

Network Inference from a Mixture of Diffusion Models for Fake News Mitigation

Karishma Sharma,¹ Xinran He,² Sungyong Seo,¹ Yan Liu¹

¹ University of Southern California

² Snap Inc.

krsharma@usc.edu, xinranhe@alumni.usc.edu, sungyons@usc.edu, yanliu.cs@usc.edu

Abstract

The dissemination of fake news intended to deceive people, influence public opinion and manipulate social outcomes, has become a pressing problem on social media. Moreover, information sharing on social media facilitates diffusion of viral information cascades. In this work, we focus on understanding and leveraging diffusion dynamics of false and legitimate contents in order to facilitate network interventions for fake news mitigation. We analyze real-world Twitter datasets comprising fake and true news cascades, to understand differences in diffusion dynamics and user behaviours with regards to fake and true contents. Based on the analysis, we model the diffusion as a mixture of Independent Cascade models (MIC) with parameters θ_T, θ_F over the social network graph; and derive unsupervised inference techniques for parameter estimation of the diffusion mixture model from observed, unlabeled cascades. Users influential in the propagation of true and fake contents are identified using the inferred diffusion dynamics. Characteristics of the identified influential users reveal positive correlation between influential users identified for fake news and their relative appearance in fake news cascades. Identified influential users tend to be related to topics of more viral information cascades than less viral ones; and identified fake news influential users have relatively fewer counts of direct followers, compared to the true news influential users. Intervention analysis on nodes and edges demonstrates capacity of the inferred diffusion dynamics in supporting network interventions for mitigation.

Introduction

Falsified information, that is generally intended to deceive people, influence public opinion and manipulate social outcomes has become a prominent topic of discussion. In 2013, the World Economic Forum regarded fake news as a rising global risk in the report entitled ‘Digital Wildfires in a Hyper-connected World’. Even though deception through falsified information has existed in the past, the increasing use and nature of social media, has made the problem much more intense and difficult to combat.

The risks associated with fake news are more significant due to the scale and reach of social media; the last decade itself has seen more than a ten-fold increase in social media usage (Perrin 2015). The major impacts of fake news

have been in social, economic and political issues around the world such as the 2016 US Presidential Elections (Allcott and Gentzkow 2017). Besides that, misleading stories discrediting the severity of climate change (Roozenbeek and van der Linden 2018), and recurrent attempts to promote fear and confusion during natural disasters (Gupta et al. 2013; Takayasu et al. 2015) cannot be neglected.

Fake news mitigation has been largely studied from the perspective of detection using content analysis, social bots analysis, and analysis of user responses/engagements to the content on social media (Sharma et al. 2019). As compared to traditional media, online social media allows decentralized dissemination and sharing of content, that can rapidly result in viral information cascades, and widespread impact of misinformation. Therefore, research in intervention strategies to mitigate fake news by monitoring or limiting such diffusions were developed in (Farajtabar et al. 2017; Goindani and Neville 2020). Farajtabar et al. derived optimal intervention intensities required at nodes in the network to accelerate diffusion of true news through external stimulation. However, facilitating network interventions such as this, requires learning diffusion dynamics of fake and true contents from observed user engagements. Here, we consider the problem of learning diffusion dynamics from observed, but unlabeled cascades of fake and true news; and leveraging the inferred dynamics to facilitate network interventions for fake news mitigation.

Contributions and Outline

In this work, we address the phenomenon of *diffusion* of fake and true contents on social media, using two real-world datasets comprising false and legitimate content cascades collected from user engagements on Twitter. Based on user behaviours in fake and true cascades, we propose a diffusion mixture model (MIC) with parameters θ_T, θ_F as a generative model of the diffusion process; and derive *unsupervised* inference techniques for parameter estimation. Unsupervised estimation is important in this domain, since the cost of acquiring labeled (fake/true) cascades is higher due to reliance on expert verification. Using the inferred parameters, we evaluate the role of different users and the network in the propagation of misinformation, and provide analysis for network interventions. The following is an outline of the contributions:

- We investigate the nature of user behaviours in response to fake and true news on Twitter datasets. Our findings indicate statistical differences in diffusion patterns of fake and true news with non-homogeneous sharing behaviours.
- We propose an unsupervised method to learn the diffusion dynamics from observed, unlabeled information cascades, under the diffusion mixture model MIC, and are the first to examine learnability guarantees in the same.
- We evaluate if fake and true cascades are separable using inferred dynamics, compared to unsupervised clustering methods based on text, user and propagation features.
- We examine characteristics of users identified as influential in spreading legitimate and fake contents using the inferred diffusion dynamics. Inferred influential fake news users have positive correlation with relative appearance in fake cascades; have relatively fewer counts of direct followers compared to influential true news users; and inferred influential users tend to have engagements in topics of more viral/larger cascades, than smaller ones.
- Intervention analysis demonstrates reduction in fake cascade size compared to other unsupervised methods. The learned diffusion dynamics are useful towards actively limiting or mitigating misinformation.

Related Work

Fake news mitigation is largely addressed as a detection (*classification*) task in existing literature. Sharma et al. (2019) classified approaches for fake news detection based on the features used for classification. Broadly, the methods focus on content or writing style analysis (Wang et al. 2014b; Khattar et al. 2019), source or bot analysis (Ferrara et al. 2016), and features from user responses/engagements on social media (Qian et al. 2018; Ma, Gao, and Wong 2017). The features from user responses are found to be informative and complementary to content or source analysis. In this work, our focus is on information diffusion on social media, to understand how fake and true contents are propagated, and learn a generative model of propagation. Farajtabar et al.; Goindani and Neville (2017; 2020) studied intervention strategies based on reinforcement learning for accelerating or limiting diffusions. However, their focus is not on learning diffusion dynamics from observations; and they assume random or known diffusion parameters. Our work on learning diffusion dynamics is therefore complementary to it, and supports different intervention strategies including these.

Network inference refers to the problem of inferring the diffusion process, under a mathematical model of propagation, from observed information cascades. It is studied under different models of propagation (Gomez-Rodriguez, Balduzzi, and Schölkopf 2011; Gomez-Rodriguez, Leskovec, and Krause 2012; Zhou, Zha, and Song 2013). The objective of network inference is to estimate the parameters of a diffusion model from observed information cascades; which might entail inferring the edges of the diffusion network, or both the edges and the strength of influence (or weights) on the edges. For instance, in the Independent Cas-

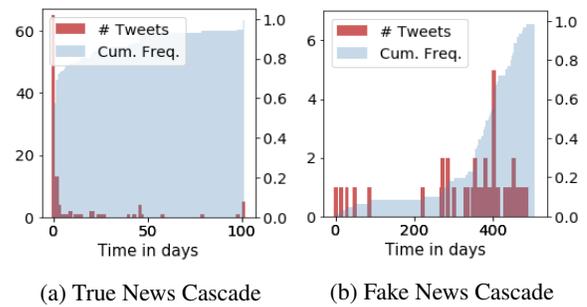


Figure 1: Example of diffusion cascades on Twitter (# tweets per day) for (a) emergency landing of an airliner in Hudson river in 2009 (b) information suggesting that the combination of Coke and Mentos can lead to death in 2006.

cade model (Kempe, Kleinberg, and Tardos 2003), for every pair of users u and v , there is a parameter $p_{u,v}$ which represents the probability with which u activates v , that is information successfully propagates from u to v . In other words it is the strength of influence between u and v . In the multivariate Hawkes process model, parameters $\alpha_{u,v} \geq 0$ model mutually-exciting nature of network activities, with conditional intensity functions capturing the instantaneous rate of future events conditioned on past events.

Most works in network inference do not address heterogeneity in strength of influence between a pair of users. Furthermore, none of them examine whether the influence is heterogeneous with regards to legitimate and fake contents. Earlier works only considered topic or time specific networks (Yang and Zha 2013; Wang et al. 2014a; He and Liu 2017); such as MultiCascades (He and Liu 2017) wherein heterogeneous diffusion models are tied together with joint network priors, and inferred from observed but labeled cascades. Our method considers a heterogeneous diffusion model for true and fake news propagation, and in contrast, we propose an unsupervised method for inference that does not require labeled cascades.

Diffusion Mixture Model

Information propagation or diffusion is widely studied using probabilistic models, in domains related to viral marketing (Domingos and Richardson 2001), and disease and epidemics (Newman 2002). Diffusion models provide a way to solve important computational problems in each domain. For instance, Domingos and Richardson addressed an important question in viral marketing, that is - to trigger a large cascade of product adoptions, who are the most influential users to target in ad campaigns? Such problems can be efficiently solved using submodular optimization under certain diffusion models such as the Independent Cascade model (Kempe, Kleinberg, and Tardos 2003). The choice of model dictates how efficient it is to optimize for important problems such as this. It also affects whether it is possible to derive analytical solutions for learning algorithms in order to infer the parameters of the diffusion model from real observed cascades. Here, we introduce the Independent

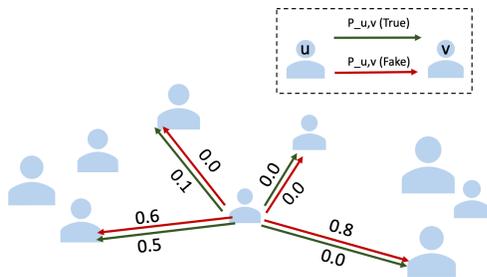


Figure 2: Diffusion mixture model (MIC).

Cascade Model, followed by our extension of the diffusion model to legitimate and fake cascades.

Cascade A cascade is defined as a time-ordered sequence of user responses/ engagements that a piece of information (content) receives, when it is circulated on a social network. It can be labeled as a true or fake news cascade, in accordance with the veracity of the content (e.g. Fig. 1).

Independent Cascade (IC) Model First, we discuss the formulation of the Independent Cascade Model studied in Kempe, Kleinberg, and Tardos. $G = (V, E)$ is the directed graph with $n = |V|$ number of nodes (users) and $m = |E|$ edges. A node is *activated* in an information cascade, if its user has an engagement with the content being propagated. Each edge $(u, v) \in E$ is associated with a parameter $p_{u,v} \in [0, 1]$. The diffusion process starts with an initial set of seed nodes assumed to be activated at the first timestep. At each following time step of the diffusion process, a node u activated at the previous time step t , independently makes a single activation attempt on each inactive neighbor v . The activation succeeds with probability $p_{u,v}$ and a node once activated remains activated in the diffusion process. The influence function σ is a function of the seed set S and $\sigma_\theta(S)$ is defined as the expected number of nodes activated by the end of the diffusion process starting at seeds S , where $\theta = \{p_{u,v} | (u, v) \in E\}$ refers to the parameter set.

Mixture of Independent Cascade (MIC)

Given a social network $G = (V, E)$, we extend the IC model to include the diffusion of both legitimate (true) and misinformation (fake) contents using separate sets of parameters $\theta_T = \{p_{uv}^T | (u, v) \in E\}$ and $\theta_F = \{p_{uv}^F | (u, v) \in E\}$, i.e., both types of contents share the same network skeleton G but with separate parameters for activation probabilities on the edges. Based on this parameterization (illustrated in Fig 2), we study the inference of the proposed diffusion mixture model parameters from observed, unlabeled cascades.

First, we formally define the inference problem for the proposed diffusion mixture model, formulated as a mixture of independent cascade models (MIC). We assume that the observed set of diffusion cascades C contains a mixture of unlabeled true and fake cascades. We study whether the diffusion process of true and fake contents can be learned directly from C , without requiring cascade labels

$\in \{\text{true/fake}\}$. This makes the inference problem more challenging, but practically more useful when collection of labeled cascades requires expert human verification.

Problem Formulation We assume π^T is the probability with which a true news cascade emerges, and $\pi^F = 1 - \pi^T$ is the probability with which a fake news cascade emerges. Let $\pi = [\pi^T, \pi^F]$ be the mixing weights of the diffusion mixture model, then each cascade $c_i \in C$ is assumed to be generated independently under MIC as follows:

1. Generated seed set $S \subseteq V$ is sampled from some unknown distribution \mathbb{P} over V .
2. Generated cascade corresponds to true or fake news based on the outcome of the random variable $h_i \sim \text{Bernoulli}(\pi^T)$; Cascade labels and mixing weights π unobserved.
3. Generated cascade is drawn from the diffusion mixture model with $c_i \sim \text{IC}(\theta_T)$ if $h_i = 1$ and $c_i \sim \text{IC}(\theta_F)$ otherwise; diffusion parameters θ_T, θ_F are unobserved.

The objective of the network inference problem thereby is to infer θ_T and θ_F and π from unlabeled cascades C .

Real Datasets and Diffusion Analysis

In the previous section, we proposed the diffusion mixture model with separate sets of parameters $\theta_T = \{p_{uv}^T | (u, v) \in E\}$ and $\theta_F = \{p_{uv}^F | (u, v) \in E\}$ for legitimate and fake contents. In this section we first answer two important questions

- Are the diffusion patterns of fake cascades significantly different from true cascades?
- Are user behaviours with respect to fake and true contents non-homogeneous?

We first analyze the diffusion patterns and investigate user behaviours in fake and true cascades. Significant differences between fake and true cascades would mean that the diffusion of fake and true contents are non-homogeneous with respect to user behaviours and should be modeled with separate parameters θ_T and θ_F . For the purpose of our analysis, we consider real world Twitter datasets described in the following section, followed by statistical hypothesis testing to analyze their diffusion characteristics. A few earlier studies such as (Kwon et al. 2013; Castillo, Mendoza, and Poblete 2011; Liu et al. 2017) identified which features of a set of hand-crafted features were most discriminative in training classifiers for detecting fake from legitimate contents. There findings indicate that features with high predictive power include - fraction of information flow from low to high-degree nodes which is higher for fake contents, multiple periodic spikes that are particular to fake contents, and greater depth to breadth ratio in the diffusion trees of fake cascades. In our analysis, we consider *temporal* and *structural* differences in diffusion cascades of fake/true news that investigate how user behaviours towards different types of contents differ.

Real-World Datasets

We utilize two publically available Twitter datasets which we refer to as Twitter-1 (Kwon, Cha, and Jung 2016)¹ and

¹<https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi%3A10.7910%2F2FDVN%2F2BFGAVZ>

Dataset	Twitter-1	Twitter-2
# Users	117,824	233,719
# Engagements	192,350	529,391
# Fake Cascades	60	498
# True Cascades	51	494
Avg T length per cascade (hr)	8,177	1,983
Avg T interval per cascade (hr)	80	65
Avg # engagements in cascade	1,733	597

Table 1: Data statistics for Twitter-1 and Twitter-2.

Twitter-2 (Ma et al. 2016)². Twitter-1 was collected during 2006-2009 and Twitter-2 from March-Dec 2015. In both datasets, topics (contents) are identified as false or legitimate from fact-checking websites like Snopes, and corresponding engagements on Twitter are obtained by keyword search related to the content. The dataset statistics are summarized in Table 1. For analysis we retain all users, and for inference, we retain users that have at least five engagements in the cascade set, resulting effectively in 3K and 7K users in the two datasets. The former contains 111 cascades and the later 992 cascades; with Twitter-1 cascades of average time length of 8177 hrs and Twitter-2 of 1983 hrs.

The datasets contain cascades in the form of time-stamped sequences of user engagements, for example, cascade $C_i = [(u_1, t_1), (u_2, t_2), \dots]$ where u_j, t_j corresponds to the engagement of user u_j at time stamp t_j with content corresponding to cascade C_i . For temporal diffusion analysis, we report statistical tests on each dataset based on the observed cascades. For structural diffusion analysis, we consider only Twitter-1, since it additionally provides follower links from which we can construct retweet structure, similar to (Kwon et al. 2013). The follower graph represents whether user A follows user B. Diffusion of a content from B to A can occur if A follows B, and B posts before A in that cascade. Therefore, we can construct the retweet graph of each cascade, from the cascade engagement sequence and the follower graph. In case A has multiple parents, the edge from the latest parent is retained.

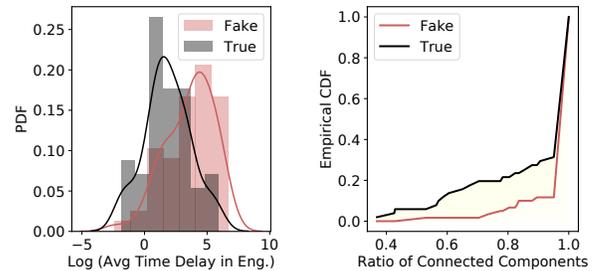
Studying User Behaviours in Fake/True Cascades

In this subsection, we study the diffusion patterns and investigate how these patterns reflect user behaviours. We conduct statistical tests to determine temporal and structural characteristics of fake and true cascades. First, we perform a two-sample t -test to verify whether the average time delay between engagements (posts) is higher in fake cascades v/s in true cascades. The first group of samples S_1 consists of the fake cascades in the datasets. The second group S_2 comprises the true cascades. The log-transform of the data is normally distributed. The null hypothesis is that there is no significant difference between the average time delay between engagements in cascades from the two groups $H_0 : \mu_f = \mu_t$. The alternate hypothesis is the average

²<https://www.dropbox.com/s/46r50ctrfa0ur1o/rumdect.zip?dl=0>

	Temporal		Structural	
	t -statistic	p -value	z -score	p -value
Twitter-1	4.9975	$\leq .00001$	1.87577	.03005
Twitter-2	12.760	$\leq .00001$	NA	NA

Table 2: Hypothesis testing results (p -values) to verify that average time between engagements is higher in fake news cascades than true news cascades (temporal); and to verify that ratio of # of connected components to total engagements is higher in fake news cascades (structural).



(a) Distribution of avg. time delay after log-transform to reduce skewness for t -test. (b) Empirical CDF of the proportion of connected components in fake/true cascades.

Figure 3: Statistical tests distributions (Twitter-1).

time delay between engagements is higher for fake cascades $H_1 : \mu_f > \mu_t$. The p -value is shown in Table 2. The null hypothesis is rejected at significance level $\alpha = 0.01$ which suggests that there is statistically significant difference between the temporal characteristics of fake and true cascades.

Second, we perform statistical significance test to examine differences in structural characteristics of the cascades. We compute the number of connected components (cc) in the retweet graph of each cascade, constructed as mentioned in the previous subsection. Then we define the proportion of connected components in a cascade $r = \frac{\text{number (cc)}}{\text{number of engagements}}$. The null hypothesis is that there is no significant difference in the proportion of connected components in the two groups of fake and true cascades $H_0 : r_f = r_t$. The alternate hypothesis is that it is higher for fake cascades $H_1 : r_f > r_t$. The data is not normally distributed, so we compute the non-parametric Mann-Whitney U test and report the z -score and p -value in Table 2. The null hypothesis is rejected at $\alpha = 0.05$ which suggests that the proportion of connected components in fake cascades is higher than in true cascades. Both statistical tests confirm that diffusion patterns differ based on the type of cascade and the user behaviours towards fake and true contents are non-homogeneous. The distribution of average time between engagements and proportion of connected components is provided in Figure 3 for Twitter-1. The distribution of avg. time between engagements for Twitter-2 cascades is similar to Twitter-1, and structural follower graph is unavailable in Twitter-2; therefore omitted.

Unsupervised Diffusion Network Inference

In this section, we consider a reduction of the problem of learning the diffusion mixture under the MIC model to the problem of learning a mixture of product distributions over a discrete domain; followed by an EM algorithm for parameter estimation of the mixture model.

PAC-Learnability Reduction

Each edge $e = (u, v)$ in the diffusion model is associated with the parameter p_e^T and p_e^F as stated earlier. Each true content cascade can be alternately represented in terms of a ‘live-edge’ graph, such that each edge $e \in E$ is *independently* declared as live with probability p_e^T and included in the graph or blocked with probability $1 - p_e^T$ and not included in the graph. The cascade is then defined by the reachability from seed set S over this graph i.e. a node is activated in the cascade iff there is a directed live edges path from S to the node. Similarly, for fake content cascades. Therefore, each edge can be represented by a random variable x_e^T and x_e^F indicating its live-edge status under the diffusion mixture model, i.e. representing whether the edge e is live or blocked in a given diffusion cascade. Naturally,

$$x_e^T \sim \text{Bernoulli}(p_e^T) \text{ and } x_e^F \sim \text{Bernoulli}(p_e^F)$$

Let X be a vector of random variables indicating live-edge status of each edge under the mixture diffusion model. According to the generative process of the diffusion mixture model, X is a mixture of $k = 2$ components X^T and X^F with mixing weights π . Therefore, X^T is then a discrete distribution over $\{0, 1\}^m$ where m is the number of edges in G and all the x_e^T are independent. Therefore X which is the mixture distribution of X^T and X^F is simply a mixture of discrete product distributions with mixing weights π . The problem is therefore reduced to learning a mixture of discrete product distributions given the live-edge graphs of the observed cascades. Mixture distributions are more generally used in recommendations systems, medicine and other applications (Feldman, O’Donnell, and Servedio 2008) and different algorithms can be used to learn the parameters of the mixture distributions, which in our case are p_e^T, p_e^F for all edges in G and mixing weights π by definition.

Theorem 1. *Given a mixture of unlabeled cascades with completely observed live-edge graphs, with diffusion parameters $\theta_M = \{p_e^M | e \in E\}$ with $M \in \{T, F\}$ and any $\epsilon, \delta > 0$, with mixing weight $\pi^M \geq \frac{\epsilon}{mn}$ we can recover in time $\text{poly}(m^2n/\epsilon) \cdot \log(1/\delta)$, a list of $\text{poly}(m^2n/\epsilon)$ many candidates, at least one of which satisfies the following bound on the influence function $\sigma_{\theta_M}(S)$ and its estimate $\hat{\sigma}_{\theta_M}(S)$ learned from the observed cascades for seed set S drawn from any distribution \mathbb{P} over nodes in G ,*

$$P_{S \sim \mathbb{P}}(|\hat{\sigma}_{\theta_M}(S) - \sigma_{\theta_M}(S)| > \epsilon) \leq \delta$$

with sample complexity $\mathcal{O}\left(\left(\frac{n^4 m^8}{\epsilon^4}\right)^3 \ln \frac{m}{\delta}\right)$ (Proof in Appx).

Parameter Estimation

We can estimate the parameters of the diffusion mixture model θ_T, θ_F and mixing weights π from unlabeled cascades, by deriving a maximum likelihood based estimation

procedure. We assume that the observed cascades record the sequence of user engagements, and the order or timestamps of user engagements (activations) are known.

Notation We use a general notation $M \in \{T, F\}$ to denote a component in the mixture model, wherein $\{T, F\}$ refer to the true and fake components of the $k = 2$ component mixture model MIC. θ_T, θ_F are the set of edge influence parameters for each component IC model in the diffusion mixture model with mixing weights $\pi = [\pi^T, \pi^F]$. That is for graph $G = (V, E)$, $\theta_T = \{p_{uv}^T | u, v \in E\}$ and $\theta_F = \{p_{uv}^F | u, v \in E\}$. We use the notation s to specify an observed sample cascade belonging to the set of cascades C . We define $C_s(t)$ as the set of nodes activated at time step t in cascade s and $t_s(v)$ as the time of activation of node v in cascade s . Also, we define $D_s(t)$ as all activated nodes up to and including time t . Let $p_{u,v}^M$ and $\hat{p}_{u,v}^M$ as the actual and estimated edge activation parameter in component M . In addition, we represent γ_s^M as the posterior probability that cascade s is generated under diffusion component M i.e. $\gamma_s^M = P(Z_s = M; \theta)$ where Z_s indicates the component to which the cascade s belongs and θ is the complete set of parameters θ_T, θ_F , and π . Applying Bayes’ rule,

$$\gamma_s^M = P(Z_s = M; \theta) = \frac{\pi^M P(s; \theta_M)}{\sum_{i=1}^k \pi^i P(s; \theta_i)} \quad (1)$$

Let $Pa(v)$ represent parents of v in G that is, $u \in Pa(v)$ if and only if $(u, v) \in E$. Similarly, $Ch(v)$ is the children of v . Let $p_s^M(v)$ be the probability with which v is activated in cascade s under diffusion component M . By the definition of the IC model, v is activated at time step $t_s(v)$ in cascade s iff at least one activation attempt of an active parent of v in s is successful. Therefore,

$$p_s^M(v) = 1 - \prod_{u \in Pa(v) \cap C_s(t_s(v)-1)} (1 - p_{u,v}^M) \quad (2)$$

In addition, let $A_{u,v} \subseteq C$ be the subset of cascades in which both u and v are activated and $t_s(v) = t_s(u) + 1$ and $B_{u,v}$ be the subset of cascades in which u is activated at some time t and v is not activated up to and including time $t + 1$.

Derivation and Algorithm We derive an expectation maximization based maximum likelihood estimation procedure. The joint log probability of cascade labels and cascades under the mixture model is,

$$\log P(C, Z; \theta) = \sum_{s \in C} \sum_{M \in k} 1_{\{Z_s=M\}} \log(\pi^M P(s; \theta_M))$$

Our goal is to maximize the expected joint log probability,

$$\begin{aligned} Q &= \mathbb{E}[\log P(C, Z; \theta)] \\ &= \sum_{s \in C} \sum_{M \in k} \gamma_s^M \log \pi^M + \sum_{s \in C} \sum_{M \in k} \gamma_s^M \log P(s; \theta_M) \end{aligned}$$

The maximization of Q with respect to π subject to constraints $\sum_M \pi^M = 1, \pi^M \geq 0$, we get, $\pi^M = \frac{1}{|C|} \sum_s \gamma_s^M$ from the first term of Q containing π^M . Now to update the estimates for edge probabilities $p_{u,v}^M$, we need to maximize

the second term of Q by differentiating Q with respect to $p_{u,v}^M$. Let $\hat{p}_{u,v}^M$ be the current estimates of edge influence parameters of edge (u, v) for component M . As stated in the notations, $C_s(t)$ is the set of activated nodes at times step t in cascade s and $t_s(v)$ is the time of activation of node v in cascade s . $p_s^M(v)$ is the probability with which v is activated in cascade s under diffusion model M . The second term of Q involves the product terms of Equation 2 which cannot be solved analytically. However, based on the definition of the IC model, it is possible to approximate Q based on the current estimates of the parameters (Gruhl et al. 2004; Saito, Nakano, and Kimura 2008). We utilize the linear approximation chosen in (Saito, Nakano, and Kimura 2008). Primarily Q can be decomposed in terms of nodes activated in a cascade and nodes not activated in a cascade. For the second case of inactive nodes, we will not need any approximation as the likelihood involves $\log(1 - p_s^M(v))$ which eliminates the product form of Equation 2. For the first case of active nodes, the form is complex because we do not know which active parent was responsible in activating a given node v . This is because, by the definition of IC, activation attempts of all parents of v activated at a given time step are arbitrarily sequenced. Therefore, following (Saito, Nakano, and Kimura 2008), we can instead approximate $p_s^M(v)$ for this case in terms of $\frac{\hat{p}_{u,v}^M}{\hat{p}_s^M(v)}$ for every active parent u since - the probability that v was activated by u should be proportional to the current estimate $\hat{p}_{u,v}^M$ of the strength of influence of u on v . Therefore, the second term of Q is as follows, where $X = C_s(t+1) \cap Ch(u)$ and $X' = Ch(u) \setminus D_s(t+1)$,

$$\sum_{s \in C} \sum_{M \in k} \gamma_s^M \sum_{t=0}^{T-1} \sum_{u \in C_s(t)} \left[\sum_{v \in X} \left[\frac{\hat{p}_{u,v}^M}{\hat{p}_s^M(v)} \log p_{u,v}^M + \left(1 - \frac{\hat{p}_{u,v}^M}{\hat{p}_s^M(v)} \right) \log(1 - p_{u,v}^M) \right] + \sum_{v' \in X'} \log(1 - p_{u,v'}^M) \right]$$

Differentiating the above with respect to $p_{u,v}^M$ and setting it to zero, and considering $Pa(v)$ represents parents of v in base graph G , $A_{u,v}$ is the subset of samples in which both u and v are activated and $t_s(v) = t_s(u) + 1$ and $B_{u,v}$ is the subset of samples in which u is activated at some time t and v is not activated up to and including time $t + 1$ we get,

$$p_{u,v}^M = \frac{1}{\sum_{s \in A_{u,v}} \gamma_s^M + \sum_{s \in B_{u,v}} \gamma_s^M} \sum_{s \in A_{u,v}} \frac{\gamma_s^M \hat{p}_{u,v}^M}{\hat{p}_s^M(v)}$$

This completes the derivation of the EM procedure with iterative updates in E and M-steps shown in Alg 1.

Relaxation Since observed cascades only contain the order of activations or time stamps at which users are activated, rather than discrete timesteps, and the edges in G are unobserved; we relax Equation 2 to deal with continuous time and let $p_s^M(v)$, the probability that v is active in s under component M equal $1 - \prod_{u \in C_s(t_s(v) - W \leq \tau < t_s(v))} (1 - p_{u,v}^M)$ where W is a lookback window and hyperparameter of the algorithm. Thus, any u activated in W before $t_v(s)$ is considered a potential parent and influencer (that can activate) v . W can be set in unit of time or in terms of number of past events.

Algorithm 1 MIC: Diffusion Mixture Parameter Estimation

Input: observed, unlabeled cascades C

Output: estimate $\hat{\theta}_M, \hat{\pi}^M, \gamma_s^M; \forall s \in C$ and $M \in \{T, F\}$

- 1: $\hat{\pi}^T, \hat{\theta}_T, \hat{\theta}_F \leftarrow \text{init} \in [0, 1]; \hat{\pi}^F \leftarrow 1 - \hat{\pi}^T.$
 - 2: **while** not converged **do**
 - 3: // E-Step
 - 4: $\gamma_s^M \leftarrow \frac{\hat{\pi}^M P(s; \hat{\theta}_M)}{\sum_{i=1}^k \hat{\pi}^i P(s; \hat{\theta}_i)}$
 - 5: $p_s^M(v) \leftarrow 1 - \prod_{u \in Pa(v) \cap C_s(t_s(v)-1)} (1 - \hat{p}_{u,v}^M)$
 - 6: // M-step
 - 7: $\hat{\pi}^M \leftarrow \frac{1}{|C|} \sum_{s \in C} \gamma_s^M;$
 - 8: $\hat{p}_{u,v}^M \leftarrow \frac{1}{\sum_{s \in A_{u,v}} \gamma_s^M + \sum_{s \in B_{u,v}} \gamma_s^M} \sum_{s \in A_{u,v}} \frac{\gamma_s^M \hat{p}_{u,v}^M}{\hat{p}_s^M(v)}$
-

Experimental Analysis on Real Datasets

Using the parameter estimation algorithm, we infer diffusion mixture MIC parameters for the real Twitter datasets described earlier. From the inferred parameters, we evaluate if fake and true cascades are separable based on inferred diffusion dynamics, compared to unsupervised baseline methods for clustering cascades. Next, we identify users that are influential in the propagation of true and fake contents, from the inferred parameters and learned diffusion model; and investigate their characteristic features from the data. Lastly, we demonstrate node and edge interventions based on the inferred diffusion dynamics and show reduction in fake cascade size compared to other baselines.

Clustering Cascades

From the inferred parameters, we can determine if an observed cascade is more likely to be considered fake or true based on the posterior probability of the cascade under each component of the mixture MIC. The predicted component for each cascade is thus obtained as $\text{argmax}_M \gamma_s^M$. This will result in two clusters of cascades. Each cluster is assigned fake or true label based on a held out one-fifth set of cascades with known labels. In Table 3, we evaluate if the fake and true cascades in the datasets are separable based on inferred diffusion dynamics, compared to unsupervised baseline methods for clustering cascades. The implemented baselines are as follows - *TruthFinder* (**TF**) (Yin, Han, and Philip 2008) is a credibility propagation algorithm that exploits conflicting sentiments between user comments to the same content. *StanceEval* (**SE**) exploits the average sentiment of users tweet texts in a cascade as a measure of its type; as it is found that fake cascades tend to elicit negative and questioning responses (Qian et al. 2018; Zhao, Resnick, and Mei 2015). *K-Means* (**KM**) clustering based on temporal and propagation features identified in (Ma et al. 2015; Kwon et al. 2013; Castillo, Mendoza, and Poblete 2011) namely, number of posts in a cascade, time length of cascade, average time gap between posts in the cascade, and fraction of most active users in the cascade. *SEIZ* (**SZ**) (Jin et al. 2013) is a rumor model proposed for unsupervised rumor detection. It partitions users as either ‘‘susceptible’’, ‘‘infected’’, ‘‘exposed’’ or ‘‘skeptical’’ with regards to the content

and models state transitions between them. The model is fit to each cascade separately by solving differential equations. They define a ratio based on the learned parameters of the rumor model for each cascade to classify it as fake or true. Lastly, we include **HIC**, where we assume a homogeneous IC model, with a single parameter value (f) shared over all edges for the fake component and another single parameter value (t) for the true component i.e. $p_{uv}^T = t \forall (u, v)$.

Summary In comparison with MIC where differences in inferred dynamics are exploited for separation; TF and SE utilize aggregate sentiments of user responses which are relatively noisy signals of veracity. We find that KM was biased towards producing a single cluster without being able to effectively separate them; HIC does not model heterogeneous influence across user pairs which limits expressibility of the model; and SZ cannot capture common patterns across cascades, as it fits separate parameters per cascade.

Analysis In terms of the distribution π estimated in MIC, we report the Mean Absolute Error (MAE) between the estimated value and the true data distribution. The data is near balanced, and the estimated $\pi = [0.44, 0.56]$ in Twitter-1 is close to the true distribution, with Mean Absolute Error (MAE) of 0.04. Twitter-2 estimated $\pi = [0.47, 0.53]$ with MAE of 0.058. Therefore, MIC outputs balanced clusters of cascade types. KMeans (KM) on the other hand produces unequal sized, biased clusters, resulting in close to random accuracy predicting most cascades to one type, with low f1. The sentiment analysis methods like StanceEval (SE) make mistakes in cases where true content evokes negative sentiments such as “Is horrified to read about the missing Air France plane” and also due to sentiment lexicons that map certain words like “missing” to negative, such as in “Air France jet missing with 228 people over Atlantic after running into thunderstorms”. This results in negatively correlated predictions below 50% depending on the sentiment patterns and content in the data. TruthFinder (TF) also utilizes user sentiments but is more robust as it accounts for conflicting relationships between users. SEIZ (SZ) is the better baseline based on rumor modeling. But it does depend on an estimated threshold for the ratio per cascade used to determine if the cascade is fake. (Jin et al. 2013) use median ratio over the set of cascades as the threshold, and any observed ratio above this threshold is considered as fake. This can be result in lower quality estimates of the threshold in datasets like Twitter-1 with fewer cascades.

We additionally compare the IC model with MIC. IC model does not have the proposed parameterization for different cascade types, and hence cannot be compared in clustering. Therefore, we report the Average Negative Log-likelihood (NLL) or loss per cascade instead, after parameter estimation using IC and MIC in the datasets. Lower NLL indicates better fit to the observed cascades. Average NLL per cascade on a 20% held-out set of cascades in Twitter-1 for MIC is 9.035 (train cascades 6.29), and for IC, it is 3757.84 (train cascades 58.45). Average NLL per cascade in Twitter-2, for MIC is 8.45 (train cascades 8.35), for IC, it is 1414.25 (train cascades 1052.97). Therefore, separate parameters to represent cascade types allows better diffusion modeling.

	Twitter-1		Twitter-2	
	F1-Score	Accuracy	F1-Score	Accuracy
TF	0.576	0.522	0.573	0.536
SE	0.535	0.531	0.388	0.469
KM	0.253	0.522	0.312	0.490
SZ	0.54±0.03	0.52±0.03	0.56±0.03	0.57±0.01
HIC	0.48±0.16	0.55±0.01	0.49±0.12	0.53±0.02
MIC	0.67±0.02	0.61±0.02	0.63±0.01	0.59±0.01

Table 3: Clustering cascade separability results.

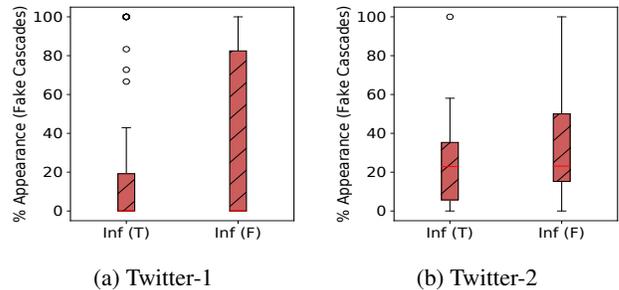


Figure 4: Results on quality of influential users selected based on the estimated diffusion parameters for true and fake news in (a) Twitter-1 and (b) Twitter-2. Inf(T) and Inf(F) are inferred influential users for true and fake news.

Influential Users Test

In this subsection, we first identify users that are influential in the propagation of true and fake contents (i.e. users that if selected as seed sets would trigger the largest cascades), using the inferred diffusion model. Top-100 Influential users for each component IC model with inferred params $\hat{\theta}_T, \hat{\theta}_F$, are selected using greedy maximization algorithm (Goyal, Lu, and Lakshmanan 2011). We use the identified influential users to further examine the quality of inferred parameters.

In Fig 4, we report the box-plot for inferred influential users based on % relative appearance in fake vs true cascades. Inferred users identified for fake news (Inf(F)) have high positive correlation with relative appearance in fake news cascades, as seen from the figure, for both datasets, in comparison with influential true news users (Inf(T)), and a uniform random sample of users (Unif.). Uniform random sample (Unif.) of users, provides insights into the degree of separation between true and fake cascade clusters in the two datasets. Compared to Twitter-1, uniformly sampled users in Twitter-2 are more likely to engage with both contents, whereas in Twitter-1 the median of the uniformly sampled users interact purely with true contents; showing potentially larger separation in Twitter-1 between its users.

Characteristics of Influential Users

Table 4 lists the features of users identified as most influential for true and fake news under estimated parameters; reported # of followers, posts from 2009 Twitter-1 snapshot.

Comp	Followers	Following	# Posts	Name	Description Tags
True News	54418	1157	24182	HuffPost	real life is news, and news is personal. Read more: https:..
	17874	0	675	TMZ	breaking biggest stories in entertainment news
	2684	2941	2144	_PSPGuru	Sending you constant news about the latest PSP news.
	1118	142	3191	FOX10News	TV news station, serving the Alabama, Florida, Gulf
	22252	23853	4621	OnlyMobileNews	We follow the latest in mobile technology news
Fake News	672	280	8200	unk.	F. SwineFlu from pork, SwineFlu zombies F. 45
	514	470	3408	Terrypooch	Fighting for liberty and justice for all F. 308
	1	0	32	08kx250f	F. xbox720 will launch before 2012 F. 9
	3273	1926	2294	unk.	F. Obama is not a natural born citizen F. 5
	8829	8362	6375	BuzzFeed	F. BigFoot, GiantCatfish Maneater, MontaukMonster F. 7

Table 4: Characteristics of most influential users inferred for propagation of fake and true news in Twitter-1.

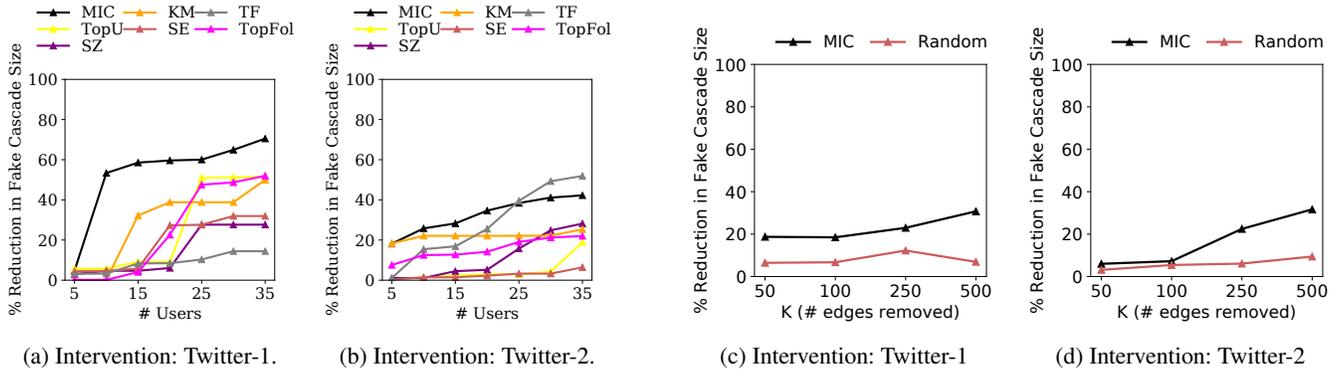


Figure 5: Intervention Analysis on Twitter-1 and Twitter-2 (a, b) Node Interventions (c, d) Edge Interventions.

Features of Inf(T) Inferred influential users identified for true news, as seen, correspond largely to accounts of known credible news and opinion websites and blogs. In terms of topic distribution, the dominant types of influential users include accounts disseminating news related to politics, entertainment, infotainment, technology updates, and local news; and tend to have large number of direct followers.

Features of Inf(F) Top influential users identified for fake news include accounts with relatively fewer counts of direct followers, compared to those for true news users. For some of these the screen name and description is unavailable from Twitter API (reported as ‘unk.’ in the table). Several of these accounts also do not have a listed description along with their screen name, unlike in the previous case of influential true news users. Therefore, we list the topic of the fake news cascades in which the users appear and their total count of engagements/appearances in the fake cascades.

The accounts influential in fake news propagation also appear among a diverse range of topics; similarly dominated by politics, technology, entertainment, and news or trending topics such as SwineFlu and current events. Interestingly, the identified top influential users appear among the larger and more viral fake cascades in the dataset such as ones corresponding to SwineFlu, Obama’s citizenship status, LadyGaga’s gender identity and technology rumors like launch of Xbox720. BuzzFeed interestingly has been his-

torically linked to unreliable journalism, especially before 2014. It appears in connection with false stories related to man-eating catfish, BigFoot and other viral false stories.

Intervention Analysis on Nodes and Edges

In Fig 5 we investigate different intervention mechanisms (mechanisms to monitor or intercept the propagation paths of fake news) leveraging the inferred diffusion dynamics, so as to limit the spread of fake news on a network.

Node Intervention In node intervention, we determine which nodes can be monitored, in order to block false contents from spreading in the network. The inferred influential users for fake news identified earlier are chosen candidates for node intervention under MIC, ranked by influence.

For offline evaluation of the intervention strategy, we utilize the available fake cascades in the datasets. First, we consider that K users are selected for intervention/monitoring. If a fake news cascade reaches any of the monitored users, it can be intercepted and removed from the network, thereby limiting its future spread. The effectiveness of the interception can be evaluated based on the % reduction in fake cascade size due to the intervention. In Fig 5a and 5b, we evaluate the proposed MIC intervention against the previously considered baselines; and we include two additional baselines **TopU** and **TopFol** that intercept users ranked by their total engagement count in the set of observed cascades / to-

tal followers count from the user profile information. For the other baselines, the selection of K users is as follows: rank users by their total engagement count in the cascades predicted as fake news cascades by the baseline method.

Edge Intervention In edge intervention, we select K edges in the network in order to intercept the propagation of fake cascades. The edges are ranked by the weight (strength of influence) p_e^F under the inferred fake component of MIC. These are the identified high transmission paths for fake news cascades and thus removed/blocked.

For offline evaluation, we again compare the percentage reduction in fake cascade size due to edge intervention with MIC, against a Random strategy that selects edges uniformly at random from the network for intervention, as shown in Fig 5c and 5d. Here the reduction is calculated over the size of fake cascades simulated over 1000 rounds under the fake component with and without the K edges removed/intercepted for intervention. The simulations are triggered from seeds sampled from users at the head of the sequence of observed fake cascades in the datasets.

Conclusion

In this work, we proposed a mixture of independent cascade models (MIC) to express and infer the diffusion dynamics of false and legitimate contents. With statistical analysis on real datasets, we confirmed notable differences in user behaviours towards fake and true contents in temporal and structural aspects of diffusion, that can be expressed with MIC. Based on that, we derived an unsupervised inference method for parameter estimation from observed unlabeled cascades, and conducted experiments on Twitter datasets with fake/true news cascades. The experiments revealed interesting analysis of the characteristics of users identified as influential in true and fake content propagation, under the inferred diffusion dynamics; and their effectiveness towards node and edge interventions to limit fake news.

Discussion and Future Work

We assumed two sets of parameters θ_T, θ_F to differentiate fake from true cascades, based on verifying that (i) differences in diffusion patterns of the two types are statistically significant in the datasets, and (ii) the datasets are built from collections of events reported during a specific period, with samples across types collected from the same data source; and no known collection biases across types.

In order to account for multiple types (such as satire, differences in political stance, source credibility or content), the mixture model *easily* generalizes to multiple types of cascades, when $k > 2$ components are initialized in Algorithm 1 (the derivation is written for the general case k).

A limitation of the current work is that it assumes a *fixed* number of components k , which need not be known a priori. In future work, this can be addressed to adaptively split and merge components starting with a large initial k , while optimizing for likelihood of the cascades. In the experiments on influential users identified based on the inferred diffusion parameters, we find that the inferred set of parameters are correlated with the two types assumed in this work, in terms

of engagements with fake and true cascades (Fig 4) and reduction in fake cascade size (Fig 5). However, although the proposed model directly generalizes to $k > 2$, we consider evaluating the model on multiple types with k unknown a priori for future work with multi-label datasets.

The runtime analysis details of the inference algorithm are provided in the *Appendix*. The runtime scales in the order of $\mathcal{O}(k|C|V^2)$ which is reduced to $\mathcal{O}(k|C|VW)$ by setting a constant window W smaller than V , where W is the window size described in Relaxation section under Parameter Estimation, V is the number of users, k is the number of components, and C is the set of cascades. This is a limitation of applying the algorithm to large-scale graphs. In future work, we can integrate dimensionality reduction techniques to reduce the number of unique user representations.

There are other possible directions of future work. The first is to provide online estimation of parameters for time evolving networks; to allow for changing dynamics due to social bots and fake accounts with manufactured and evolving social connections. A second interesting direction is to leverage diffusion network inference to better understand polarization and existence of echo chambers, and its impact on the spread of misinformation - whether polarization fuels misinformation, and can interventions to mitigate one phenomenon support the other phenomenon.

Acknowledgments

The work is supported in part by NSF Research Grant IIS-1619458 and IIS-1254206. The views and conclusions are those of the authors and should not be interpreted as representing the social policies of the funding agency, or the U.S. Government.

Appendix

Proof for Theorem 1

Proof. Each edge (coordinate) j has associated bernoulli variables x_j^i with parameter p_j^i for component i in the k -component mixture distribution. The pairwise coordinate means then are defined as follows,

$$\text{corr}(j, j') = \mathbb{E}[x_j x_{j'}] = \sum_{i=1}^k \pi^i p_j^i p_{j'}^i, \quad 1 \leq j < j' \leq m \quad (3)$$

The sample estimate of $\text{corr}(j, j')$ can be obtained directly from the observed live-edge graphs of unlabeled cascades. By the reduction to learning mixtures of discrete product distributions, given the sample estimates of the pairwise coordinate means, the parameters p_j^i and π can be estimated using algorithm Weights and Means (WAM) (Feldman, O'Donnell, and Servedio 2008) for learning mixture distributions. We restate lemmas in (Feldman, O'Donnell, and Servedio 2008; Chen et al. 2016) used in the proof for completeness, with notations used in the reduction.

Lemma 2 ((Feldman, O'Donnell, and Servedio 2008)). *For $k = \mathcal{O}(1)$ and any $\epsilon', \delta' > 0$, WAM runs in time $\text{poly}(m/\epsilon') \cdot \log(1/\delta')$ and outputs a list of $\text{poly}(m/\epsilon')$*

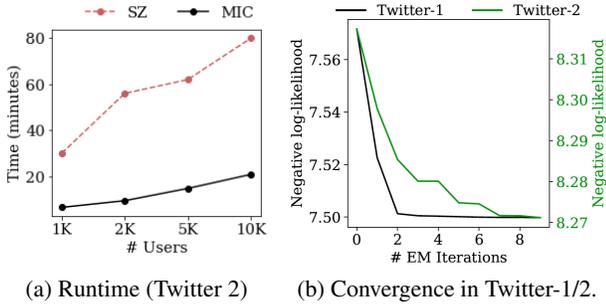


Figure 6: Runtime and Convergence Analysis.

many candidates, at least one of which (with probability at least $1 - \delta'$) satisfies the following,

$$|\hat{\pi}^i - \pi^i| \leq \epsilon', \forall i \quad \text{and} \quad |\hat{p}_e^i - p_e^i| \leq \epsilon', \forall \pi^i \geq \epsilon'$$

Lemma 3 (Lemma 4 in (Chen et al. 2016)). *Given graph G and parameter space ϑ such that $\forall \theta_1, \theta_2 \in \vartheta$, $\|\theta_1 - \theta_2\|_\infty \leq \epsilon_0$, then, $\forall S \subseteq V$,*

$$|\sigma_{\theta_1}(S) - \sigma_{\theta_2}(S)| \leq mn\epsilon_0$$

Using the above lemmas and setting $\epsilon_0 = \frac{\epsilon}{mn}$, $\delta' = \delta$ and $\epsilon' = \frac{\epsilon}{mn}$, the sample complexity for the desired influence function estimate is obtained. WAM requires sample estimates for $\mathbb{E}[x_j x_{j'}]$ for all $1 \leq j < j' \leq m$ to be within an additive accuracy of $\epsilon_{matrix} = \left(\frac{\epsilon'}{m^2}\right)^{(k+1)}$. $x_j x_{j'} \in \{0, 1\}$ and therefore is Bernoulli distributed with some parameter say $p_{jj'}$ equal to $\mathbb{E}[x_j x_{j'}]$. Let $\hat{p}_{jj'}$ be the sample estimate for $\mathbb{E}[x_j x_{j'}]$ calculated from the observed cascades. Since each observed cascade is independently generated, we can compute the sample complexity of estimating $\mathbb{E}[x_j x_{j'}] = p_{jj'}$ within additive accuracy of ϵ_{matrix} given the observed cascades. Applying chernoff bounds, we get $P(|\hat{p}_{jj'} - p_{jj'}| \geq \epsilon_{matrix}) \leq \delta_{matrix}$ with number of observed samples being at least $\frac{2 + \epsilon_{matrix}}{\epsilon_{matrix}^2} \ln \frac{2}{\delta_{matrix}}$. Applying union bound, we get $P(|\hat{p}_{jj'} - p_{jj'}| \geq \epsilon_{matrix}) \leq \delta_{matrix} m(m-1)/2$ for all $j, j' \in [m]$. Setting $\delta_{matrix} = \frac{2\delta}{m(m-1)}$, we get with probability at least $1 - \delta$, $|\hat{p}_{jj'} - p_{jj'}|$ is within additive accuracy of ϵ_{matrix} for all j, j' and the sample complexity is $O\left(\left(\frac{n^4 m^8}{\epsilon^4}\right)^{k+1} \ln \frac{m}{\delta}\right)$. \square

Runtime Analysis

In Fig 6, the runtime analysis of MIC vs. baseline SEIZ (SZ) on Twitter-2 are provided. The baseline SEIZ is run with time interval of 24hours and cut-off time of 10K hours, and it runs differential equation solvers for each cascade, to fit the data with parameters specific to each cascade. The runtimes are evaluated and compared on Intel(R) Xeon(R) CPU E5-2630 v3 @ 2.40GHz on single thread in python. Multi-threading, parallelization is left to future implementations.

Follower Graph	Value
# Edges	27K
# Active users	3K
Avg Out-degree	6.54
Max Out-degree	126
Median Out-degree	3
Avg In-degree	6.55
Max In-degree	137
Median In-degree	2
# strongly connected comps	810
# weakly connected comps	35

Table 5: Follower Graph Statistics in Twitter-1.

The runtime analysis is conducted on Twitter-2, since is the larger of the two datasets (with more users and more cascades), so that runtime can be analyzed with respect to different user sizes. We implemented vectorized computations and pre-computed users and cascades needed in the likelihood computation at the start of the EM iterations which improves computational efficiency, and reduces impact of number of cascades on runtime due to vectorization.

The EM estimation in Algorithm 1 is trained till convergence, i.e. the change in likelihood is smaller than 0.01. The lookback window W , discussed in Section Relaxation under Parameter Estimation, is set to 10 past events. The value of W impacts computational time and should be set to a constant smaller than V , that is the number of users. In the experiments, we set W with line search in the range $\{5, 10, 15\}$ based on cross validation for computational efficiency.

The EM converges within few iterations. The worst-case runtime complexity per EM iteration is $O(k|C|V^2)$ where V is the number of users and C is the set of cascades and k is the number of components, and by setting W to a constant smaller than V , the complexity reduces to $O(k|C|VW)$.

Additional Data Statistics

We provide the follower graph statistics in Table 5 available in Twitter-1 (Kwon et al. 2013). The follower graph was used for structural diffusion analysis; and is a directed graph between the active users considered in the dataset, which appear at least five times in the cascades set; as described in the Section on real-world datasets. The direction of the edge from A to B indicates that A follows B. The table provides degree distribution and connected components statistics.

References

- Allcott, H.; and Gentzkow, M. 2017. Social media and fake news in the 2016 election. *Journal of Economic Perspectives* 31(2): 211–36.
- Castillo, C.; Mendoza, M.; and Poblete, B. 2011. Information credibility on twitter. In *Proceedings of the 20th international conference on World wide web*, 675–684. ACM.
- Chen, W.; Lin, T.; Tan, Z.; Zhao, M.; and Zhou, X. 2016. Robust influence maximization. In *Proceedings of the*

- 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 795–804. ACM.
- Domingos, P.; and Richardson, M. 2001. Mining the network value of customers. In *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*, 57–66. ACM.
- 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 *Proceedings of the 34th International Conference on Machine Learning*, volume 70, 1097–1106. PMLR.
- Feldman, J.; O’Donnell, R.; and Servedio, R. A. 2008. Learning mixtures of product distributions over discrete domains. *SIAM Journal on Computing* 37(5): 1536–1564.
- Ferrara, E.; Varol, O.; Davis, C.; Menczer, F.; and Flammini, A. 2016. The rise of social bots. *Communications of the ACM* 59(7): 96–104.
- Goindani, M.; and Neville, J. 2020. Social reinforcement learning to combat fake news spread. In *Uncertainty in Artificial Intelligence*, 1006–1016. PMLR.
- Gomez-Rodriguez, M.; Balduzzi, D.; and Schölkopf, B. 2011. Uncovering the Temporal Dynamics of Diffusion Networks. In *Proceedings of the 28th International Conference on International Conference on Machine Learning, ICML’11*, 561–568. Madison, WI, USA: Omnipress. ISBN 9781450306195.
- Gomez-Rodriguez, M.; Leskovec, J.; and Krause, A. 2012. Inferring networks of diffusion and influence. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 5(4): 21.
- Goyal, A.; Lu, W.; and Lakshmanan, L. V. 2011. Celf++: optimizing the greedy algorithm for influence maximization in social networks. In *Proceedings of the 20th international conference companion on World wide web*, 47–48. ACM.
- Gruhl, D.; Guha, R.; Liben-Nowell, D.; and Tomkins, A. 2004. Information diffusion through blogspace. In *Proceedings of the 13th international conference on World Wide Web*, 491–501. ACM.
- Gupta, A.; Lamba, H.; Kumaraguru, P.; and Joshi, A. 2013. Faking sandy: characterizing and identifying fake images on twitter during hurricane sandy. In *Proceedings of the 22nd international conference on World Wide Web*, 729–736. ACM.
- He, X.; and Liu, Y. 2017. Not enough data?: Joint inferring multiple diffusion networks via network generation priors. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, 465–474. ACM.
- Jin, F.; Dougherty, E.; Saraf, P.; Cao, Y.; and Ramakrishnan, N. 2013. Epidemiological modeling of news and rumors on twitter. In *Proceedings of the 7th Workshop on Social Network Mining and Analysis*, 8. ACM.
- Kempe, D.; Kleinberg, J.; and Tardos, É. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, 137–146. ACM.
- Khattar, D.; Goud, J. S.; Gupta, M.; and Varma, V. 2019. MVAE: Multimodal Variational Autoencoder for Fake News Detection. In *The World Wide Web Conference*, 2915–2921. ACM.
- Kwon, S.; Cha, M.; and Jung, K. 2016. Rumor detection over varying time windows. doi:10.7910/DVN/BFGAVZ. URL <https://doi.org/10.7910/DVN/BFGAVZ>.
- Kwon, S.; Cha, M.; Jung, K.; Chen, W.; and Wang, Y. 2013. Prominent features of rumor propagation in online social media. In *2013 IEEE 13th International Conference on Data Mining*, 1103–1108. IEEE.
- Liu, Y.; Jin, X.; Shen, H.; and Cheng, X. 2017. Do rumors diffuse differently from non-rumors? A systematically empirical analysis in Sina Weibo for rumor identification. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, 407–420. Springer.
- Ma, J.; Gao, W.; Mitra, P.; Kwon, S.; Jansen, B. J.; Wong, K.-F.; and Cha, M. 2016. Detecting Rumors from Microblogs with Recurrent Neural Networks. In *IJCAI*, 3818–3824.
- Ma, J.; Gao, W.; Wei, Z.; Lu, Y.; and Wong, K.-F. 2015. Detect rumors using time series of social context information on microblogging websites. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*, 1751–1754. ACM.
- Ma, J.; Gao, W.; and Wong, K.-F. 2017. Detect Rumors in Microblog Posts Using Propagation Structure via Kernel Learning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, 708–717.
- Newman, M. E. 2002. Spread of epidemic disease on networks. *Physical review E* 66(1): 016128.
- Perrin, A. 2015. Social media usage. <https://www.pewinternet.org/2015/10/08/social-networking-usage-2005-2015>. Accessed: 2021-04-11.
- Qian, F.; Gong, C.; Sharma, K.; and Liu, Y. 2018. Neural User Response Generator: Fake News Detection with Collective User Intelligence. In *IJCAI*, volume 3834, 3840.
- Roozenbeek, J.; and van der Linden, S. 2018. The fake news game: actively inoculating against the risk of misinformation. *Journal of Risk Research* 0(0): 1–11. doi: 10.1080/13669877.2018.1443491. URL <https://doi.org/10.1080/13669877.2018.1443491>.
- Saito, K.; Nakano, R.; and Kimura, M. 2008. Prediction of information diffusion probabilities for independent cascade model. In *International Conference on Knowledge-Based and Intelligent Information and Engineering Systems*, 67–75. Springer.
- 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*.
- Takayasu, M.; Sato, K.; Sano, Y.; Yamada, K.; Miura, W.; and Takayasu, H. 2015. Rumor diffusion and convergence

during the 3.11 earthquake: a Twitter case study. *PLoS one* 10(4): e0121443.

Wang, S.; Hu, X.; Yu, P. S.; and Li, Z. 2014a. MMRate: inferring multi-aspect diffusion networks with multi-pattern cascades. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, 1246–1255. ACM.

Wang, Z.; Dong, W.; Zhang, W.; and Tan, C. W. 2014b. Rumor source detection with multiple observations: Fundamental limits and algorithms. In *ACM SIGMETRICS Performance Evaluation Review*, volume 42, 1–13. ACM.

Yang, S.-H.; and Zha, H. 2013. Mixture of mutually exciting processes for viral diffusion. In *International Conference on Machine Learning*, 1–9.

Yin, X.; Han, J.; and Philip, S. Y. 2008. Truth discovery with multiple conflicting information providers on the web. *IEEE Transactions on Knowledge and Data Engineering* 20(6): 796–808.

Zhao, Z.; Resnick, P.; and Mei, Q. 2015. Enquiring minds: Early detection of rumors in social media from enquiry posts. In *Proceedings of the 24th International Conference on World Wide Web*, 1395–1405. International World Wide Web Conferences Steering Committee.

Zhou, K.; Zha, H.; and Song, L. 2013. Learning social infectivity in sparse low-rank networks using multi-dimensional Hawkes processes. In *Artificial Intelligence and Statistics*, 641–649.