

Socially Fair Mitigation of Misinformation on Social Networks via Constraint Stochastic Optimization

Ahmed Abouzeid¹, Ole-Christoffer Granmo¹, Christian Webersik², Morten Goodwin¹

¹Centre for Artificial Intelligence Research (CAIR), University of Agder, Norway

²Centre for Integrated Emergency Management (CIEM), University of Agder, Norway
{ahmed.abouzeid, ole.granmo, christian.webersik, morten.goodwin}@uia.no

Abstract

Recent social networks' misinformation mitigation approaches tend to investigate how to reduce misinformation by considering a whole-network statistical scale. However, unbalanced misinformation exposures among individuals urge to study fair allocation of mitigation resources. Moreover, the network has random dynamics which change over time. Therefore, we introduce a stochastic and non-stationary knapsack problem, and we apply its resolution to mitigate misinformation in social network campaigns. We further propose a generic misinformation mitigation algorithm that is robust to different social networks' misinformation statistics, allowing a promising impact in real-world scenarios. A novel loss function ensures fair mitigation among users. We achieve fairness by intelligently allocating a mitigation incentivization budget to the knapsack, and optimizing the loss function. To this end, a team of Learning Automata (LA) drives the budget allocation. Each LA is associated with a user and learns to minimize its exposure to misinformation by performing a non-stationary and stochastic walk over its state space. Our results show how our LA-based method is robust and outperforms similar misinformation mitigation methods in how the mitigation is fairly influencing the network users.

Introduction

From a computation perspective, there are many approaches to combat the dissemination of misinformation¹. Recently, (De Beer and Matthee 2020) illustrated some of the main techniques for classifying misinformation content and how these approaches can be applied in different scenarios. However, classification methods tend to be offline and limited to particular social network features to be learned, such as linguistics and the local political context (Lazer et al. 2018). Furthermore, such classification models have a potential for *False Positive* matches, which may violate human rights conventions by misjudging and questioning individuals credibility and controlling free speech (Özdan 2021). On the other hand, recent work proposed intervention-based

resolutions as an online approach to mitigate the circulation of misinformation on social media platforms. Such an approach is considered more convenient since it facilitates better collaboration between humans and technology by providing learned misinformation mitigation strategies instead of black-box classification models. For example, (Farajtabar et al. 2017) proposed a reinforcement learning-based optimization method which provides a strategy to decrease the difference between misinformation and true content exposures in Twitter, given that such misinformation exposure was dominating the network. The purpose was to mitigate the effect of misinformation on network users by incentivizing the latter to spread true information. A similar method was developed to facilitate decentralized and faster computation, as proposed by (Abouzeid et al. 2021).

The latter approach introduces a light-weight decentralized computation that reduces the optimization sample space and utilizes Learning Automata (LA) that learn from reinforcement feedback (Narendra and Thathachar 1974). However, the method was evaluated according to the decrease in difference between the dominating misinformation and the incentivized true content, averaging over the whole network. The problem with such an approach is that there would be real-world scenarios where some individuals need mitigation efforts more than others, while a sub-network individuals would be already protected from high misinformation exposures. Therefore, we believe a more socially fair intervention and allocation of mitigation resources should be introduced under the framework of (Abouzeid et al. 2021).

This paper proposes a robust LA-based decentralized mitigation method that addresses a wide range of possible unbalanced exposures to either misinformation or true content, seeking robustness on a variety of a social network's statistics. Our contribution is threefold:

- We propose a novel learning scheme for an LA learning in a stochastic and non-stationary environment. The randomness comes from an information diffusion model based on point processes (Laub, Taimre, and Pollett 2015), while the non-stationarity comes from the temporal changes that occur over the whole network when an individual user responds to incentivization. This non-stationarity is particularly intricate due to the hidden dependencies in the information diffusion model. The LA task is to construct a network of individual automata

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹The term misinformation is sometimes used to refer to all forms of fake news/content. However, in some literature, misinformation is defined as the unintentional spread of false content while disinformation is the on-purpose spread. In this paper, we refer to all forms of false content as misinformation.

on top of the social network. Each individual automaton is associated with a single user and performs a constraint Knapsack optimization via a random walk (Pearson 1905) over the automaton state space.

- We propose a novel loss function to ensure that true content incentivization budget is fairly assigned according to individual users exposure needs. To this end, the problem is defined as a stochastic and non-stationary multi-agent Knapsack (Nicosia, Pacifici, and Pferschy 2009) optimization problem.
- We introduce two evaluation metrics (Achieved Mitigation and Achieved Fairness) to measure the efficiency and robustness of the proposed misinformation mitigation algorithm on different social network's statistics. And we evaluate how our proposed technique is more socially fair compared to the proposed approach in (Abouzeid et al. 2021). We conduct our empirical experiments on both synthetic and real-world social networks. Software source code and data are available at (<https://github.com/Ahmed-Abouzeid/MMSS>).

Preliminaries

Information Diffusion Modelling. In order to apply intervention-based resolutions to misinformation mitigation, an information diffusion model is required to simulate the social network which to intervene with. The simulation is considered because intervention with the actual social media platforms is not feasible. We simulate the process of information diffusion by employing a Multivariate Hawkes process (MHP) as practiced by (Goindani and Neville 2020), (Abouzeid et al. 2021), and (Farajtabar et al. 2017). An MHP is a multivariate stochastic process (Chen 2016) which models the occurrence of temporal or spatio-temporal asynchronous events by capturing the mutual-excitation (dependencies) between these events. To model the social network dynamics, each user is represented by two Hawkes processes (HP), one for misinformation dissemination behavior, and the other for true content. The associated user HPs generate estimated random counts for both information types, given some behaviour observation in the past (e.g., estimating number of re/tweeted events given historical dependency). These counts indicate the intensity of the process at a specific time realization. Hence, An HP can be defined with its conditional intensity function λ . The intensity function has two main components: base intensity μ , and exponential decay function g over an adjacency matrix A . The formal explanation of the conditional intensity function is given by:

$$\lambda_i(t_r|H^{t_r}) := \mu_i + \sum_{t_s < t_r} g(t_r - t_s). \quad (1)$$

Where μ is the base intensity that models some external motivation to propagate some content (independent from inferred relationships in data). On the other hand, g is some kernel function over the observed history H^{t_r} associated with user i from the discrete time realization t_s prior to time t_r . g is concerned with the history of some influence matrix A_i , where $A_{ij} = 1$ if there is an influence indicating that user i influences user j , and $A_{ij} = 0$ if not. We used an exponential

decay kernel function $g = A_i e^{-wt}$ as practiced by (Farajtabar et al. 2017), where w is the decay factor which represents the rate for how the influence is reduced over time. For all users, the base intensity vector μ , and the influence matrix A can be estimated using maximum likelihood as proposed in (Ozaki 1979). To simulate all users behaviours for each content type, an MHP is created, given that different intensity rates are generated at different discrete time realizations. Hence, at each realization, each user behaviour is simulated as an estimated number of events (misinformation or true content) to be generated. We set the interval window between realizations to two hours. The HP simulation algorithm adopted in this study follows the modified thinning algorithm introduced by (Ogata 1981). See Appendix A.1 (Abouzeid et al. 2022) for a detailed explanation of the simulation evaluation metric.

Mitigated Diffusion. The core idea behind misinformation mitigation is by introducing the true information to the network through incentivization. Therefore, users associated true content HPs are modified. Hence, let x_i be the incentivization amount decided for user i , and the modified HP for mitigation purposes can be redefined by:

$$\lambda_i(t_r|H^{t_r}) := x_i + \mu_i + \sum_{t_s < t_r} g(t_r - t_s). \quad (2)$$

Related Work

Misinformation Impact. According to (Bradshaw and Howard 2017), at least 50% of the world's countries suffer from organized political manipulation campaigns over social media. Other examples of misinformation can be observed during the Ebola outbreak in West Africa, which was believed to be three times more worse than the previous Ebola outbreaks (Jin et al. 2014). Therefore, research on the role of online media and border-free passing through messages became an emerging topic of interest in scientific communities. Furthermore, investigation on such a topic is more complicated and requires different perspectives of analysis. For example, recent studies (Rampersad and Althiyabi 2020) argued that the influence of social media on accepting political misinformation may differ depending on age, culture or gender. Such social studies actively investigated the social impact of misinformation propagation on different social media platforms such as Reddit, Facebook, and Twitter. Novel views on the problem emerged recently. For instance, recent investigations reported that deliberation contexts promoted in social media overcome false information about health (Pulido et al. 2020). An example of such deliberation can be viewed as a counterfactual campaigns to spread true health information against the spread of misinformation as practiced for the COVID-19 case on Twitter by (Abouzeid et al. 2021).

Misinformation Detection. The spread of fake news on social media has been initially considered as the intentional dissemination of false content in news articles (Allcott and Gentzkow 2017). Progressively, others gave attention to the broader range of the problem (Sharma et al. 2019; Shu et al. 2017). Moreover, rumor detection (Zhang et al. 2015),

malicious accounts classification (Zannettou et al. 2019; Shao et al. 2018), and the causal aspects of misinformation (Abouzeid et al. 2019) have been discussed. However, the majority of these methods are highly depending on linguistic or local features which cause a lack of generality in the final resolution. To the best of our knowledge, it is hard to solve the problem in real-time or without data selection-bias concerns (Ousidhoum, Song, and Yeung 2020). (Wasike 2013) expressed similar moral concerns since fake news detection resolutions are judgemental by nature. Therefore, the need for safer online strategies that would lead to more generic and authentic resolutions is critically desirable.

Knapsack Optimization. The utilization of Learning Automaton (LA) with Knapsack optimization problems is widely approached in the literature. For instance, (Granmo et al. 2007) worked on optimizing the allocation of polling resources for web page monitoring when the monitoring capacity is restricted. In web page monitoring systems, the system may involve n web pages that are updated on different time intervals. Hence, to avoid involving all web pages including the ones with no updates, the system must determine the most important web pages only, without exceeding the monitoring capacity. The work utilized a team of learning automata, where each automaton is involved with a particular web page and learns its importance to a Knapsack total value. Similarly, (Yazidi and Hammer 2018) dealt with a Stochastic Non-linear Fractional Equality Knapsack (NFEK) problem which is a fundamental resource allocation problem based on incomplete and noisy information. In the latter work, they proposed an optimal resolution to the resource allocation problem using a continuous LA without mapping the Knapsack materials onto a binary hierarchy. In such work, the proposed LA had a Reward-Inaction (R-I) learning scheme which only updates the LA actions (transitions) probabilities when rewarded. (Ulker and Tongur 2017) worked on another combinatorial optimization problem for Knapsack with a proposed Migrating Birds Optimization (MBO) algorithm to solve a 0-1 knapsack problem (Fréville 2004).

Hawkes Processes. The utilization of Hawkes processes-based intervention strategies was effectively presented on minimizing-risk problems. For example, (Gupta et al. 2018) worked on the problem of invasive species spreading to new areas which threatens the stability of ecosystems and causes major economic losses. The latter study proposed a novel approach to minimize the spread of an invasive species given a limited intervention budget, where the spread of species was modelled by a Hawkes process and the minimization task was considered a constraint Knapsack optimization problem.

Methodology

Learning Automata Network. A Learning Automaton (LA) is a stochastic model suitable for learning in random environments (Narendra and Thathachar 1974). The LA learns by interacting with the random environment, and updates its actions or state transitions according to the stochastic signal from the environment. Depending on the automa-

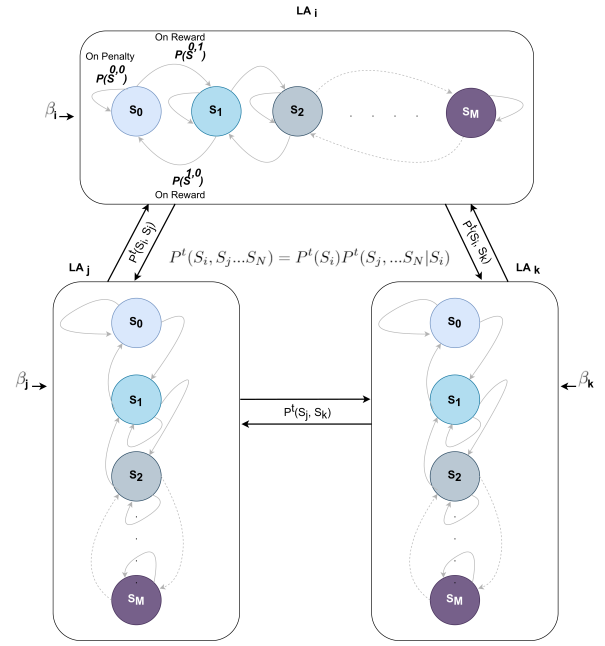


Figure 1: The proposed LA network and the underlying multivariate Markov chain architecture for three automata.

ton design and architecture, the task is to find either an optimum/sub-optimum action or state. The LA seeks convergence to such state or action, eventually. The advantage of utilizing an LA-based optimization is due to its decentralized and easy implementation. An LA defined by its stochastic state transitions can be formally defined as a Markov Process (Ames 1989). Therefore, to reach equilibrium over all LA, we build a network of LA, each performs a random walk over a finite and discrete state space, where the individual optimum or sub-optimum states will be the recommended incentivization values for a misinformation mitigation campaign. The individual random walks together form as a multidimensional joint random walk (Marquioni 2019) modelled by a multivariate Markov chain (Gotzamani et al. 2018). Figure 1 demonstrates the proposed LA network and the underlying multivariate Markov chain (e.g., three automata with M states, each.), where the joint state transitions and their probabilities are derived by the individual automata state transitions which are dictated by a reward signal β .

Learning State Transition. An individual LA_i has a state space with memory depth M , where $M > 0$. If LA_i is in a state S_i^k where $0 < k < M$, then, it has a three possible state transitions: $S_i^{k,k-1}, S_i^{k,k}, S_i^{k,k+1}$ indicating going to left, staying at same state, and moving to the right, respectively. In order to reach an optimum or sub-optimum state S_i^* , LA_i needs to learn the probabilities of its state transitions until it converges. Consequently, the optimum or sub-optimum S_i^* value will be the recommended incentivization value x_i^* to modify the information diffusion model with (See Equation 2). LA_i could have only two possible state transitions: $S_i^{k,k}, S_i^{k,k+1}$ or $S_i^{k,k}, S_i^{k,k-1}$, when $k = 0$ or

$k = M$, respectively. At each interaction step t , the probability of LA_i being in a next state depends on its present state and the transition direction a_i^t . With a uniform initial state transitions probabilities, LA_i determines the next state S_i^{t+1} and updates its state transition probability distribution vector π_i according to the below:

$$\delta : S_i^t, a_i^t, \beta_i^t \rightarrow S_i^{t+1}, \pi_i^{t+1}. \quad (3)$$

Where π_i states probabilities are updated with regard to their rewarded visits frequency, and a_i^t represents the applied state transition $a_i^t = S_i^{k,j}$, where k, j are neighbor state indices and $k = j$ if it was a recurrent state transition. Based on a_i^t and the environment stochastic reward β_i^t , LA_i conducts a random step move over its state space. For instance, if $a_i^{t_0} = S_i^{k,k+1}$, the state transition function δ commits the transition $S_i^k \rightarrow S_i^{k+1}$ only if $\beta_i^{t_0} = 0$, and rolls it back if $\beta_i^{t_0} = 1$. Consequently, $\pi_i^{t_1} = [0^{k,k}, 1^{k,k+1}, \dots]^{t_1}$ or $\pi_i^{t_1} = [1^{k,k}, 0^{k,k+1}, \dots]^{t_1}$, respectively. We denote v_i^t and w_i^t as how many times a transition was rewarded ($\beta_i = 0$) and performed for LA_i up to interaction step t , respectively. Hence, For the state indices k, j , when $k = j + 1$, state transition probabilities are updated as the below:

$$P^{t+1}(S_i^{k,j}) = \frac{v_i^t(S_i^{k,j})}{w_i^t(S_i^{k,j})}, \quad (4)$$

$$P^{t+1}(S_i^{j,k}) = \frac{1 - P^{t+1}(S_i^{k,j})}{2}, \quad (5)$$

$$P^{t+1}(S_i^{k,k}) = \frac{1 - P^{t+1}(S_i^{k,j})}{2}, \quad (6)$$

$$\text{where } P^{t+1}(S_i^{k,j}) + P^{t+1}(S_i^{j,k}) + P^{t+1}(S_i^{k,k}) = 1. \quad (7)$$

Since each LA performs a random walk over its state space through a stochastic state transition, then the optimization problem is solved by the multidimensional joint random walk over the automata network. Furthermore, the individual state transitions are dependent to each others due to the shared knapsack capacity and the inter-connected influence in their environment rewards. Therefore, the probability of a particular automata network state is calculated as the joint probability of the individual automata current states. Hence the joint probability can be calculated as the below, where N is the network size:

$$P^t(S_i, S_j \dots S_N) = P^t(S_i)P^t(S_j, \dots S_N | S_i). \quad (8)$$

LA Environment. To learn incentivization values for the social network's users, all users' associated LA interact with a Knapsack which evaluates how valuable the current LA state (incentive) for the mitigation campaign. The Knapsack evaluation is individual to each user behaviour on the network. Users behaviours are modeled through a multivariate Hawkes process (MHP). Hence, the LA environment has the following main properties.

- **Stochastic:** which is due to the randomness of each HP itself, which generates random counts for each user events (e.g., re/tweets).

- **Non-stationary:** which occurs because of the dependencies between users HP generated events. For instance, when both users i, j have an explicit or implicit dependency, a particular incentivization value $x_i = 0$ might not be optimum for user i but could be optimum when the incentivization value $x_j > 0$. Since the latter could cause user i to be fairly exposed to true content without the need to increase for x_i (incentivize user i).

To reinforce the learning of targeted state values. each individual LA_i will receive a reward signal β_i from its Knapsack environment where $\beta_i \in \{1, 0\}$, indicating a penalty, or reward Knapsack signal, respectively. The final committed state transition for an LA_i is driven by the reward signal β_i . For instance, if LA_i randomly walks towards the right and received a reward, it commits the transition and updates its current state. However, if LA_i receives a penalty, it rolls back the transition and stays at its recent current state before that transition. The state update mechanism also works if LA_i randomly walks to the left direction. These random walks probabilities in both directions are learned according to Equations 4, 5. On the other hand, recurrent state transitions probabilities are updated according to Equation 6 until converging to a state where the probabilities of performing random walks in both directions became almost 0. The detailed information about how the reward signal β_i is calculated for each direction of an LA_i random walk:

$$(\rightarrow) \beta_i(m_i, \Phi) := \begin{cases} 1, & \text{if } m_i > 0 \vee \Phi = 1 \\ 0, & \text{otherwise} \end{cases}, \quad (9)$$

$$(\leftarrow) \beta_i(m_i, \Phi) := \begin{cases} 1, & \text{if } m_i > 0 \\ 0, & \text{otherwise} \end{cases}, \quad (10)$$

$$\text{subject to } m_i = \frac{\Delta \mathcal{F}(x_i)}{\Delta x_i}, \text{ where } \Delta x_i > 0. \quad (11)$$

Where m_i is the slope of a fairness loss function \mathcal{F} for the associated user i and Φ indicates either the Knapsack is currently full ($\Phi = 1$) or not ($\Phi = 0$). The Knapsack initial capacity starts with 0 and increased or decreased according to each individual LA state transition, while the current Knapsack capacity is shared across the LA network. Given that $x_i = S_i : i \leq M$, since the mitigation incentive x_i over time is represented by the current LA state where such an LA has M states. The above definition of the environment reward for the proposed random walk state transitions ensures converging to optimum or sub-optimum mitigation incentive values. Figure 2 shows an example of our proposed LA state transitions mechanism where the optimization environment is non-stationary and stochastic. However, the LA managed to find a sub-optimal state value.

Fairness Loss Function. To achieve fair mitigation, we need to consider each individual user exposures to both misinformation and true content. Each user exposure associated with a content type is calculated as how much impact that content has on the user. Therefore, the ratio between true and misinformation impact for each user is considered. Hence, a more skewed initial distribution of these ratios will acquire

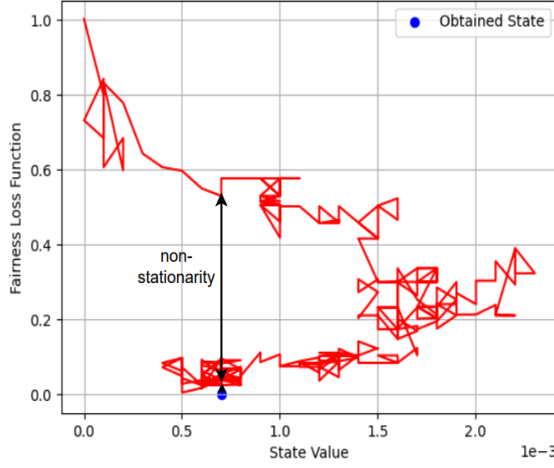


Figure 2: Finding global minima example for an individual LA random walk over a stochastic and non-stationary HP-based Knapsack response.

a fair mitigation strategy to assign the incentivization budget according to user needs, without wasting the budget on users with already high exposures to true content. During the intervention, a ratio $R_i < 1$ means that user i is more exposed to misinformation. Alternatively, a ratio $R_i > 1$ indicates that user i incentivization is not necessary since the latter has already high level of true content exposures. The exposure values used in R_i were calculated as proposed by (Abouzeid et al. 2021), see Appendix A.2 (Abouzeid et al. 2022) for more details. Below, we define our proposed fair misinformation mitigation loss function:

$$\min \mathcal{F}(X) := \sum_i^N \mathcal{F}(x_i), \text{ where } \mathcal{F}(x_i) := \sum_{j=0}^n (1 - R_j^{x_i})^2, \quad (12)$$

$$\text{subject to } \sum_{i=1}^N x_i, \text{ where } x_i \in [0, C]. \quad (13)$$

Where N represents the number of network users and n is the number of adjacent users connected to user i , where user i is considered adjacent to itself. Therefore, j is the index represents i and all its adjacent over the summation. $R_j^{x_i}$ represents the updated ratio between true content and misinformation after applying the incentivization value x_i to the true content HP diffusion model associated with user i . As noticed in Equation 12, we square the subtraction $1 - R_j^{x_i}$ to maintain positive values in the interval $[0, \infty)$, while the task is to minimize the loss function as much closer to 0 as possible (See Figure 2). It is important to highlight that the total loss is calculated through the achieved individual loss of each user during the allocation of incentives (e.g., associated LA and its current state value). That means the total loss ensures optimum or sub-optimum assigned incentivization values over X , where X can be viewed as the set of all automata current states. Eventually, the consumption of all

Algorithm 1: Fair misinformation mitigation.

Input: $\mu_i^{t_0}, \pi_i^{t_0}, R^{x_i=0}, \mathcal{F}_i^{t_0}(x_i = 0), \forall i : u_i \in U, c^{t_0}$, and N where $|U| = N$.

Output: $S_i^*, \forall i : u_i \in U$, where $|S^*| = N$.

```

1: Let  $t = 1$ .
2: while  $\neg(\pi_{all}^t \leftarrow \pi_{all}^*)$  do
3:   for  $i \leftarrow 1$  to  $N$  do
4:     if  $\pi_i^t \neq \pi_i^*$  then
5:        $a_i^t \leftarrow \max[P(S_i^{k,j}), P(S_i^{k,k}), P(S_i^{j,k})]^t$ .
6:        $S_i^t \leftarrow a_i^t$ .
7:        $x_i \leftarrow S_i^t$ .
8:        $\Delta x_i \leftarrow \text{abs}(S_i^t - S_i^{t-1})$ .
9:        $\sum_{j=0}^n (R_j^{x_i}) \leftarrow \lambda(x_i)$ .
10:       $\mathcal{F}^t(x_i) \leftarrow \sum_{j=0}^n (1 - R_j^{x_i})^2$ .
11:       $\Delta F(x_i) \leftarrow \text{abs}(F(x_i)^t - F(x_i)^{t-1})$ .
12:       $m_i = \frac{\Delta F(x_i)}{\Delta x_i}$ .
13:       $\beta_i^t \leftarrow \beta_i^t(m_i, \Phi)$ .
14:       $S_i^{t+1}, \pi_i^{t+1} \leftarrow \delta(S_i^t, a_i^t, \beta_i^t)$ .
15:     else
16:       continue.
17:     end if
18:   end for
19:    $t \leftarrow t + 1$ .
20: end while
21: return  $S^*$ .
```

incentivization values (LA states) must not exceed the bound C , which represents the Knapsack capacity.

Misinformation Mitigation. To obtain the optimum or sub-optimum learned states vectors of N automata, we initialize each individual LA_i with an initial state transition probability vector $\pi_i^{t_0}$, and the initial ratio $R^{x_i=0}$ where no incentivization values yet to be added to the associated estimated base intensity $\mu_i^{t_0}$ of the relevant HP. Eventually, the initial fairness loss function $\mathcal{F}_i^{t_0}(x_i = 0)$ is calculated while the Knapsack is initially empty $c^{t_0} = 0$. The mitigation algorithm then iterates over the whole LA network until it converges to all optimum or sub-optimum state probability vectors. Then, converged states values are suggested as incentivization values for the underlying associated users on the network. The details of the misinformation mitigation procedure is shown in Algorithm 1.

Experimental Setup

In our experiments we design six synthetic social networks $\{syn1, syn2, syn3, \dots, syn6\}$. Each with a unique statistical misinformation exposure distribution among users. The six networks represent the possible real-world scenarios where some user groups might be highly exposed to misinformation more than other groups on the social network. Moreover, some individuals in these groups might be also highly exposed to misinformation more than others from the same group. Allowing for these possible scenarios in our experiments should stress the evaluation of robustness for a fair misinformation mitigation resolution. We design our synthetic networks by randomly generate variant true informa-

Network	Knapsack size	Overall Misinformation
Syn1	0.06	17.00%
Syn2	0.06	58.00%
Syn3	0.06	88.50%
COVID-19	0.06	89.50%
Syn4	0.18	11.75%
Syn5	0.18	47.25%
Syn6	0.18	86.50%
COVID-19	0.18	89.50%

Table 1: Configuration details of fair misinformation mitigation experiments on the proposed social networks.

tion and misinformation event counts on both user and network levels. Then, we set different bounds on these synthetic exposures to maintain a variety of statistics for each network. Eventually, we run our resolution on a real-world social network used in (Abouzeid et al. 2021) as another benchmark. The real-world network is a COVID-19 social network and annotated for ordinary and false re/tweets from Twitter on the 28th of March, 2020. The collected re/tweets focused on discussions about COVID-19. The criteria for the misinformation annotation was if any propagated content urged the public for using false drugs (Tsfaye et al. 2020) without any official statements from the health authorities at that time. Within each of our experiments, we consider different mitigation incentivization budget for the Knapsack capacity to evaluate for different levels of constraints. Due to the randomness of experiments, we run each for multiple times and take the average as an estimate of the final outcome. Table ?? shows the configuration of our experiments, where all networks have 200 users. For the selection of hyper-parameters values in all experiments, see Appendix A.3 (Abouzeid et al. 2022).

Evaluation

Uniform-baseline. To highlight the need for a fair misinformation mitigation method, we make an analogy with a uniform allocation of the incentivization budget. For instance, if all or almost network users are equally exposed to misinformation than true content, a uniform distribution of incentivization budget is theoretically an optimum fair mitigation strategy. We refer to the latter as **Case-0**. However, the more the two content types were unbalanced on the network, the more challenging for a budget uniform distribution to achieve the desired mitigation results. For example if only 20% of network users were exposed to misinformation, a uniform incentivization becomes a waste for 80% of the budget, which might cause no mitigation at all since 20% of the budget becomes insufficient to maintain $R = 1$ for the targeted users. We refer to the latter as **Case-1**. Another form of skewness is when the majority of users are exposed to misinformation but a subset of them are significantly more exposed to misinformation than others, in such scenario, the uniform method will suffer as well, since these subset of users will need more incentivization than others. We refer to the latter as **Case-2**. It is important to highlight that the purpose of the HP information diffusion model is to predict future behaviours. Therefore, the initial distribution of

misinformation exposures before any future intervention is unknown, and a robust incentivization is mandatory to overcome all the potential misinformation percentages.

AVG-LA-baseline. We further investigate how our LA-based resolution performs against current existed LA-based methods (Abouzeid et al. 2021). We refer to the latter as **AVG-LA**, while we refer to our proposed method as **Fair-LA**.

Mitigation Efficiency. To evaluate for robustness on multiple social networks’ scenarios, we introduce a mitigation efficiency metric which is calculated as per the below:

$$1 - \frac{a}{b}. \quad (14)$$

Where a and b are the misinformation percentages after and before mitigation, respectively. According to our synthetic social networks’ different setups (See Table ??), **Case-1** can be observed in syn1 and syn4, while **Case-2** can be observed in syn3, and syn6. As concluded from Figure 3, our proposed **Fair-LA** outperforms both **AVG-LA** and **Uniform** methods in most of the scenarios, especially in **Case-1**. Moreover, when **Case-2** occurs, **Fair-LA** still outperforms other methods when the Knapsack capacity C was larger. From our statistical analysis on the COVID-19 network with 200 users, we observed almost a scenario equivalent to **Case-0**. Therefore, the **Uniform** method performs better than others. However, we can observe how the efficiency gap is reduced between **Fair-LA** and **Uniform** when the Knapsack capacity is more restricted. Eventually, the STD error in the achieved mitigation efficiency percentages for **Fair-LA** is significantly lower than **AVG-LA** which also shows how our proposed method is more stable.

Fairness Error. Since our proposed loss function (See Equation 12) is considered a general fairness concept, we measure how fair the distribution of incentivization budget among all methods by calculating a normalized total loss. Figure 4 shows how our proposed method significantly achieved less fairness error among other methods in all scenarios with stable STD error as well. Consequently, that resulted in not consuming the whole incentivization budget by our method. See Appendix A.4 (Abouzeid et al. 2022) for more details about how **Fair-LA** is wisely consuming the Knapsack capacity.

Learning Bias. In the context of our work, a learning bias means unnecessary incentivization values to be assigned based on incomplete evaluation of users’ needs due to the non-stationary problem. To reduce such bias, we considered a relatively small learning rate (the automaton state increase/decrease value) that ensured all users will be visited almost equal times before consuming the whole budget. Moreover, the fairness error ensured that no user will consume more than its needs from the budget. Eventually, political polarization would reshape how the learned incentives could actually cause mitigation. Hence, modeling the polarized responses to incentives should be integrated with our resolution in the future work.




Figure 3: Mitigation efficiency on different social network scenarios. Left image: $C=0.06$, right image: $C=0.18$.




Figure 4: Normalized fairness error on different social network scenarios. Left image: $C=0.06$, right image: $C=0.18$.

Desired Mitigation Baseline. As demonstrated, the idea of misinformation mitigation is to introduce counter information by incentivizing users to propagate it on the network. However, a question remains about to which extent a mitigation should be considered enough. In other words, what if an equal exposure of counter information to misinformation is not enough to maintain authenticity on the network. In such scenario, we propose a balance factor parameter, where the ratios in Equation 12 are considered fair only when approaching some balance. For instance, if the desired counter information exposure needed to be twice the amount of misinformation exposure per each user, then, the balance factor is set to 2 and the fairness of the ratio R is interpreted accordingly. See Appendix A.2 (Abouzeid et al. 2022).

Computation Speed. Due to the criticality of the misinformation problem, time is an important factor when evaluating misinformation mitigation resolutions. The complete comparison between **AVG-LA** and **Fair-LA** regarding their computation speed is given in Appendix A.5 (Abouzeid et al. 2022).

Large Scaled Networks. Appendix A.6 (Abouzeid et al. 2022) demonstrates how our method could be scaled on larger networks when sampling techniques are adopted to reduce the optimization space without sacrificing the mitigation efficiency.

Conclusion

This paper proposed a socially fair approach to misinformation mitigation on social networks. We introduced different synthetic social networks to generate diversity in scenarios where fairness will be critical to how we consume mitigation resources. Unlike other methods, where the fairness perspective was not considered and therefore the social networks which were evaluated were not diverse enough. However, as a limitation in our work, we did not consider the problem of non-responding users in a detailed manner. For instance, some users might be extremely polarized to respond to our incentivization even if their associated HP was responsive. Therefore, we believe that a model for political polarization can be integrated with our proposed method in the future.

References

- Abouzeid, A.; Granmo, O. C.; Webersik, C.; and Goodwin, M. 2019. Causality-based Social Media Analysis for Normal Users Credibility Assessment in a Political Crisis. In *2019 25th Conference of Open Innovations Association (FRUCT)*, 3–14. IEEE.
- Abouzeid, A.; Granmo, O.-C.; Webersik, C.; and Goodwin, M. 2021. Learning Automata-based Misinformation Mitigation via Hawkes Processes. *Information Systems Frontiers*, 1–20.
- Abouzeid, A.; Granmo, O.-C.; Webersik, C.; and Goodwin, M. 2022. *Socially Fair Mitigation of Misinformation on Social Networks via Constraint Stochastic Optimization*. arXiv.
- Allcott, H.; and Gentzkow, M. 2017. Social media and fake news in the 2016 election. *Journal of economic perspectives*, 31(2): 211–36.
- Ames, C. 1989. The Markov process as a compositional model: A survey and tutorial. *Leonardo*, 22(2): 175–187.
- Bradshaw, S.; and Howard, P. 2017. *Troops, trolls and troublemakers: A global inventory of organized social media manipulation*. Oxford Internet Institute.
- Chen, Y. 2016. Thinning algorithms for simulating point processes. *Florida State University, Tallahassee, FL*.
- De Beer, D.; and Matthee, M. 2020. Approaches to Identify Fake News: A Systematic Literature Review. In *International Conference on Integrated Science*, 13–22. Springer.
- Farajtabar, M.; Yang, J.; Ye, X.; Xu, H.; Trivedi, R.; Khalil, E.; Li, S.; Song, L.; and Zha, H. 2017. Fake news mitigation via point process based intervention. In *International Conference on Machine Learning*, 1097–1106. PMLR.
- Fréville, A. 2004. The multidimensional 0–1 knapsack problem: An overview. *European Journal of Operational Research*, 155(1): 1–21.
- Goindani, M.; and Neville, J. 2020. Social reinforcement learning to combat fake news spread. In *Uncertainty in Artificial Intelligence*, 1006–1016. PMLR.
- Gotzamani, K.; Georgiou, A.; Andronikidis, A.; and Kamvysi, K. 2018. Introducing multivariate Markov modeling within QFD to anticipate future customer preferences in product design. *International Journal of Quality & Reliability Management*.
- Granmo, O.-C.; Oommen, B. J.; Myrer, S. A.; and Olsen, M. G. 2007. Learning automata-based solutions to the non-linear fractional knapsack problem with applications to optimal resource allocation. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 37(1): 166–175.
- Gupta, A.; Farajtabar, M.; Dilkina, B.; and Zha, H. 2018. Discrete Interventions in Hawkes Processes with Applications in Invasive Species Management. In *IJCAI*, 3385–3392.
- Jin, F.; Wang, W.; Zhao, L.; Dougherty, E.; Cao, Y.; Lu, C.-T.; and Ramakrishnan, N. 2014. Misinformation propagation in the age of twitter. *Computer*, 47(12): 90–94.
- Laub, P. J.; Taimre, T.; and Pollett, P. K. 2015. Hawkes processes. *arXiv preprint arXiv:1507.02822*.
- Lazer, D. M.; Baum, M. A.; Benkler, Y.; Berinsky, A. J.; Greenhill, K. M.; Menczer, F.; Metzger, M. J.; Nyhan, B.; Pennycook, G.; Rothschild, D.; et al. 2018. The science of fake news. *Science*, 359(6380): 1094–1096.
- Marquioni, V. M. 2019. Multidimensional elephant random walk with coupled memory. *Physical Review E*, 100(5): 052131.
- Narendra, K. S.; and Thathachar, M. A. 1974. Learning automata-a survey. *IEEE Transactions on systems, man, and cybernetics*, (4): 323–334.
- Nicosia, G.; Pacifici, A.; and Pferschy, U. 2009. On Multi-Agent Knapsack Problems. In *CTW*, 44–47.
- Ogata, Y. 1981. On Lewis’ simulation method for point processes. *IEEE transactions on information theory*, 27(1): 23–31.
- Ousidhoum, N.; Song, Y.; and Yeung, D.-Y. 2020. Comparative evaluation of label agnostic selection bias in multilingual hate speech datasets. In *Proceedings of the 2020 conference on empirical methods in natural language processing (EMNLP)*, 2532–2542.
- Ozaki, T. 1979. Maximum likelihood estimation of Hawkes’ self-exciting point processes. *Annals of the Institute of Statistical Mathematics*, 31(1): 145–155.
- Özdan, S. 2021. The Right to Freedom of Expression Versus Legal Actions Against Fake News: A Case Study of Singapore. In *The Epistemology of Deceit in a Postdigital Era*, 77–94. Springer.
- Pearson, K. 1905. The problem of the random walk. *Nature*, 72(1867): 342–342.
- Pulido, C. M.; Ruiz-Eugenio, L.; Redondo-Sama, G.; and Villarejo-Carballido, B. 2020. A new application of social impact in social media for overcoming fake news in health. *International journal of environmental research and public health*, 17(7): 2430.
- Rampersad, G.; and Althiyabi, T. 2020. Fake news: Acceptance by demographics and culture on social media. *Journal of Information Technology & Politics*, 17(1): 1–11.
- Shao, C.; Ciampaglia, G. L.; Varol, O.; Yang, K.-C.; Flammini, A.; and Menczer, F. 2018. The spread of low-credibility content by social bots. *Nature communications*, 9(1): 1–9.
- Sharma, K.; Qian, F.; Jiang, H.; Ruchansky, N.; Zhang, M.; and Liu, Y. 2019. Combating fake news: A survey on identification and mitigation techniques. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(3): 1–42.
- Shu, K.; Sliva, A.; Wang, S.; Tang, J.; and Liu, H. 2017. Fake news detection on social media: A data mining perspective. *ACM SIGKDD Explorations Newsletter*, 19(1): 22–36.
- Tesfaye, W.; Abrha, S.; Sinnollareddy, M.; Arnold, B.; Brown, A.; Matthew, C.; Oguoma, V. M.; Peterson, G. M.; and Thomas, J. 2020. How do we combat bogus medicines in the age of the COVID-19 pandemic? *The American Journal of Tropical Medicine and Hygiene*, 103(4): 1360–1363.

- Ulker, E.; and Tongur, V. 2017. Migrating birds optimization (MBO) algorithm to solve knapsack problem. *Procedia computer science*, 111: 71–76.
- Wasike, J. 2013. Social media ethical issues: role of a librarian. *Library Hi Tech News*.
- Yazidi, A.; and Hammer, H. L. 2018. Solving stochastic nonlinear resource allocation problems using continuous learning automata. *Applied Intelligence*, 48(11): 4392–4411.
- Zannettou, S.; Caulfield, T.; Setzer, W.; Sirivianos, M.; Stringhini, G.; and Blackburn, J. 2019. Who let the trolls out? towards understanding state-sponsored trolls. In *Proceedings of the 10th acm conference on web science*, 353–362.
- Zhang, Q.; Zhang, S.; Dong, J.; Xiong, J.; and Cheng, X. 2015. Automatic detection of rumor on social network. In *Natural Language Processing and Chinese Computing*, 113–122. Springer.