and lower ($\underline{\mu}$) bounds on a non-submodular objective function to provide data-dependent approximation guarantees. Specifically, given upper & lower bounding functions that are submodular, we can obtain solutions S^U and S^L resp. with approximation guarantees.

Lower Bound. We make the following observation to motivate our choice of lower bound objective $\underline{\mu}(\cdot)$. The supermodular behaviour of $\mu(\cdot)$ arises from the combined effort of seed nodes in S_M which individually would not activate a target node v. That is, together they are able to block the paths from S_F to v such that the mitigating campaign's disadvantage, in the form of meeting events, is overcome. E.g., in Fig. 1, suppose $S_F = \{v_1\}$. Then neither one of v_2, v_7 by itself can activate v_8 in M. However, $S_M = \{v_2, v_7\}$ can. Thus, to eliminate the possibility of such *coordination*, we define a lower bound function $\underline{\mu}_X$ that only measures the maximum mitigation achieved by any node in S_M when it acts as a singleton seed set. Formally, we define $\underline{\mu}_X(S_M) = \sum_v \max_{u \in S_M} \rho_X(v, \{u\})$ and $\underline{\mu}(S_M) = \mathbb{E}[\underline{\mu}_X(S_M)]$. Note, $\underline{\mu}_X$ clearly lower bounds μ_X since $\max_{u \in S_M} \rho_X(v, \{u\}) \leq \rho_X(v, S_M)$.

Lemma 1. *$\underline{\mu}(S_M)$ is submodular.*

Upper Bound. An obvious candidate for $\overline{\mu}(\cdot)$ is to forego the meeting events associated with campaign M and enforce an M-dominant tie-breaking rule. The resulting model reduces to the CIC model under which previous results ensure that the resulting objective is submodular. The existence of meeting events only acts to "hinder" the mitigation and thus without them the mitigation would reach every node sooner, thus increasing reward. Further, M-dominant tie-breaking rule ensures all propagation paths shared by the two campaigns are won by the mitigation. However, we develop a tighter upper bounding function so as to improve the data-dependent approximation guarantees.

Consider a possible world X and call all edges touched by the propagation of campaign F in X, *critical* edges $E_C \subseteq E$. Next, construct a modified possible world X' by removing all meeting events for campaign M on critical edges. That is, $h_e^M = 1$ for all critical edges $e \in E_C$. Next, define an *overlap* indicator variable $\mathbb{I}_v^{OL}$ where $\mathbb{I}_v^{OL} = 0$ iff the collection of paths from S_F to v is edge-disjoint from the collection of paths from S_M to v in X'. Intuitively, if there is no overlap then there is no opportunity for several seeds from S_M to "collude" together to achieve a larger reward than if they were to act alone. Finally, we define a modified reward function $\rho'_X(\cdot)$ that upper bounds $\rho_X(\cdot)$ by lifting the reward to its maximum value for the case that both campaigns reach some node u within its activation window. Denote $t_v^{A,X}$ as the first step in which v meets with a node in campaign A in possible world X, where $A \in \{F, M\}$. We similarly define $\rho'_X(v, S_M) = 0$ when $v \notin R_F^X$, and otherwise as

$$\rho'_X(v, S_M) = \begin{cases} 2 & \text{if } t_v^{M,X} \leq t_v^{F,X} + \tau_v \\ 0 & \text{if } t_v^{M,X} > t_v^{F,X} + \tau_v. \end{cases} \tag{2}$$

Then, we define

$$\overline{\mu}_X(S_M) = \sum_v \max_{u \in S_M} \begin{cases} \rho_{X'}(v, \{u\}) & \text{if } \mathbb{I}_v^{OL} = 0 \\ \rho'_{X'}(v, \{u\}) & \text{if } \mathbb{I}_v^{OL} = 1 \end{cases} \tag{3}$$

and $\overline{\mu}(S_M) = \mathbb{E}[\overline{\mu}_X(S_M)]$.

Lemma 2. *$\overline{\mu}(S_M)$ is submodular.*

Reverse Delayed Reward Sets. State-of-the-art solutions for the IM problem are based on the concept of *Reverse Reachable (RR)* sets. We denote a possible world under the IC model by W.

Definition 1 (Reverse Reachable Set [42]). *The reverse reachable (RR) set for a root node v in W is the set of nodes that can reach v in W. That is, for each node u in the RR set, there is a directed path from u to v in W.*

The connection between RR sets and a node's activation is formalized in the following equivalence lemma. It is the key result that underpins the approximation guarantees of [30, 41, 42] by establishing an unbiased estimator for the influence objective.

Lemma 3 (Activation Equivalence [5]). *For any seed set S and node v, the probability that an influence propagation process from S can activate v equals the probability that S overlaps an RR set for v.*

For our objective function, we seek an analog to the RR set definition in the TCIC setting. Importantly, since RR sets are only concerned with the coverage status of a node u w.r.t. an RR set R, the presence of u in R implies this condition is satisfied. By contrast, our analog must be able to express the *reward* associated with each node present in the set. To overcome this challenge, we introduce the notion of a *Reverse Delayed Reward (RDR)* set for a node v in a possible world X of our TCIC model. RDR sets augment the traditional RR sets by including *delayed reward* information associated with each node in the RDR set. They can be viewed as a *weighted* version of RR sets.

Definition 2 (Reverse Delayed Reward Set). *The reverse delayed reward (RDR) set for node v in X is the set of pairs $(u, \rho_X(v, \{u\}))$ of nodes that can reach v in X and their associated reward (interpreted as a weight). For each node u in the RDR set, there exists a path P in X from u to v for which all tie-breaks along P are won by M when $\{u\}$ is initially activated in M, achieving reward $\rho_X(v, \{u\})$, where $\rho_X(v, \{u\})$ is the reward computed in possible world X.*

Note that RDR sets are defined w.r.t. a fixed seed set S_F given by the problem instance. E.g., in Fig. 1, let all propagation probabilities be 1 and $S_F = \{v_0\}$. Then the RDR set for node v_3 in this possible world is $\{(v_{10}, 1), (v_{14}, 1), (v_2, 1)\}$. An M campaign started at any other node either does not reach v_3 or reaches it too late. We make the following important observation: the delayed-distance $dd_X(u, v)$ from u to v in X is necessary, but not sufficient, information for determining the delay of M reaching v w.r.t. F. The simultaneous propagation of F and M in X can lead to interactions resulting in a node's AW opening before either campaign would

have triggered it when propagating independently. As a result, this may lead to a delay that is not simply $dd_X(u,v) - dd_X(S_F,v)$. Therefore, we cannot compute the reward associated with node u directly from the delayed-distance. A random RDR set is defined in a similar fashion to a random RR set where the "root" node v is chosen at random from G. Since the mitigation objective is not submodular (Theorem 1), we cannot leverage RDR sets directly for maximizing $\mu(.)$. Instead, we will establish a connection between RDR sets and the bounding functions $\overline{\mu}(\cdot)$ and $\underline{\mu}(\cdot)$.

Reward Equivalence. We say a set S *covers* an RDR set R with weight $\omega_R > 0$, provided $\exists u \in S$ such that the pair (u, ω_R) appears in R and ω_R is the largest weight over all nodes $u \in S$. Abusing notation, we write this as $S \cap R = \omega_R$. If there are no pairs (u, ω_R) in R such that $u \in S$, then we define $\omega_R = 0$ and say S *does not cover* R. Next, we establish reward equivalence for the two bounds.

Lemma 4. *Let S_M be a fixed set of nodes, and v be a fixed node. Suppose that we generate an RDR set $\underline{R}$ for v in a possible world X. Let ϱ_1 be the probability that S_M covers $\underline{R}$ with weight $\omega_{\underline{R}}$, and ϱ_2 be the probability that S_M, when used as a seed set for campaign M, achieves a reward $\omega_{\underline{R}}$ at v in a propagation process on G w.r.t $\underline{\mu}(\cdot)$. Then, $\varrho_1 = \varrho_2$.*

Lemma 5. *Let S_M be a fixed set of nodes, and v be a fixed node. Suppose that we generate an RDR set $\overline{R}$ for v in a possible world X' where X' is the modified possible world constructed from possible world X sampled from G. Let ϱ_1 be the probability that S_M covers $\overline{R}$ with weight $\omega_{\overline{R}}$, and ϱ_2 be the probability that S_M, when used as a seed set for campaign M, achieves a reward $\omega_{\overline{R}}$ at v in a propagation process on G w.r.t. $\overline{\mu}(\cdot)$. Then, $\varrho_1 = \varrho_2$.*

Lemmas 4 and 5 extend Lemma 3 by establishing a connection between the probability of a node receiving a particular reward value and RDR coverage weights. As remarked in § 2.1, in place of step reward, we could use arbitrary non-increasing reward functions and find corresponding bounding functions needed for SA using the step functions corresponding to Riemann sums [16] traditionally used in numerical integration (more details in [37]).

4 IMPORTANCE SAMPLING

In this section, we describe how importance sampling can be used in our framework to reduce the sample complexity by analyzing the variance of the random variables associated with our unbiased estimator. Having established reward equivalence for both our upper and lower bounding functions, all of the analysis in this section applies to both functions. Therefore, to simplify the exposition, we describe the idea behind importance sampling and the resulting reverse sampling framework for a single abstracted objective $\sigma(\cdot)$ which could be instantiated as the upper or the lower bounding function, $\underline{\mu}(\cdot)$ or $\overline{\mu}(\cdot)$, of the mitigation objective.

Unbiased Estimators. Unlike the IM problem, where all nodes are candidates to be influenced, in the the MM problem, only those nodes that are influenced by the misinformation are candidates to contribute reward. Unlike targeted IM [24], these nodes are not known a priori nor can they be precomputed. As such, the uniform sampling approach leveraged by random RR sets for the IM problem is not directly useful for our setting. In theory, we can apply Rejection Sampling (RS) by selecting source nodes v for our random RDR sets uniformly at random from G and define the corresponding random RDR set as empty if $v \notin R_F^X$. However, RS is best suited when the target probability is high and becomes less practical as the events become rarer.

A more sophisticated approach that yields improved sample efficiency for estimating rare events is Importance Sampling (IS). IS has been successfully leveraged for the *targeted* IM problem [24]. IS estimates the expected value of a function f in a probability space P via sampling from another *proposal distribution* Q, then re-weights the samples by an *importance factor* for unbiased estimation. When applying IS, ideally Q is chosen to support efficient sampling.

IS for RDR Sets. To apply IS in our setting we select the root for a random RDR set uniformly at random from the set R_F^X. This ensures that the corresponding RDR set is non-empty. In other words, let P be the probability space of RDR sets generated when the root node is selected uniformly at random from G. Then, we define Q as a subspace of P that corresponds to the space of *only* non-empty samples of P, i.e., those RDR sets for which the root $v \in R_F^X$. It remains to define an appropriate importance factor to ensure we have an unbiased estimator for $\sigma(S_M)$. We let INF_F denote the expected number of activated *non-seed* nodes due to seed set S_F in the absence of any mitigating campaign. In other words, INF_F is the expected influence of campaign F in the absence of any mitigating campaign while ignoring the activation of seed nodes. Consider a set S and a random RDR set $R_i(v)$ rooted at v generated with IS as defined above. Define the following random variable:

$$Y_i(S) = \begin{cases} S \cap R_i(v) = \omega_{R_i(v)} & \text{if } S \text{ covers } R_i(v) \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

Then, we have the following lemma.

Lemma 6. *Given a random RDR set $R_i(v)$ generated with importance sampling rooted at v, for any set $S \subseteq V$, we have, $\sigma(S) = \mathbb{E}[Y_i(S)] \cdot INF_F$.*

Lemma 6 states that we can estimate the expected reward of the mitigation campaign using random RDR sets generated with IS. Let $\mathcal{R}$ be a collection of θ random RDR sets generated with IS and let $\mathcal{W}_{\mathcal{R}}(S)$ be the total weight of RDR sets in $\mathcal{R}$ covered by a node set S. Then, based on Lemmas 4, 5 and 6, we can prove:

Corollary 1. $\mathbb{E}\left[\frac{\mathcal{W}_{\mathcal{R}}(S)}{\theta}\right] \cdot INF_F = \sigma(S)$

Concentration Bounds. Next, we analyze the random variables associated with random RDR sets generated using IS. In particular, we show they have smaller variances than random RDR sets generated by RS and, as a consequence, fewer samples are required by our reverse sampling framework. Define the random variable $Z_i(S) = \frac{Y_i(S) \cdot INF_F}{n}$. Notice that the means of $Y_i(S)$ and $Z_i(S)$ are $\mathbb{E}[Y_i(S)] = \frac{\sigma(S)}{INF_F}$ and $\mathbb{E}[Z_i(S)] = \mathbb{E}[Y_i(S)] \cdot \frac{INF_F}{n} = \frac{\sigma(S)}{n}$ respectively. If we construct a set of random variables $Z_1(S), \ldots, Z_\theta(S)$, observe that $\frac{n}{\theta} \sum_{i=1}^{\theta} Z_i(S)$ is an empirical estimate of $\sigma(S)$. An important challenge is that INF_F is #P-hard to compute. We overcome this challenge by computing an approximation of INF_F, denoted $\hat{INF_F}$, and define the random variable $\hat{Z}_i(S) = \frac{Y_i(S) \cdot \hat{INF_F}}{n}$ where $\mathbb{E}[\hat{Z}_i(S)] = \frac{\sigma(S)}{n} \frac{\hat{INF_F}}{INF_F}$. Notice that estimating INF_F is a standard

influence estimation task, where the associated random variables form a *martingale* [42]. Thus, we can leverage existing solutions for the influence estimation problem [29] to efficiently compute an (ϵ, δ)-approximation of INF_F. Suppose we have computed a value for $\hat{INF}_F$ with error ϵ' that holds with probability $1-\delta'$. In order to ease the exposition, we introduce the following definition.

Definition 3. *The* misinformation sampling error ratio *is defined as* $\Gamma = \frac{\hat{INF}_F}{INF_F}$ *where* $(1-\epsilon') \leq \Gamma \leq (1+\epsilon')$ *holds with probability at least* $1-\delta'$.

The variance of $\hat{Z}_i(S)$ satisfies the following inequality.

Proposition 2. $\text{Var}[\hat{Z}_i(S)] \leq 2\Gamma \frac{\sigma(S)}{n} \frac{\hat{INF}_F}{n}$.

Now, we seek a form of Chernoff bounds for random variable $\hat{Z}_i(S)$ so as to provide performance guarantees of our reverse sampling framework. A key requirement is to account for the error associated with the estimation of INF_F. This challenge, unique to our setting, does not arise in previous applications of IS to targeted IM [24] since the target set of nodes is *pre-defined*. We make use of martingale-based concentration bounds in the following.

Definition 4 (Martingale). *A sequence of random variables* $Y_1, Y_2, Y_3, \ldots$ *is a* martingale *if and only if* $\mathbb{E}[|Y_i|] < +\infty$ *and* $\mathbb{E}[Y_i | Y_1, Y_2, \ldots, Y_{i-1}] = Y_{i-1}$ *for any* i.

It is straightforward to show that the random variables $\hat{Z}_i(S)$ form a martingale. Thus, the Chernoff bounds for martingales (see [42], Lemma 2) let us derive the following concentration bounds for the random variables $\hat{Z}_i(S)$ associated with our RDR sets generated with IS, by plugging in the variance derived above. Note, we assume that $\hat{INF}_F \leq \frac{n}{2}$ as a necessary boundary condition for the derivation of our concentration bounds. In our experiments, the condition always held since an unrealistic number of misinformation seeds would be required to influence over half the networks considered.

Lemma 7. *Given a fixed collection of* θ *RDR sets* $\mathcal{R}$ *constructed with importance sampling and seed set* S, *let* $\Lambda(S) = \frac{\hat{INF}_F}{n} \mathcal{W}_{\mathcal{R}}(S)$ *be the normalized weighted coverage of* S *in* $\mathcal{R}$. *For any* $\lambda > 0$,

$$\Pr\left[\Lambda(S) - \sigma(S) \cdot \Gamma \frac{\theta}{n} \geq \lambda\right] \leq \exp\left(\frac{-\lambda^2}{\frac{2}{3}\lambda + 4\sigma(S)\Gamma \frac{\theta}{n} \frac{\hat{INF}_F}{n}}\right)$$

$$\Pr\left[\Lambda(S) - \sigma(S) \cdot \Gamma \frac{\theta}{n} \leq -\lambda\right] \leq \exp\left(\frac{-\lambda^2}{4\sigma(S)\Gamma \frac{\theta}{n} \frac{\hat{INF}_F}{n}}\right)$$

We will make use of the above concentration bounds to derive the approximation guarantees of our sampling framework and to establish appropriate parameter settings.

5 REVERSE SAMPLING FRAMEWORK

Recently, Tang et al. [41] introduced the *OPIM* approach to the online version of the IM problem. Interestingly, an adaptation of *OPIM* to the traditional IM problem yields state-of-the-art performance. Unlike *IMM* [42], which uses the same collection $\mathcal{R}$ of RR sets for constructing the solution seed set S^* and deriving its approximation guarantees, *OPIM* generates a solution on one collection of RR sets $\mathcal{R}_1$ and then derives its approximation guarantees using $\mathcal{R}_1$ and an independent collection of RR sets $\mathcal{R}_2$. In particular, the concentration bounds leveraged by *OPIM* require that S^* be a fixed seed set independent of the RR sets on which it is being evaluated. Intuitively, we can think of $\mathcal{R}_1$ as a set of *nominators* that nominate S^* as the IM solution and $\mathcal{R}_2$ as the set of *assessors* that determine whether S^* is a good enough solution. Notice that, if S^* is not independent of $\mathcal{R}_2$, then the evaluation of S^* could be biased.

We make the following important observation motivating our framework: *the combination of nominators and assessors leveraged by the* OPIM *algorithm does not depend on any particular properties of the IC diffusion model, as long as it satisfies activation equivalence* (Lemma 3). Thus, in view of Lemmas 4 and 5, which are weighted versions of Lemma 3 for our RDR sets in the TCIC model, we can employ a nominator-assessor framework to derive approximation guarantees for our upper and lower bounding mitigation functions. Thus, our technique and results apply to any propagation models satisfying reward equivalence.

MM Solution. Our framework for finding solutions to the upper and lower bounding objectives to the MM problem, *NAMM (Nominators and Assessors for Misinformation Mitigation)* is presented in Algorithm 1. During its execution, *NAMM* invokes the standard greedy algorithm for weighted maximum coverage to obtain a size-k seed set S^*. All of the analysis in this section applies to both of the upper or lower bounding functions, $\underline{\mu}(\cdot)$ or $\overline{\mu}(\cdot)$, of the mitigation objective. Therefore, to simplify the exposition, we refer to a single abstracted objective $\sigma(\cdot)$ in our algorithm description. The approximation guarantee of *NAMM* relies on two critical lemmas that establish an upper bound on the mitigation of the optimal solution S^o ($\sigma^u(S^o)$) and a lower bound on the mitigation of the current solution ($\sigma^l(S^*)$). After establishing an upper bound, N_{max}, on the number of RDR sets required in the worst-case, we show that *NAMM* achieves a $(1-1/e-\epsilon)$-approximation by leveraging our newly defined concentration bounds (Lemma 7). Importantly, the error associated with estimating INF_F must be carefully accounted for in deriving upper and lower bounds to ensure the desired approximation guarantees.

Algorithm 1 NAMM

Input: $G, \epsilon \geq 0, 0 < \delta < 1, k$
Output: An $(1-1/e-\epsilon)$-optimal solution S^*

$\delta' \leftarrow \frac{\delta}{9}, \epsilon' \leftarrow \frac{\epsilon}{2}, \Delta \leftarrow \delta - \delta'$
compute $\hat{INF}_F$; an (ϵ', δ')-approximation of INF_F
set N_{max} according to (5)
$N_0 = N_{max} \cdot \epsilon^2 \frac{LB}{n}$;
generate two collections, $\mathcal{R}_1$ and $\mathcal{R}_2$, of random RDR sets where $|\mathcal{R}_1| = |\mathcal{R}_2| = N_0$;
$i_{max} = \lceil \log_2(\frac{N_{max}}{N_0}) \rceil$;
for $i \leftarrow 1$ **to** i_{max} **do**
    $S^* \leftarrow \texttt{WeightedMaxCover}(\mathcal{R}_1, k, n)$
    compute $\sigma^l(S^*)$ and $\sigma^u(S^o)$ by (6) and (7) respectively, setting $\delta_1 = \delta_2 = \Delta/(3i_{max})$
    $\alpha \leftarrow \sigma^l(S^*)/\sigma^u(S^o)$
    if $\alpha \geq (1-1/e-\epsilon)$ **or** $i = i_{max}$ **then**
        return S^*
    double the sizes of $\mathcal{R}_1$ and $\mathcal{R}_2$ with new random RDR sets

Deriving N_{max}. Tang et al. [42] derived a threshold for the maximum number of RR sets required by their IMM algorithm to ensure that a $(1-1/e-\epsilon)$ approximation guarantee holds with probability at least $1-\delta$. We derive the corresponding threshold for the MM problem. Notably, the derivation requires a careful accounting of the error associated with estimating INF_F.

Lemma 8. *Let $\mathcal{R}$ be a collection of random RDR sets and S^* be a size-k seed set generated by applying the greedy algorithm on $\mathcal{R}$. For fixed ϵ, ϵ', and δ, if $\delta' \leq \frac{\delta}{9}$ and $|\mathcal{R}| \geq \frac{8n(3+\epsilon')(1-1/e)[\ln \frac{9}{4\delta} + \ln \binom{n}{k}]}{3 \cdot OPT[\epsilon(1+\epsilon') - 2\epsilon'(1-1/e)]^2}$ then S^* is a $(1-1/e-\epsilon)$-approximate solution to OPT with at least $1-\delta$ probability.*

Define $\Delta = \delta - \delta'$ and let $LB \leq OPT$ be a lower bound of the optimal mitigation. Based on Lemma 8, we define

$$N_{max} = \frac{8n(3+\epsilon')(1-1/e)[\ln \frac{27}{4\Delta} + \ln \binom{n}{k}]}{3 \cdot LB[\epsilon(1+\epsilon') - 2\epsilon'(1-1/e)]^2}, \tag{5}$$

which is an upper bound on the number of RDR sets needed to guarantee a $(1-1/e-\epsilon)$-approximation w.p. $\geq 1-\Delta/3$.

Deriving LB. In *OPIM*, a crude lower bound of k is used which ensures the number of iterations is bounded by $O(\log n)$. By contrast, due to the objective of the MM problem, the same lower bound is no longer valid. To tackle this, we adopt ideas from the classical *Maximum Influence Arborescence (MIA)* [51] approach to the IM problem to derive a lower bound on OPT. The MIA framework assumes that influence only travels via the paths of maximum influence in the network and leverages the resulting structures to estimate influence spread. A *MIA* structure constructed from a seed set S computes an activation probability $ap(v)$ for each node v in the arborescence. Our idea is to use a *MIA* to lower bound the influence of S_F. Notice that the $ap(v)$ value gives a direct means of estimating the expected reward achieved at v by selecting v as a seed for the mitigating campaign. In particular, we select the top-k nodes at depth 1 in a *MIA* constructed for S_F, ranked by activation probabilities. Thus, the expected mitigation of the selected nodes S_{LB} is at least $LB \geq 2 \cdot \sum_{v \in S_{LB}} ap(v)$. We replace OPT in Equation 5 with LB and note that this setting ensures that the number of iterations is still bounded by $O(\log n)$.

Deriving $\sigma^l(S^*)$. Let $\Lambda_2(S^*) = \frac{I\hat{N}F_F}{n}\mathcal{W}_{\mathcal{R}_2}(S^*)$ be the weighted coverage of S^* in $\mathcal{R}_2$ and $\theta_2 = |\mathcal{R}_2|$. We have the following result.

Lemma 9. *For any $\delta \in (0,1)$ where $a = \ln(1/\delta)$, we have*

$$\Pr\left[\sigma(S^*) \geq \left(\left(\sqrt{\Lambda_2(S^*) + \frac{25a}{36}} - \sqrt{a}\right)^2 - \frac{a}{36}\right)\frac{n}{\theta_2(1+\epsilon')}\right] \geq 1-\delta.$$

Based on Lemma 9 and a parameter δ_2 to be discussed shortly, we set

$$\sigma^l(S^*) = \left(\left(\sqrt{\Lambda_2(S^*) + \frac{25\ln(1/\delta_2)}{36}} - \sqrt{\ln(1/\delta_2)}\right)^2 - \frac{\ln(1/\delta_2)}{36}\right)\frac{n}{\theta_2(1+\epsilon')} \tag{6}$$

Deriving $\sigma^u(S^o)$. We establish an upper bound of $\sigma(S^o)$ from the weighted coverage of S^* in $\mathcal{R}_1$, denoted as $\Lambda_1(S^*) = \frac{I\hat{N}F_F}{n}\mathcal{W}_{\mathcal{R}_1}(S^*)$, by leveraging the property of the greedy algorithm that ensures $\Lambda_1(S^*) \geq (1-1/e)\Lambda_1(S^o)$.

Lemma 10. *Let $\theta_1 = |\mathcal{R}_1|$. For any $\delta \in (0,1)$ where $a = \ln(1/\delta)$, we have*

$$\Pr\left[\sigma(S^o) \leq \left(\sqrt{\frac{\Lambda_1(S^*)}{1-1/e} + a} + \sqrt{a}\right)^2 \frac{n}{\theta_1(1-\epsilon')}\right] \geq 1-\delta.$$

Based on Lemma 10 and a parameter δ_1 to be discussed shortly, we set

$$\sigma^u(S^o) = \left(\sqrt{\frac{\Lambda_1(S^*)}{1-1/e} + \ln(1/\delta_1)} + \sqrt{\ln(1/\delta_1)}\right)^2 \frac{n}{\theta_1(1-\epsilon')} \tag{7}$$

Putting It Together. The reason *NAMM* ensures a $(1-1/e-\epsilon)$-approximation with probability at least $1-\delta$ is as follows. First, the algorithm has at most i_{max} iterations. In each of the first $i_{max}-1$ iterations, Algorithm 1 generates a size-k seed set S^* and derives $\sigma^l(S^*)$ and $\sigma^u(S^o)$ from $\mathcal{R}_2$ and $\mathcal{R}_1$, respectively, setting $\delta_1 = \delta_2 = \Delta/(3i_{max})$. Then, it computes $\alpha \leftarrow \sigma^l(S^*)/\sigma^u(S^o)$ as the approximation guarantee of S^*. By Lemmas 9 and 10, and conditioning on the event $(1-\epsilon') \leq \Gamma \leq (1+\epsilon')$, α is correct with at least $1 - 2\Delta/(3i_{max})$ probability. By the union bound, it has at most $\frac{2\Delta}{3}$ probability to output an incorrect solution in the first $i_{max}-1$ iterations. Meanwhile, in the last iteration, it returns a seed set S^* generated by applying the greedy algorithm on $\mathcal{R}_1$, with $|\mathcal{R}_1| \geq N_{max}$. By Equation 5 and conditioning on the event $(1-\epsilon') \leq \Gamma \leq (1+\epsilon')$, this ensures that S^* is an $(1-1/e-\epsilon)$-approximation with at least $1-\Delta/3$ probability. Therefore, the probability that *NAMM* returns an incorrect solution is at most $\frac{2\Delta}{3} + \Delta/3 + \delta' = \delta$ leading to the following regarding Algorithm 1.

Theorem 2. *Algorithm 1 returns a $(1-1/e-\epsilon)$-approximate solution for $\underline{\mu}$ as well as $\overline{\mu}$ with at least $1-\delta$ probability.*

Runtime. The computation overhead of *NAMM* consists of (i) the generation of RDR sets, (ii) the execution of greedy, and (iii) the computation of $\sigma^l(S^*)$ and $\sigma^u(S^o)$. Both (ii) and (iii) are linear in the number of RDR sets, therefore the time complexity of *NAMM* is $O(\sum_{R \in \mathcal{R}_1 \cup \mathcal{R}_2} EPT)$, where EPT is the expected runtime required to generate an RDR set. Next, we compare EPT to the expected runtime required to generate a traditional RR set used for the IM problem allowing us to gauge the time complexity of the MM problem relative to the IM problem. Let INF_1 be the largest expected influence of any size-1 node set in G under the IC model. We have the following lemma.

Lemma 11. *The expected runtime to generate a RDR set is given by $EPT = O\left(INF_F + \frac{m}{n} \cdot \left(INF_1\right)^2\right)$.*

In contrast, the expected runtime to generate an RR set is $EPT_{RR} = O(\frac{m}{n} \cdot INF_1)$. Thus, in terms of EPT_{RR}, the expected runtime to generate an RDR set is $O(INF_F + EPT_{RR} \cdot INF_1)$. The additional factor of INF_1 in the runtime expression is explained by observing that RDR generation has to potentially perform tie-breaking (linear in INF_1) on every node reverse-reachable from the root of the RDR set. Meanwhile, the additional INF_F term follows from the necessary computation of the set of nodes influenced by F.

Table 2: Network Statistics.

	Gnutella	Flixster	D-B	D-M	DBLP
# nodes	6.3K	7.6K	23.3K	34.9K	317K
# edges	20.8K	75.7K	141K	274K	2.1M
avg. deg	3.3	9.43	6.5	7.9	6.62
type	dir	undir	dir	dir	undir

Sandwich Algorithm. Based on the above results, the *SA* algorithm is given in Algorithm 2, which returns the seed set $S^* \in \{S^L, S^U\}$ that leads to the largest objective value w.r.t. the original objective. The solution produced by *SA* has the following performance guarantee due to [27].

Lemma 12 ([27]). *$\mu(S^*) \geq \beta \cdot (1 - 1/e) \cdot \mu(S^o)$, where $\beta = \max\left\{\frac{\underline{\mu}(S^U)}{\overline{\mu}(S^U)}, \frac{\underline{\mu}(S^o)}{\overline{\mu}(S^o)}\right\}$ and S^o is the optimal MM solution.*

Algorithm 2 Sandwich Approximation($\underline{\mu}, \overline{\mu}, G, k$)

1: $S^L \leftarrow NAMM(\underline{\mu}, G, k)$
2: $S^U \leftarrow NAMM(\overline{\mu}, G, k)$
3: **return** $S^* = \text{argmax}_{S \in \{S^L, S^U\}} \mu(S)$

Anytime Algorithm. Often, in MM applications, it is advantageous to be able to know how close to the optimum the current solution at any time is, so we can stop if desired. The OPIM algorithm developed for IM is such an anytime algorithm. Unlike IM, the non-submodular objective of MM complicates the development of an anytime algorithm. We overcome this challenge as follows.

Lemmas 9 and 10 are the key ingredients required for an anytime algorithm for the MM problem. We can modify Algorithm 1 to stop at a user-specified timestamp and produce a seed set along with its current approximation assurance. However, due to the non-submodular behaviour of the mitigation objective, the anytime algorithm has to generate RDR sets for both the upper and lower bounding functions *simultaneously* in order to apply Algorithm 2 on the resulting anytime solutions. Formally, for each of the submodular bounding functions, given the collections $\mathcal{R}_1$ and $\mathcal{R}_2$ of RDR sets that have been generated so far, our anytime algorithm derives a size-k seed set S^* and derives $\sigma^l(S^*)$ and $\sigma^u(S^o)$ from $\mathcal{R}_2$ and $\mathcal{R}_1$, respectively, and returns $\alpha \leftarrow \sigma^l(S^*)/\sigma^u(S^o)$ as the approximation guarantee of S^*. By Lemmas 9 and 10, α is correct with at least $1 - \delta_1 - \delta_2$ probability. Therefore, setting $\delta_1 = \delta_2 = \delta/2$ suffices to ensure the failure probability does not exceed δ.

6 EXPERIMENTS

Setup. We perform experiments on 5 real networks. The seeds for campaign F are generated in two ways. First, we sample a small number of nodes from the top-k most influential nodes in the network to simulate the spread of fake news by a few popular users in the network. Second, we sample a larger number of users at random from the network to simulate the coordinated spread of misinformation by several bots or newly created puppet accounts. *Figures for each network depict the results for influential fake seeds (left) and random fake seeds (right).* In our experiments, the mitigation achieved by the final solution of each algorithm is evaluated by 20K Monte Carlo simulations. All experiments are performed on a Linux machine with Intel Xeon 2.6 GHz CPU and 128 GB RAM.

Networks. Table 2 summarizes the networks and their characteristics. Flixster is mined from a social movie site and a strongly connected component is extracted. Gnutella is a peer-to-peer file sharing network. Douban is a Chinese social network, where users rate books, movies, music, etc. All movie and book ratings of the users in the graph are crawled separately to derive two datasets from book and movie ratings: Douban-Book and Douban-Movie. Finally DBLP, a peer collaboration network, is a large network that we use to test scalability.

Baselines. Since there is no previous algorithm explicitly addressing the model and reward function considered in this paper, we compare the mitigation achieved by *NAMM* against the following baselines. *IMM* [42] is a state-of-the-art solution for the IM problem. The Influential baseline selects seeds in decreasing order of expected influence. The Proximity baseline selects seed nodes from the out-neighbors of the fake seeds, where a preference is given to those nodes connected by a high probability edge. Random is a baseline method which selects the seeds randomly.

Default Parameters. Following previous works [27, 39, 41, 42, 44, 46, 47] we set edge probabilities for $e = (u, v)$ to $1/indeg(v)$, where $indeg(v)$ is the in degree of node v. Unless otherwise specified, we use $\epsilon = 0.1$ and $\delta = 1/n$ as our default for all methods.

Meeting Probabiltiies. In the same way as [8], where meeting events were introduced, we sample meeting delays from a geometric distribution. We calibrate the meeting probabilities for campaign M based on observations made in [49]. The authors investigate the temporal characteristics of the propagation of true and false news over all the fact-checked cascades that spread on Twitter from 2006 to 2017 totalling over 126,000 cascades. The classification into truth or falsehood was established by six independent fact-checking organizations. The resulting cascades were compared on depth (number of retweet hops from the source tweet), size (number of users involved in the cascade), maximum breadth, and structural virality. The authors observed that falsehood diffused significantly farther, faster, deeper, and more broadly than the truth. In particular, it was observed that truth rarely reached more than 1500 users and it took the truth about six times as long as falsehood to reach 1500 people. Based on these observations, the meeting delay distribution is parameterized by success probability $m(e) = 1/6$ so that, on average, the misinformation propagates 6× faster than the truth.

Activation Windows. The activation window lengths are generated by a two-step procedure. First, we simulate if a user reads the linked content following observations in [13] on real-world click-through behaviours. The authors present a large scale, unbiased study of social clicks by gathering a month of web visits to online resources mentioned in Twitter. Their dataset covers 2.8 million shares and 9.6 million actual clicks. The authors estimate that a majority (59%) of URLs mentioned on Twitter are not clicked at all. Informed by these observations, we first flip a biased coin with probability 0.6 and set the AW length to 0 if the flip succeeds.

Meanwhile, non-zero AW lengths are generated based on observations made in [28] on real-world reading times of social media users. The authors utilize audience behaviour metrics provided by the web analytics firm Parse.ly covering 117 million anonymized,

complete cellphone interactions with 74,840 articles from 30 news websites in the month of September 2015. Overall, short-form news stories (< 1000 words) represent 76% of the articles while long-form articles (≥ 1000 words) account for the remaining 24% of articles. The authors observe that on short-form and long-form articles users spend on average 57 seconds and 123 seconds reading, respectively. Further, a more fine-grained breakdown of average reading times reveals a distribution that closely resembles a geometric distribution. Based on these observations, we generate a sample distribution of reading times from a geometric distribution parameterized by $p = 1/57$ with probability 0.76 and $p = 1/123$ with probability 0.24 reflecting the distribution of short-form and long-form articles. Finally, we apply the bias-corrected maximum likelihood estimator for geometric distributions [9] to learn a AW length parameterization of 74 seconds.

Finally, AW lengths must be normalized to the length of a single hop in our model, which corresponds to the base propagation rate of misinformation measured in seconds. We learned a normalization factor from a collection of retweet cascades of misinformation crawled from Twitter during Oct. 10–Nov. 10, 2020. We extract the distribution of retweet intervals from the cascades and clean the distribution by removing the effect of bots that are programmed to instantly retweet particular accounts, by removing intervals of length less than 3 seconds. We observed that the resulting distribution closely resembles a geometric distribution. As a result, we apply the bias-corrected maximum likelihood estimator for geometric distributions [9] to learn a base propagation rate of 200 seconds for campaign F.

Mitigation Results. Across all datasets our *NAMM* algorithm significantly outperforms the baselines in terms of mitigation achieved. Among the baselines, *IMM* typically outperforms the rest and even produces the best mitigation on Douban-Movie under random fake seeds illustrating that, in some specific scenarios, a "blanket" approach that tries to spread the truth blindly to as much of the network as possible may outperform a more targeted approach. Fig. 3, 4, & 5 also show lower bounds on the data-dependent approximation guarantee (Lemma 12) achieved by *NAMM*. Specifically, we compute $\frac{\mu(S^U)}{\bar{\mu}(S^U)}$ which acts as a lower bound on the data-dependent guarantee of *NAMM*. Across all datasets we observe that $\beta > 0.6$ and that β is typically better when the fake seeds are chosen at random. In Fig. 6 we plot the reward breakdown achieved by *NAMM* on Gnutella and Douban-Book where the blue (red) component refers to the the fraction of nodes for which a reward of 2 (1) was achieved. We observe a strong tendency for reward 2 nodes to dominate the total reward breakdown. In particular, on Gnutella, over 99% of the reward is due to winning outright under both influential and random fake seeds.

Running Time. Next, we investigate the running time achieved when leveraging importance sampling (IS) compared to rejection sampling (RS). We observe three interesting parameter domains exhibiting different behaviours.

First, on the medium datasets (Fig. 7, middle two), we observe that RS is typically faster than IS. To explain, consider the competing mechanisms at play that contribute to running time. On the one hand, RS typically requires more iterations of *NAMM* to terminate due to the possibility of generating "empty" RDR sets that do not contribute reward signal. On the other hand, IS incurs additional runtime overhead as it is required to maintain and manage significantly more state than RS. Recall, under IS, the root of the RDR set is selected from R_F^X. Thus, in order to faithfully represent the possible world, the traversal tree from this first phase must be constructed and stored for every sample generated under IS. Further, in order to ensure efficient edge liveness lookups in the backward traversal, the edges traversed by F must be sorted, incurring additional overheard. Note, we also tested storing these live edges in a hashmap, but observed an increase in runtime.

Interestingly, this trend is reversed in one parameter domain and mixed in another. First, while *NAMM* terminates as soon as a $(1 - 1/e - \epsilon)$ approximation guarantee is achieved for each of the upper and lower bounding functions, the algorithm can be run beyond such a threshold. In the left plots of Fig. 7 we observe that there is not a clear winner when the guarantee α is increased beyond $(1 - 1/e - \epsilon)$. Second, on the large dataset (right of Fig. 7) we see that IS outperforms RS for larger seed set sizes. In this regime, the cost of additional iterations of *NAMM* outweighs the cost of additional state management.

Table 3: Effect of varying meeting delays on Flixster.

	FSTop				FSRan			
ML	*NAMM*	Inf	Prox	Ran	*NAMM*	Inf	Prox	Ran
6	(709.34)	(461.70)	(162.93)	(35.72)	(186.43)	(101.63)	(74.56)	(8.75)
5	+7.81	+14.19	+2.56	-1.40	+2.12	+2.95	+0.22	-1.51
4	+17.17	+24.01	+2.68	-9.14	+4.31	+6.41	-0.03	+1.63
3	+14.11	+30.53	+2.44	+9.18	+2.5	+6.95	+0.24	-0.08
2	+39.62	+50.85	+6.47	+33.22	+8.67	+12.58	+0.40	-1.14
1	+76.38	+117.34	+13.39	-29.27	+17.63	+27.29	+0.54	+4.59

Alternative Parameter Settings. We conduct experiments that consider how the ability to mitigate the spread of fake news is impacted by varying the reading probabilities, activation window lengths and meeting delays. In particular, we are interested in determining if mitigation can improve when users are more likely to click through and read an article, spend more time reading and considering the content and/or are exposed to true information with a similar propagation rate as false information.

Table 3 shows the effects of reducing the meeting delay disadvantage incurred by campaign M. In particular, we see that as the propagation rate of truth and misinformation approach an equal footing, there is a steady increase in the mitigation achieved by *NAMM* and Inf. Furthermore, we observe that the marginal gain achieved increases as the propagation rate of campaign M approaches that of F. Meanwhile, we observe that the Random and Prox baselines do not exhibit monotone behaviour when varying the meeting delay. Table 4 shows the effects of elongating the AW length. We see a similar trend where the mitigation gain of *NAMM* grows at an increasing rate with longer AW lengths. Meanwhile, the baselines all exhibit non-monotone behaviour with varying AW lengths and lower total gains. Finally, we did not observe a noticeable change in mitigation when varying the reading probability. To explain, notice that even when the reading probability succeeds (i.e., AW length is positive), the AW length sampled is unlikely to exceed a few hops in the propagation model. As such, there is little effect on the overall mitigation as the reading probability goes to 1.

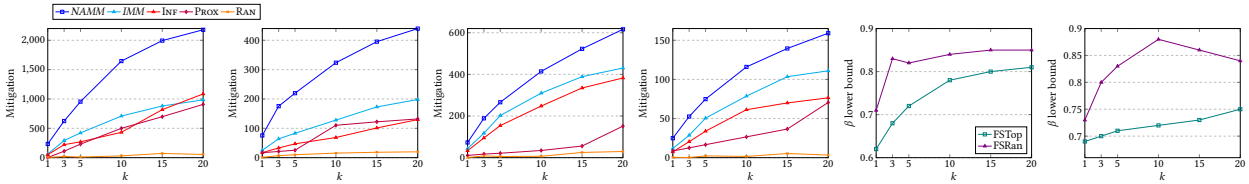

Figure 3: Mitigation on Gnutella (left) & Flixster (middle) with 10 (50) influential (random) fake seeds. Lower bound of β for *SA* (right).

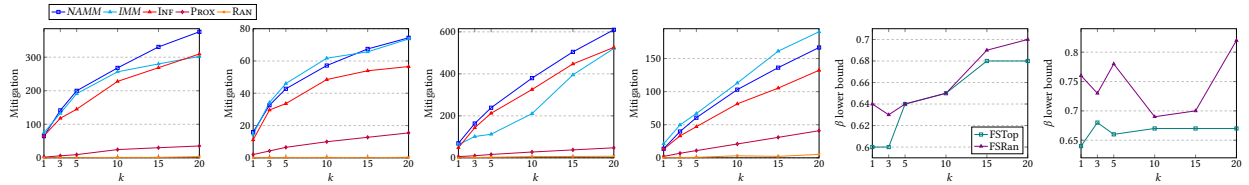

Figure 4: Mitigation on D-B (left) & D-M (middle) with 50 (200) influential (random) fake seeds. Lower bound of β for *SA* (right).

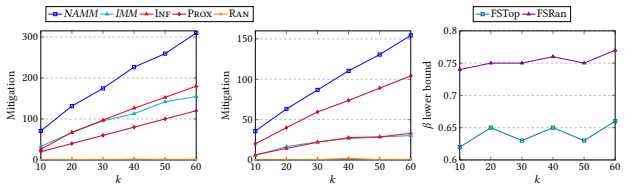

Figure 5: Mitigation on DBLP with 50 (200) influential (random) fake seeds and lower bound of β.

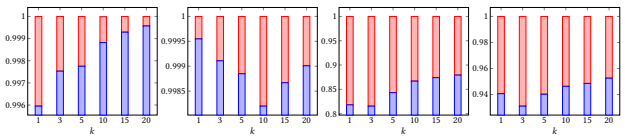

Figure 6: Reward breakdown on Gnutella (left) & Douban-Book (right) with influential & random fake seeds respectively.

Table 4: Effect of varying activation window length on Flixster.

	FSTop				FSRan			
AW	*NAMM*	INF	PROX	RAN	*NAMM*	INF	PROX	RAN
30	(719.47)	(374.44)	(150.28)	(25.52)	(191.36)	(74.87)	(70.30)	(3.71)
60	+10.51	+12.72	-0.76	-10.04	+0.12	-0.03	+0.45	+4.87
120	+24.41	-8.10	-0.51	+7.39	+4.47	+1.87	-0.23	-3.73
240	+48.56	+7.94	+2.23	+4.68	+8.6	+1.95	+0.53	+0.47
480	+73.59	+15.70	+0.24	-3.70	+17.07	+3.06	-0.71	+1.47

In Fig. 8 and 9 we show the mitigation achieved under two alternative models: ego-centric meeting events and fixed edge probabilities respectively. First, under ego-centric meeting events, we consider $m(u,v) = c/(d_{out}(u)+c)$, since it is reasonable to deem that the more friends u has, the less probable that u could meet a certain individual in one time unit. Here c is a smoothing constant and we set it according to [8] which introduced meeting events in a single campaign setting. We find that the mitigation behaviour closely matches the results without ego-centric meeting events (Fig. 3) both in terms of outperforming the baselines and in absolute reward values. Second, we consider fixing edge probabilities to $p = 0.1$ and continue to observe superior performance by *NAMM*.

Robustness. In real deployment, we may not have the exact ground truth values of the parameters. How robust are our solutions in the face of possibly imperfect temporal parameter settings? Also, how

Table 5: Jaccard similarity of seed sets for ground truth, 5/10% perturbed, and no temporal parameter values on Flixster.

	FSTop			FSRan		
k	P5	P10	CIC	P5	P10	CIC
1	1.00	1.00	1.00	1.00	1.00	1.00
3	1.00	1.00	0.50	1.00	1.00	0.50
5	1.00	1.00	0.55	1.00	1.00	0.67
10	1.00	0.91	0.82	1.00	1.00	0.67
15	0.88	0.94	0.67	0.94	0.94	0.67
20	1.00	0.86	0.74	0.86	0.86	0.63

Table 6: Jaccard similarity of seed sets for ground truth, 5/10% perturbed, and no temporal parameter values on Douban-Book.

	FSTop			FSRan		
k	P5	P10	CIC	P5	P10	CIC
1	1.00	1.00	1.00	1.00	1.00	1.00
3	1.00	1.00	1.00	1.00	1.00	0.50
5	1.00	1.00	1.00	0.84	0.84	0.67
10	0.91	0.91	0.67	1.00	1.00	0.67
15	1.00	1.00	0.94	1.00	1.00	0.67
20	0.90	0.90	0.74	0.95	0.90	0.78

do our solutions compare with models such as CIC which have no temporal parameters? In Tables 5 and 6 we show the Jaccard similarity between seed sets selected under a fixed ground truth, corresponding to the default parameter settings of TCIC, and those selected when the temporal parameters are perturbed and/or ignored. Specifically, we generate solutions after perturbing each of the meeting delay, activation window length, reading probability, and base propagation rate parameters with 5 & 10% Gaussian noise (P5 and P10). Additionally, we consider solutions generated with all temporal parameters "turned off", i.e., all meetings delays set to 1 and activation window lengths of 0, which reduces TCIC to CIC. We observe that perturbed solutions retain high similarity with the ground truth seed sets across both influential and random fake seeds. Furthermore, the similarity is always greater compared to seed sets generated under CIC, highlighting the importance of capturing the differential propagation rates of truth and misinformation. Finally, the relative drop in mitigation for solutions generated under CIC is up to 20% and 4.62% on Flixster and Douban-Book respectively.

7 RELATED WORK

Influence Maximization. The IM problem was formulated as a discrete optimization problem by Kempe et al. [19] where the

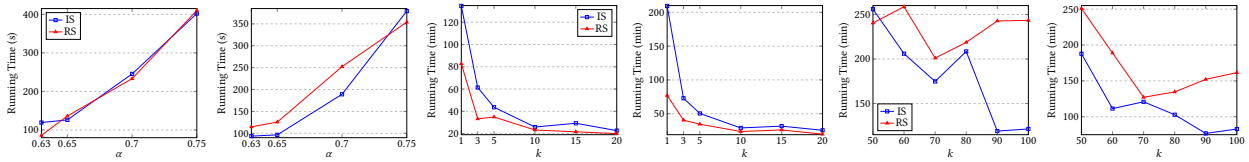

Figure 7: Running time (left) under varying α's on Flixster, (middle) on D-M and (right) DBLP with 50 (200) influential (random) fake seeds.

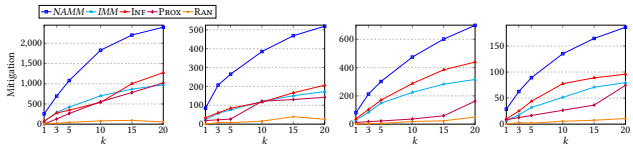

Figure 8: Mitigation under ego-centric meeting delays on Gnutella (left) & Flixster (right).

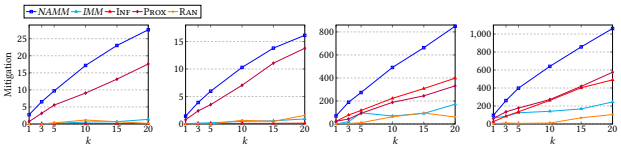

Figure 9: Mitigation under fixed propagation probabilities on Gnutella (left) & Flixster (right).

independent cascade (IC) and linear threshold (LT) models were introduced. Since then, various aspects of IM, and its variants, have been extensively studied (see [7, 23] for surveys). State-of-the-art IM solutions [15, 30, 41–43] rely on reverse sampling for their efficiency. Our edge-level time-delayed propagation is similar to the diffusion dynamics captured by the IC-M model of [8] set in a single campaign propagation setting to model the log-in and log-out behaviour of users. The concept of a time-sensitive reward function was considered in [21, 26] to better model time-sensitive information such as product sales. IM under competition is studied in [3, 25, 27] among others. Finally, models that distinguish between awareness and adoption have been considered in [2, 27].

Misinformation Mitigation. The MM problem was first studied under an independent cascade model by Budak et al. [6] and under a linear threshold model in [11, 14]. In both settings, the objective is shown to be monotone and submodular, thus the greedy algorithm provides a $(1 - 1/e)$-approximation. Subsequently, there have been a number of works that either study variants of the classical MM problem or improve the running time of the greedy approach. The related problem of determining the budget required to reach a threshold mitigation level is investigated in [11, 33, 34].

Each of [36, 39, 40, 44, 46, 47] considers MM variants under models that ignore the temporal nature of misinformation propagation. Roughly speaking, they parallel the improvements made in reverse sampling frameworks developed for the IM problem. Most relevant, [40] uses a competitive IC model augmented with meeting events that are shared by *both* campaigns. Thus, the observed difference in propagation rates between fake and true information is not captured. Notably, none of the above studies incorporates activation windows or a reward function that is time-critical.

Finally, related versions of the MM problem have been investigated by other communities including crowd-sourced mitigation [22, 48], epidemiology [20, 35, 38, 50, 52, 53], and ML [12, 45]. In particular, the epidemiology community focuses on *preemptive* strategies, without propagation for the mitigating side, that limit the susceptibility of a network, and consider propagation models without the temporal notions we study. The ML approaches aim to learn mitigation strategies trained on a fixed input graph which limits their transferability to new graph instances. Recently, Juul et al. [17] repeat the comparison of the spread of fake and true news conducted by [49] on two subsampled datasets with exactly the same size distribution. Interestingly, they find that *under these conditions*, the propagation characteristics become indistinguishable, which has important consequences for misinformation detection. We note that our work focuses on the general case where the propagation differences observed by Vosoughi et al. [49] do hold. In particular, the authors of [17] confirm that false and true information propagate at different rates when the naturally occurring cascade size distributions remain unaltered.

8 CONCLUSION AND FUTURE WORK

In this paper, we address a major shortcoming of existing MM propagation models by introducing the TCIC model, which captures important temporal aspects of fake news diffusion and formulate a time-sensitive variant of the MM problem. We prove our mitigation objective is non-submodular and develop submodular upper and lower bounding functions to sandwich the objective and provide data-dependent approximation guarantees. Finally, we propose a reverse sampling framework that provides $(1 - 1/e - \epsilon)$-approximate solutions to our bounding functions and present an anytime version of our approach. Using experiments over five datasets, we demonstrate that our *NAMM* algorithm outperforms various baselines including those that are oblivious to time-critical aspects.

The techniques developed here are applicable to other problem settings. In the filter bubble problem, multiple conflicting opinions propagate in a social network and the goal is to ensure balanced exposure. As the set of users adopting any particular opinion is stochastic, our *NAMM* algorithm can be applied to help obtain a more balanced exposure. Further, revisit the "classic" competitive IM problem, where realistically the diffusion rates of companies/brands may differ owing to differences in reputation and marketing strategies. In this case, *NAMM* can be applied for the competitive IM problem from a follower perspective. Notice that previous techniques for competitive IM do not apply to this setting.

For simplicity of exposition, we considered a fixed set of fake seeds. However, our proposed solution retains its guarantees even when the fake seeds are not known exactly but are chosen from a distribution. It is interesting to study the scenario where the fake seed set is dynamically evolving. This is closely tied with adaptive influence maximization, and we leave it for future work.

REFERENCES

[1] George M Beal and Joe M Bohlen. 1956. *The diffusion process.* Technical Report.

[2] Smriti Bhagat, Amit Goyal, and Laks VS Lakshmanan. 2012. Maximizing product adoption in social networks. In *Proceedings of the fifth ACM international conference on Web search and data mining*. 603–612.

[3] Shishir Bharathi, David Kempe, and Mahyar Salek. 2007. Competitive influence maximization in social networks. In *WINE'07* (San Diego, CA, USA). Springer-Verlag, Berlin, Heidelberg, 306–311. http://dl.acm.org/citation.cfm?id=1781894.1781932

[4] Andrew An Bian, Joachim M Buhmann, Andreas Krause, and Sebastian Tschiatschek. 2017. Guarantees for greedy maximization of non-submodular functions with applications. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 498–507.

[5] Christian Borgs, Michael Brautbar, Jennifer Chayes, and Brendan Lucier. 2014. Maximizing social influence in nearly optimal time. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*. SIAM, 946–957.

[6] C. Budak, D. Agrawal, and A. El Abbadi. 2011. Limiting the spread of misinformation in social networks. In *WWW'11*.

[7] Wei Chen, Laks V. S. Lakshmanan, and Carlos Castillo. 2013. *Information and Influence Propagation in Social Networks*. Morgan & Claypool Publishers. 1–177 pages.

[8] Wei Chen, Wei Lu, and Ning Zhang. 2012. Time-critical influence maximization in social networks with time-delayed diffusion process.. In *AAAI*, Vol. 2012. 1–5.

[9] David R Cox and E Joyce Snell. 1968. A general definition of residuals. *Journal of the Royal Statistical Society: Series B (Methodological)* 30, 2 (1968), 248–265.

[10] Ullrich KH Ecker, Stephan Lewandowsky, and David TW Tang. 2010. Explicit warnings reduce but do not eliminate the continued influence of misinformation. *Memory & cognition* 38, 8 (2010), 1087–1100.

[11] Lidan Fan, Zaixin Lu, Weili Wu, Bhavani Thuraisingham, Huan Ma, and Yuanjun Bi. 2013. Least cost rumor blocking in social networks. In *Distributed Computing Systems (ICDCS), 2013 IEEE 33rd International Conference on*. IEEE, 540–549.

[12] Mehrdad Farajtabar, Jiachen Yang, Xiaojing Ye, Huan Xu, Rakshit Trivedi, Elias Khalil, Shuang Li, Le Song, and Hongyuan Zha. 2017. Fake news mitigation via point process based intervention. In *International Conference on Machine Learning*. PMLR, 1097–1106.

[13] Maksym Gabielkov, Arthi Ramachandran, Augustin Chaintreau, and Arnaud Legout. 2016. Social clicks: What and who gets read on Twitter?. In *Proceedings of the 2016 ACM SIGMETRICS international conference on measurement and modeling of computer science*. 179–192.

[14] Xinran He, Guojie Song, Wei Chen, and Qingye Jiang. 2012. Influence blocking maximization in social networks under the competitive linear threshold model. In *Proceedings of the 2012 SIAM International Conference on Data Mining*. SIAM, 463–474.

[15] Keke Huang, Sibo Wang, Glenn Bevilacqua, Xiaokui Xiao, and Laks VS Lakshmanan. 2017. Revisiting the stop-and-stare algorithms for influence maximization. *Proceedings of the VLDB Endowment* 10, 9 (2017), 913–924.

[16] Deborah Hughes-Hallett, Andrew M Gleason, and William G McCallum. 2020. *Calculus: Single and multivariable*. John Wiley & Sons.

[17] Jonas L. Juul and Johan Ugander. 2021. Comparing information diffusion mechanisms by matching on cascade size. *Proceedings of the National Academy of Sciences* 118, 46 (2021). https://doi.org/10.1073/pnas.2100786118 arXiv:https://www.pnas.org/content/118/46/e2100786118.full.pdf

[18] Shlomo Kalish. 1985. A new product adoption model with price, advertising, and uncertainty. *Management science* 31, 12 (1985), 1569–1585.

[19] David Kempe, Jon Kleinberg, and Éva Tardos. 2003. Maximizing the spread of influence through a social network. In *KDD'03*. https://doi.org/10.1145/956750.956769

[20] Elias Boutros Khalil, Bistra Dilkina, and Le Song. 2014. Scalable diffusion-aware optimization of network topology. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. 1226–1235.

[21] Arijit Khan. 2016. Towards time-discounted influence maximization. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*. 1873–1876.

[22] Jooyeon Kim, Behzad Tabibian, Alice Oh, Bernhard Schölkopf, and Manuel Gomez-Rodriguez. 2018. Leveraging the crowd to detect and reduce the spread of fake news and misinformation. In *Proceedings of the eleventh ACM international conference on web search and data mining*. 324–332.

[23] Yuchen Li, Ju Fan, Yanhao Wang, and Kian-Lee Tan. 2018. Influence maximization on social graphs: A survey. *IEEE Transactions on Knowledge and Data Engineering* 30, 10 (2018), 1852–1872.

[24] Yuchen Li, Dongxiang Zhang, and Kian-Lee Tan. 2015. Real-time targeted influence maximization for online advertisements. *Proceedings of the VLDB Endowment* (2015).

[25] Yishi Lin and John CS Lui. 2015. Analyzing competitive influence maximization problems with partial information: An approximation algorithmic framework. *Performance Evaluation* 91 (2015), 187–204.

[26] Bo Liu, Gao Cong, Dong Xu, and Yifeng Zeng. 2012. Time constrained influence maximization in social networks. In *2012 IEEE 12th international conference on data mining*. IEEE, 439–448.

[27] Wei Lu, Wei Chen, and Laks VS Lakshmanan. 2015. From competition to complementarity: comparative influence diffusion and maximization. *Proceedings of the VLDB Endowment* 9, 2 (2015), 60–71.

[28] Amy Mitchell, Galen Stocking, and Katerina Eva Matsa. 2016. Long-form reading shows signs of life in our mobile news world. *Pew Research Center* 5 (2016).

[29] Hung T Nguyen, Tri P Nguyen, Tam N Vu, and Thang N Dinh. 2017. Outward influence and cascade size estimation in billion-scale networks. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 1, 1 (2017), 1–30.

[30] Hung T Nguyen, My T Thai, and Thang N Dinh. 2016. Stop-and-stare: Optimal sampling algorithms for viral marketing in billion-scale networks. In *Proceedings of the 2016 International Conference on Management of Data*. 695–710.

[31] Nam P. Nguyen, Guanhua Yan, My T. Thai, and Stephan Eidenbenz. 2012. Containment of Misinformation Spread in Online Social Networks. In *Proceedings of the 4th Annual ACM Web Science Conference* (Evanston, Illinois) *(WebSci '12)*. ACM, New York, NY, USA, 213–222. https://doi.org/10.1145/2380718.2380746

[32] Brendan Nyhan and Jason Reifler. 2010. When corrections fail: The persistence of political misperceptions. *Political Behavior* 32, 2 (2010), 303–330.

[33] Canh V Pham, Quat V Phu, and Huan X Hoang. 2018. Targeted misinformation blocking on online social networks. In *Asian Conference on Intelligent Information and Database Systems*. Springer, 107–116.

[34] Canh V Pham, Quat V Phu, Huan X Hoang, Jun Pei, and My T Thai. 2019. Minimum budget for misinformation blocking in online social networks. *Journal of Combinatorial Optimization* 38, 4 (2019), 1101–1127.

[35] B Aditya Prakash, Lada Adamic, Theodore Iwashyna, Hanghang Tong, and Christos Faloutsos. 2013. Fractional immunization in networks. In *Proceedings of the 2013 SIAM International Conference on Data Mining*. SIAM, 659–667.

[36] Akrati Saxena, Wynne Hsu, Mong Li Lee, Hai Leong Chieu, Lynette Ng, and Loo Nin Teow. 2020. Mitigating misinformation in online social network with top-k debunkers and evolving user opinions. In *Companion Proceedings of the Web Conference 2020*. 363–370.

[37] Michael Simpson, Farnoosh Hashemi, and Laks V. S. Lakshmanan. 2022. Misinformation Mitigation under Differential Propagation Rates and Temporal Penalties. https://doi.org/10.48550/ARXIV.2206.11419

[38] Michael Simpson, Venkatesh Srinivasan, and Alex Thomo. 2016. Clearing contamination in large networks. *IEEE Transactions on Knowledge and Data Engineering* 28, 6 (2016), 1435–1448.

[39] Michael Simpson, Venkatesh Srinivasan, and Alex Thomo. 2020. Reverse Prevention Sampling for Misinformation Mitigation in Social Networks. In *23rd International Conference on Database Theory (ICDT 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik.

[40] Chonggang Song, Wynne Hsu, and Mong Li Lee. 2017. Temporal influence blocking: Minimizing the effect of misinformation in social networks. In *2017 IEEE 33rd International Conference on Data Engineering (ICDE)*. IEEE, 847–858.

[41] Jing Tang, Xueyan Tang, Xiaokui Xiao, and Junsong Yuan. 2018. Online processing algorithms for influence maximization. In *Proceedings of the 2018 International Conference on Management of Data*. 991–1005.

[42] Youze Tang, Yanchen Shi, and Xiaokui Xiao. 2015. Influence maximization in near-linear time: A martingale approach. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*. ACM, 1539–1554.

[43] Youze Tang, Xiaokui Xiao, and Yanchen Shi. 2014. Influence Maximization: Near-optimal Time Complexity Meets Practical Efficiency. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data* (Snowbird, Utah, USA) *(SIGMOD '14)*. ACM, New York, NY, USA, 75–86. https://doi.org/10.1145/2588555.2593670

[44] Amo Tong, Ding-Zhu Du, and Weili Wu. 2018. On misinformation containment in online social networks. In *Advances in neural information processing systems*. 341–351.

[45] Guangmo Tong. 2020. StratLearner: Learning a Strategy for Misinformation Prevention in Social Networks. In *Advances in Neural Information Processing Systems*, Vol. 33. 15546–15555. https://proceedings.neurips.cc/paper/2020/file/b2f627fff19fda463cb386442eac2b3d-Paper.pdf

[46] Guangmo Tong, Weili Wu, Ling Guo, Deying Li, Cong Liu, Bin Liu, and Ding-Zhu Du. 2017. An efficient randomized algorithm for rumor blocking in online social networks. *IEEE Transactions on Network Science and Engineering* 7, 2 (2017), 845–854.

[47] Guangmo Amo Tong and Ding-Zhu Du. 2019. Beyond uniform reverse sampling: A hybrid sampling technique for misinformation prevention. In *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*. IEEE, 1711–1719.

[48] Nguyen Vo and Kyumin Lee. 2018. The rise of guardians: Fact-checking url recommendation to combat fake news. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. 275–284.

[49] Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *Science* 359, 6380 (2018), 1146–1151.

[50] Manh M Vu and Huan X Hoang. 2017. Minimizing the spread of misinformation on online social networks with time and budget constraint. In *2017 9th International Conference on Knowledge and Systems Engineering (KSE)*. IEEE, 160–165.

[51] Chi Wang, Wei Chen, and Yajun Wang. 2012. Scalable influence maximization for independent cascade model in large-scale social networks. *Data Mining and Knowledge Discovery* 25, 3 (2012), 545–576.

[52] Hong Wu, Zhijian Zhang, Yabo Fang, Shaotang Zhang, Zuo Jiang, Jian Huang, and Ping Li. 2021. Containment of rumor spread by selecting immune nodes in social networks. *Mathematical Biosciences and Engineering* 18, 3 (2021), 2614–2631.

[53] Yao Zhang and B Aditya Prakash. 2015. Data-aware vaccine allocation over large networks. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 10, 2 (2015), 1–32.