Identifying Coordinated Accounts on Social Media through Hidden Influence and Group Behaviours

Karishma Sharma*

University of Southern California, USA
krsharma@usc.edu

Emilio Ferrara

University of Southern California, USA
emiliofe@usc.edu

Yizhou Zhang*

University of Southern California, USA
zhangyiz@usc.edu

Yan Liu

University of Southern California, USA
yanliu.cs@usc.edu

ABSTRACT

Disinformation campaigns on social media, involving coordinated activities from malicious accounts towards manipulating public opinion, have become increasingly prevalent. Existing approaches to detect coordinated accounts either make very strict assumptions about coordinated behaviours, or require part of the malicious accounts in the coordinated group to be revealed in order to detect the rest. To address these drawbacks, we propose a generative model, AMDN-HAGE (Attentive Mixture Density Network with Hidden Account Group Estimation) which jointly models account activities and hidden group behaviours based on Temporal Point Processes (TPP) and Gaussian Mixture Model (GMM), to capture inherent characteristics of coordination which is, accounts that coordinate must strongly influence each other's activities, and collectively appear anomalous from normal accounts. To address the challenges of optimizing the proposed model, we provide a bilevel optimization algorithm with theoretical guarantee on convergence. We verified the effectiveness of the proposed method and training algorithm on real-world social network data collected from Twitter related to coordinated campaigns from Russia's Internet Research Agency targeting the 2016 U.S. Presidential Elections, and to identify coordinated campaigns related to the COVID-19 pandemic. Leveraging the learned model, we find that the average influence between coordinated account pairs is the highest. On COVID-19, we found coordinated group spreading anti-vaccination, anti-masks conspiracies that suggest the pandemic is a hoax and political scam.

CCS CONCEPTS

• **Computing methodologies** → **Machine learning**.

KEYWORDS

Coordinated Influence Campaigns, Disinformation, Social Media, Fake News, Temporal Point Process

*These authors contributed equally. They are also the corresponding authors.

This work is licensed under a Creative Commons Attribution International 4.0 License.

KDD '21, August 14–18, 2021, Virtual Event, Singapore
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8332-5/21/08.
<https://doi.org/10.1145/3447548.3467391>

ACM Reference Format:

Karishma Sharma, Yizhou Zhang, Emilio Ferrara, and Yan Liu. 2021. Identifying Coordinated Accounts on Social Media through Hidden Influence and Group Behaviours. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '21), August 14–18, 2021, Virtual Event, Singapore*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3447548.3467391>

1 INTRODUCTION

In recent times, the persistent abuse of social media for spreading disinformation and influencing public opinion and social outcomes has become an increasingly pressing problem [20]. It has largely been used as a tactic to influence elections [2], public perception on social issues such as social distancing policies related to COVID-19 [21] and other local and global events. The issue has gained even more relevance during the ongoing COVID-19 pandemic, where increased reliance on social media for information related to healthcare and policies, has made it an easy target for large-scale disinformation campaigns [21]. The earliest reported cases of disinformation campaigns surfaced when investigations by the U.S. congress found coordinated accounts operated by Russia's "Internet Research Agency" (IRA) or Russian "troll farm" on social media to influence the 2016 U.S. Election by promotion of disinformation, and politically divisive narratives [2]. Identification of coordinated accounts used to manipulate social media is critical.

To address this crucial task, earlier approaches have tried to uncover coordinated accounts based on their *individual* behaviours, from account features and participation in disinformation promotion [1], or collective *group* behaviours, such as activities synchronized in time [4, 17]. However, they all face inherent limitations. The methods relying on disinformation cues or automated behaviours to detect coordinated actors fall short in detecting human operated accounts that instead use persuasion, defamation, and polarization to manipulate opinions (characteristics noted in coordinated influence operations [13]).

Existing methods exploiting collective behaviours heavily rely on assumed features of coordinated behaviour e.g. similar sequence of hashtags as coordination signature or activities synchronized in time [17, 24, 28]. However such features can be inconsistent [30], limiting generalization to unseen accounts. Moreover, reliance on hand-crafted coordination signatures can only capture a limited range of behaviours, and are ineffective at reducing false positives, making strict assumptions on coordinated groups which need not hold true. Fig. 1 shows an example of accounts officially suspended

Account 1	Account 2
2020-03-06 19:22:40 RT @Jali_Cat: 🔵#COVID—19: a deep state attempt?! ...	2020-03-07 01:58:57 RT @Jali_Cat: 🔵#COVID—19: a deep state attempt?! ...
2020-07-04 03:55:37 '!!Where are the law agencies and @realDonaldTrump administration regarding @CDCgov @CDCDirector DELIBERATELY exaggerating & running up COVID19 #s!?! WHERE!?! ...	2020-07-08 04:05:33 '!!Where are the law agencies and @realDonaldTrump administration regarding @CDCgov @CDCDirector DELIBERATELY exaggerating & running up COVID19 #s!?! WHERE!?! ...
2020-07-09 15:28:40 RT @Jali_Cat: Where is @TheJusticeDept?? Where is @HHSGov?? Where is @DHS_Wolf??	2020-07-14 18:20:07 RT @Jali_Cat: Where is @TheJusticeDept?? Where is @HHSGov?? Where is @DHS_Wolf??

Figure 1: Coordinated accounts suspended by Twitter in COVID-19 data (also identified with proposed method). Example tweets from these accounts show political conspiracies, promoted in coordination. The time difference of their coordinated activity varies from less than 6 hours in one instance to more than half a week in the other.

by Twitter and their activities observed in COVID-19 data (also identified as coordinated with our method). The time differences of coordinated activities observed from the accounts is diverse, with less than 6 hours in one case to more than half a week in another.

In this work, to address these shortcomings we propose to model the following inherent characteristics of coordination:

- Strong hidden influence. If accounts coordinate to amplify social media posts or target specific individuals, there should be a strong hidden (latent) influence between their activities. Significant recent computational propaganda is produced and operated by political operatives and governments [25].
- Highly concerted activities. The collective behaviours of coordinated accounts should be collectively anomalous, from other normal accounts on the network with less organized activity patterns (i.e., observations that deviate from normal when generated by a different mechanism [10]).

To capture them, we propose **AMDN-HAGE (Attentive Mixture Density Network with Hidden Account Group Estimation)**, an unsupervised generative model for identification of coordinated accounts, which jointly models account activities and hidden account groups based on Neural Temporal Point Processes (NTPP) and Gaussian Mixture Model (GMM). To learn the latent interactions or influence between account activities, we model the distribution of future activities conditioned on past activities of all accounts with temporal differences, and jointly capture collective anomalous behaviour by simultaneously learning the group membership of accounts. To address the joint optimization, we propose a bilevel optimization algorithm using both stochastic gradient descent and expectation-maximization, from observed activity traces.

As AMDN-HAGE directly learns account representations and hidden groups from activity traces with only account ids and activity timestamps in an unsupervised manner, it does not require

knowledge of partially uncovered accounts in coordinated groups, or predefined individual features, although they can also be plugged in easily if necessary. Furthermore, unlike the existing models relying on strict group behaviour hypotheses like timestamp synchronization, our model only assumes that coordinated accounts are an anomalous group with concerted activities, allowing more data-driven identification of coordinated groups. In addition to above advantages on effectiveness, by incorporating an explicit attention module, we enable our model to learn and output the strength of latent interactions between accounts on the network. In general, our contributions are the following.

- We propose AMDN-HAGE to detect coordinated campaigns from collective group behaviours inferred from account activities based on NTPP and GMM.
- We provide a bilevel optimization algorithm for joint learning of activity trace modeling (NTPP) and social group modeling (GMM) with theoretical and empirical guarantee.
- Extensive experiments on real-world Twitter data about known coordinated campaign verify our method is highly effective on coordination detection. From learned model we find identifiable patterns of coordinated accounts such as the influence between them is highest but also decay fastest.
- We apply the method to identify unknown coordinated campaigns in COVID-19 data, and find coordinated group of Spanish and English accounts (NoMask, NoVaccine, NoALaVaccuna, NoAlNuevoOrdenMundia, QAnon) about no masks, no vaccine, no new world order conspiracies, opposing Bill Gates, that suggest COVID-19 is a hoax and political scam.

2 RELATED WORK

The problem of disinformation and social media abuse has reduced trust in online platforms. Efforts to combat false and misleading information have been widely studied in many different contexts [20] ranging from detection of false information from content features and the responses to it on social media, to understanding the diffusion patterns [19] and accounts involved in its spread [7]. Different from social bots and individual malicious accounts, a growing area of research is the detection of coordinated accounts (coordinated disinformation or influence campaigns), also called “troll farms”, orchestrated by group of human and/or bot accounts that are made to work jointly to promote disinformation or other narratives [12]. The main techniques to identify coordinated accounts are:

Individual behaviours. Existing works mainly uses two kinds of individual behaviours. The first one is the participation in disinformation spreading [18, 19]. For instance, Ruchansky et al. propose a fake news detection model that assigns a suspiciousness score to accounts based on their participation in fake news cascades [18]. The second kind are individual characteristics such as deceptive linguistic features, number of shared links, hashtags and device of posting and cross-platform activity [1, 11, 29]. Apart from above pre-defined features, the activity traces of troll accounts have been found useful for understanding malicious behaviours. In recent work, the tweet, retweet and reply patterns of Twitter accounts are utilized to infer the incentives or rewards behind their activities, formulated as an inverse reinforcement learning problem [12]. Based on the estimated rewards, the authors found that the behaviours of

trolls was different from regular users as they appeared to perform their activity regardless of the responses.

Collective behaviours. Approaches that examine the collective or group behaviours as a whole to detect anomalous malicious accounts are related to our approach. Cao et al. and Gupta et al. cluster accounts that take similar actions around the same time, based on the assumption that malicious account activities are synchronized in time [4, 9]. Other works cluster or partition a account similarity graph defined over hand-crafted features assumed to be indicative of coordinated behaviours, including the sequence of hashtags or articles shared collectively by a large group of accounts [17, 24]. The significant limitation of such approaches is that the assumption on synchronization or hand-crafted features used to define coordination might not hold. In contrast, we propose to automatically learn and detect coordinated behaviours from observed account activities by learning the latent account interactions.

3 TASK DEFINITION AND PRELIMINARIES

3.1 Task Definition

In this section, we introduce the task of detecting coordinated accounts in social networks from collective group behaviours of the accounts. Coordinated campaigns are orchestrated efforts where accounts collude to inorganically spread and amplify the spread of specific narratives for opinion manipulation, and the task we address is to identify such coordinated accounts.

In this work, we propose AMDN-HAGE (Attentive Mixture Density Network with Hidden Account Group Estimation), an unsupervised generative model for coordination detection. It jointly models account **activity traces** and latent **account groups**, to learn collective group behaviours from observed account activities, and detect coordinated accounts with collective anomalous behaviours.

Activity traces: The only input we consider are the activity traces of accounts on the social network. An activity trace can be represented as a sequence of events ordered in time, which can be formulated as $C_s = [(u_1, t_1), (u_2, t_2), (u_3, t_3), \dots, (u_n, t_n)]$. Each tuple (u_i, t_i) corresponds to an activity by account u_i at time t_i . The activities represent account actions on the network such as posting original content, re-sharing, replying, or reacting to other posts.

In order to provide platform/language independent detection, we do not include the type of action, or features such as content of the post, and account metadata, although additional available features can be easily incorporated in the method. The basic input is the most easily available for any social network. Furthermore, the only assumption we make on the coordinated campaign is that:

- Compared to normal accounts, the number of coordinated accounts is quite small (i.e., collectively anomalous).
- Coordinated users have highly concerted activity patterns

Hidden Account Group: In real social networks, accounts with similar activities form social groups, that can constitute normal communities as well as coordinated groups. Supposing that there are N groups in the account set U , we can define a membership function $M : U \rightarrow \{1, \dots, N\}$, which projects each account u_i to its group index. This membership in many cases is hidden or unknown [3]. Acquiring M can help us identify collective anomalous group behaviours to detect coordinated groups. In this work, we aim to learn the hidden groups from only the observed activity traces.

3.2 Temporal Point Process

A temporal point process (TPP) is a stochastic process whose realization is a sequence of discrete events in continuous time $t \in \mathbb{R}^+$ [5]. The history of events in the sequence up to time t are generally denoted as $H_t = \{(u_i, t_i) | t_i < t, u_i \in \mathcal{U}\}$ where $\mathcal{U}$ represents the set of event types (here, accounts). The conditional intensity function $\lambda(t|H_t)$ of a point process is defined as the instantaneous rate of an event in an infinitesimal window at time t given the history i.e. $\lambda(t|H_t)dt = \mathbb{E}[dN(t)|H_t]$ where $N(t)$ is the number of events up to time t . The conditional density function of the i^{th} event can be derived from the conditional intensity [6] as

$$p(t|H_t) = \lambda(t|H_t) \exp\left(-\int_{t_{i-1}}^t \lambda(s|H_t)ds\right) \quad (1)$$

In social network data, the widely used formulation of the conditional intensity is the multivariate Hawkes Process (HP) [32], defined as $\lambda_i(t|H_t) = \mu_i + \sum_{t_j < t} \alpha_{i,j} \kappa(t - t_j)$, where $\lambda_i(t|H_t)$ is the conditional intensity of event type i at time t with base intensity $\mu_i > 0$ and mutually triggering intensity $\alpha_{i,j} > 0$ capturing the **influence** of event type j on i and κ is a decay kernel to model *influence decay* over time. μ and α are learnable parameters. Since HP's fixed formulation and few learnable parameters limit its expressive power, recent works propose to model the intensity function with neural networks [6, 14, 16, 22, 31, 33].

4 COORDINATION DETECTION METHOD

In order to capture the latent influence between account's activities, and collective behaviours of coordinated groups, as well as diversity in coordinated activities from such accounts, we introduce the proposed model AMDN-HAGE.

AMDN-HAGE consists of two components: an Attentive Mixture Density Network (AMDN) that models observed activity traces as a temporal point process and a Hidden Account Group Estimation (HAGE) component that models account groups as mixture of multiple distributions. An overview is shown in Figure 2. The two components share the account embedding layer and reflect the complete generative process that the accounts are first drawn from multiple hidden groups and then interact with each other so that activity traces are observed. Using the observed activity traces, we can learn the generative model by maximizing the likelihood function of the joint model, and acquire not only account embeddings but also a activity trace model and group membership function. Denoting the account embeddings as E , the parameters in AMDN as θ_a and the parameters in HAGE as θ_g , the joint likelihood function can be written as:

$$\begin{aligned} \log p(C_s, U; \theta_g, \theta_a, E) &= \log p(C_s|U; \theta_g, \theta_a, E) + \log p(U; \theta_g, \theta, E) \\ &= \log p(C_s|U; \theta_a, E) + \log p(U; \theta_g, E) \end{aligned} \quad (2)$$

$p(C_s|U; \theta_a, E)$ is the probability density that the activity traces are observed given a known account set, and $p(U; \theta_g, E)$ is the probability density that we observe the account set drawn from the latent hidden social groups. With the account embeddings and the learned membership, we obtain collectively anomalous groups as latent groups with anomalous distributions having small variance or size compared to the rest of the accounts, to detect coordination.

Figure 2: Architecture of proposed (AMDN-HAGE) to model conditional density function of account activities and hidden groups on social media.

Table 1: Summary of neural point process models

Model	Flexible intensity function	Closed-form likelihood	Long-range dependencies	Interpretable Influence
HP [32]	no	y	y	y
RMTPP [6]	limited	y	limited	no
FullyNN [16]	y	y	limited	no
LogNormMix [22]	y	y	limited	no
SAHP [31]	limited	no	y	y
THP [33]	y	no	y	y
AMDN	y	y	y	y

4.1 Modeling Account Activities

In this section, we introduce the AMDN component (Attentive Mixture Density Network) to model account activities for coordination detection. AMDN consists of two parts: a history encoder and an event decoder. Suppose we are modeling the activity (u_i, t_i) , the history encoder represents all the activities that happened before t_i as a vector H_{t_i} . Then the event decoder, which is the conditional density function, predicts (u_i, t_i) based on the history representation H_{t_i} and the known account set. This encoder-decoder architecture models activities with likelihood $p(C_s|U; \theta_a, E)$ factorized as:

$$\log p(C_s|U; \theta_a, E) = \sum_{i=1}^L [\log p_{\theta_a, E}(t_i|H_{t_i}) + \log p_{\theta_a, E}(u_i|H_{t_i})] \quad (3)$$

We provide architecture details in the following paragraphs.

4.1.1 AMDN architecture and training. In Table 1, we summarize existing point process models (detailed in Appx.); these models suffer from different drawbacks with a trade-off on flexible intensity function (better expressive power), closed-form likelihood (reducing gradient noise in training) and interpretable influence (explicit influence score on event pairs). For modeling coordinated accounts, above properties are all useful. Thus, we develop AMDN, which has

all the above properties. We use masked self-attention [23] (with additional temporal encoding for handling irregular inter-event times) to encode the event history for interpretable influence between past and future events (alternative to the recurrent neural network used in [22]), but still use a log-normal mixture distribution as event decoder to model the conditional density of the next event given the history (similar to [22]), achieving all properties.

Encoding event sequences with masked self-attention. Let $\tau \in \mathbb{R}^+$ represent inter-event time, $p(\tau|H_\tau)$ the conditional density. History H_τ is encoded with a neural network to automatically extract useful features, similar to other neural point process models.

Masked self-attention with position encoding. For interpretable influence of past event on future events, we encode the event sequence with masked self-attention [23].

$$A = \sigma(QK^T / \sqrt{d}) \quad \text{and} \quad H_{attn} = AV \quad (4)$$

$$Q = XW_q, \quad K = XW_k, \quad V = XW_v$$

with masked attention weights A of pairwise influence from previous events, input sequence representation $X \in \mathbb{R}^{L \times d}$ (L sequence length, d feature dimension), and learnable weights W_q, W_k, W_v . At the end, we apply layer normalization, dropout and feed-forward layer to H_{attn} to get output $H_{out} \in \mathbb{R}^{L \times d}$. To maintain ordering of the history events, we represent the position information of the i -th event as an m -dim position encoding $PE_{pos=i}$ with trigonometric integral function following [23].

Account and time encoding. We also represent account type and time (u_i, t_i) using learnable account embedding matrix $E \in \mathbb{R}^{|U| \times m}$ and translation-invariant temporal kernel functions [26], using feature maps ϕ with multiple periodic functions of different frequencies ω to embed inter-event times.

$$\phi_\omega(t) = [\sqrt{c_1}, \dots, \sqrt{c_{2j}} \cos(j\pi t/\omega), \sqrt{c_{2j+1}} \sin(j\pi t/\omega) \dots] \quad (5)$$

$$\phi(t) = [\phi_{\omega_1}(t), \phi_{\omega_2}(t) \dots \phi_{\omega_k}(t)]^T \quad (6)$$

The input X_i (to the attention head) of the i^{th} event (u_i, t_i) , is a concatenation of event, position and temporal embedding,

$$X_i = [Eu_i, PE_{pos=i}, \phi(t_i - t_{i-1})] \quad (7)$$

Event history context vector. The attention mechanism gives us representations of each event using attention over events prior to it. We can use the representation of the last event or a recurrent network layer over the attention outputs $H_{out} \in \mathbb{R}^{L \times d}$ to summarize events histories into context vectors $C \in \mathbb{R}^{L \times d}$ where L is event sequence length. Each $c_i \in C$ is a context vector encoding history of events up to t_i i.e. history H_{t_i} of the temporal point process.

Conditional probability density function. With the encoded event history (context vector), the event decoder (learnable conditional density function $p(\tau|H_\tau)$) is used to generate the distribution of the next event time conditioned on the history.

While we can choose any functional form for $p(\tau|H_\tau)$, the only condition is that it should be a valid PDF (non-negative, and integrate to 1 over $\tau \in \mathbb{R}_+$). To maintain a valid PDF, exponential or other distributions with learnable parameters are generally used in point process models [6, 22, 32]. We define the PDF as mixture

of log-normal distributions since the domain of $\tau \in \mathbb{R}_+$ is non-negative (as in [22]), and mixture distributions can approximate any density on $\mathbb{R}$ arbitrarily well [22] not restricted to exponential or other monotonic functions. The conditional PDF is defined as,

$$p(\tau_i | w_i, \mu_i, s_i) = \sum_{k=1}^K w_i^k \frac{1}{\tau s_i^k \sqrt{2\pi}} \exp\left(-\frac{(\log \tau_i - \mu_i^k)^2}{2(s_i^k)^2}\right) \quad (8)$$

$$w_i = \sigma(V_w c_i + b_w), \quad s_i = \exp(V_s c_i + b_s), \quad \mu_i = V_\mu c_i + b_\mu \quad (9)$$

where the mixture weights w_i , means μ_i and stddevs s_i are parameterized by extracted context history c_i and learnable V, b .

The encoder-decoder parameters (denoted jointly as θ_a) and the learnable account embeddings (E) can be learned using maximum likelihood estimation (for log-likelihood defined in Eq. 3) from observed activity sequences and account set, and trained with gradient back-propagation as $\theta_a^*, E^* = \arg\max_{\theta_a, E} \log p(C_s | U; \theta_a, E)$.

4.2 Modeling Hidden Groups

For modeling the hidden social groups from the observed activity traces, we model the hidden groups of accounts as a mixture of N Gaussian multivariate distributions (GMM) in the account embedding space.

Formally, the i -th social group is modeled as a Gaussian distribution $\mathcal{N}(\mu_i, \Sigma_i)$ where μ_i is the cluster mean and Σ_i is the covariance matrix. Since the group membership of accounts is unknown, we assume that the account embeddings are drawn from the mixture, as embedding E_{u_j} of account u_j is distributed as $\sum_i p(i) \mathcal{N}(E_{u_j}; \mu_i, \Sigma_i)$, where $p(i)$ represents prior probability of group i .

Hidden group estimation: A notable difference from general Gaussian mixture models, is that we define the GMM over the learnable account embeddings (*latent* space), as compared to over *observed* variables. Therefore, the optimization and learning of AMDN-HAGE requires bilevel optimization for jointly learning the model parameters and account embeddings (as discussed in the next section). The model proposed here aims to capture latent or hidden groups from activity traces, rather than from observed account features. This is because coordination patterns (like, activities of coordinated accounts might strongly influence each other) remain in activity traces and provide information about their collective behaviours, enabling identification of coordinated hidden groups.

For modeling distinct groups of coordinated and normal accounts, we used tied covariance Σ for all groups. Denoting the parameters (means, the shared covariance, and prior) of GMM as θ_g , the log-likelihood of social group modeling over accounts U is:

$$\begin{aligned} \sum_{j=1}^{|U|} \log p(u_j; \theta_g, E) &= \sum_{j=1}^{|U|} \log \sum_{i=1}^N p(u_j, i; \theta_g, E) \\ &= \sum_{j=1}^{|U|} \log \sum_{i=1}^N p(i) \mathcal{N}(E_{u_j}; \mu_i, \Sigma) \end{aligned} \quad (10)$$

Algorithm 1 Training Algorithm for AMDN-HAGE

Require: Activity traces (C_s), Account set (U)

Ensure: Generative model (θ_a, θ_g and E)

- 1: $\theta_a^{(0)}, E^{(0)} \leftarrow \arg\max_{\theta_a, E} \log p(C_s | U; \theta_a, E)$
 - 2: Set i as 1 {Iteration index}.
 - 3: **while** not converged **do**
 - 4: $\theta_g^{(i)} \leftarrow \arg\max_{\theta_g} \log p(U; E^{(i-1)}, \theta_g)$ using EM algorithm
 - 5: $\theta_a^{(i)}, E^{(i)} \leftarrow \arg\max_{\theta_a, E} \log p(C_s, U; \theta_g^{(i)}, \theta_a, E)$ using SGD or its variants
 - 6: $i \leftarrow i + 1$.
 - 7: **end while**
-

4.3 Jointly Learning

For joint learning, we maximize following log-likelihood:

$$\begin{aligned} \log p(C_s, U; \theta_g, \theta_a, E) &= \log p(C_s | U; \theta_a, E) + \log p(U; \theta_g, E) \\ &= \sum_{i=1}^L [\log p_{\theta_a, E}(t_i | H_{t_i}) + \log p_{\theta_a, E}(u_i | H_{t_i})] + \sum_{j=1}^{|U|} \log p(u_j; \theta_g, E) \end{aligned} \quad (11)$$

Above loss function has a trivial infinite asymptotic solution where all embedding vectors equal to the mean of a cluster and $\det(\Sigma) = 0$. To avoid this solution, we constraint $\det(\Sigma)$ to be greater than a small constant λ , lower bounding the loss function.

The log-likelihood (Eq 11) is a function of parameters θ_a, θ_g, E . In this, the optimization of the second term (Gaussian mixture of the shared latent embeddings), is a constrained optimization problem. Therefore, directly optimizing the joint likelihood with Stochastic Gradient Descent (SGD) or its variants like ADAM does not respect the constraints on mixture weights (normalized non-negative) and covariance (positive definite), leading to invalid log-likelihood in training (ablation study of loss with ADAM is in the expt section).

To address above disadvantages, we provide an equivalent bilevel optimization formulation to solve the joint learning problem:

$$\theta_a^*, E^* = \arg\max_{\theta_a, E} [\log p(C_s | U; \theta_a, E) + \max_{\theta_g} (\log p(U; \theta_g, E))] \quad (12)$$

$$\theta_g^* = \arg\max_{\theta_g} \log p(U; \theta_g, E^*) \quad (13)$$

Above bilevel optimization can be solved with iterative optimization. In each iteration, we first freeze E and θ_a , then estimate θ_g with EM algorithm. After that, we freeze θ_g and use SGD (or its variant) to optimize E and θ_a . Since the log-likelihood of *latent* embeddings is optimized in the second term (Eq. 11), we require initialization of embeddings by pre-training E and θ_a on observed sequences by maximizing the first term in the objective function, before jointly optimizing the two terms. Detailed algorithm is in Algorithm 1.

A generic concern to such alternating optimization algorithm is its convergence. Therefore, we give following theoretic guarantee that the proposed algorithm leads the model to converge at least on a local minimum or a saddle point with appropriate selection of the gradient based optimizer (denoting $L(\theta_a, E, \theta_g)$ as the negative log-likelihood (loss function)).

THEOREM 1. *Our proposed optimizing algorithm will converge at a local minimum or a saddle point if in any iteration i the neural network optimizer satisfies following conditions:*

- Given the frozen $\theta_g^{(i)}$ acquired by EM algorithm in iteration i , the neural network optimization algorithm we applied in converges at a local minimum or or a saddle point $(\theta_a^{(i)}, E^{(i)})$
- $L(\theta_a^{(i)}, E^{(i)}, \theta_g^{(i)}) \leq L(\theta_a^{(i-1)}, E^{(i-1)}, \theta_g^{(i)})$, where $\theta_a^{(i-1)}$ and $E^{(i-1)}$ are the starting points in iteration i

The proof can be found in the Appendix. Theoretically, the two conditions can be guaranteed when our loss function is L-smooth and we apply standard Gradient Descent Algorithm with learning rate lower than $\frac{1}{L}$ [15]. But in practice, since finding strict local minimum is not as important as training speed and generalization, we can alternatively apply Adam or other variants.

5 EXPERIMENT RESULTS

We verified the effectiveness of the proposed method AMDN-HAGE and training algorithm on real datasets collected from Twitter related to coordinated accounts by Russia’s Internet Research Agency, and for identification and analysis of coordination in COVID-19. We provide data collection, baselines and model variants. Details of implementation and experiment settings, are in the Appendix.¹

5.1 Data Collection

5.1.1 Russia’s Internet Research Agency (IRA) coordinated campaign. 2752 Twitter accounts were identified by the U.S. Congress² as coordinated accounts operated by the Russian agency (referred to as “troll farm”) to manipulate the U.S. Election in 2016. The social media posts (activities) of subset of these accounts were available through paid Twitter API access by the academic community to study coordinated account detection [2, 12]. We obtained the collected dataset from Luceri et al. which contains 312 of the Russian coordinated accounts (referred to as *coordinated “trolls”*) with their 1.2M tweets, and also includes 1713 normal accounts that participated in discussion about the U.S. Election during which the coordinated trolls were active (collected based on election related keywords using Twitter API); accounts in the collected dataset were active accounts with at least 20 active and passive tweets³ [12].

Activity traces: Activity traces C_s are constructed from the tweets, as any account’s posts and subsequent engagements from others (retweets, replies to the post) forming a time-ordered sequence of activities. The account set (2025 accounts) and activity traces are utilized to train AMDN-HAGE, and a held-out 15% of sequence subsets are used as validation loss of the model.

5.1.2 COVID-19 pandemic. Due to concerns around disinformation and social media abuse around COVID-19, we collect social media posts from with keywords related to COVID-19, using Twitter’s streaming API service from March 1 to July 22, 2020. We use the collected 119,298 active accounts, that have at least twenty active and passive collected tweets, and their 13.9M tweets.

Unknown coordination: The COVID-19 data does *not* have any labeled coordinated groups, unlike the IRA dataset. But it is important to examine if we can uncover any unknown coordinated

campaigns. We run AMDN-HAGE on the full 119k account set with their activity traces as in the IRA dataset, and examine tweets from identified coordinated groups, and the overlap of accounts with suspended Twitter accounts (manually suspended by Twitter for violations of platform policies). There can be violating accounts that have not yet been found and suspended by Twitter. Also, accounts that are suspended can be due to various reasons (e.g. spam, automation, multiple accounts) but Twitter suspensions are not restricted to coordinating accounts. Thus we cannot use Twitter suspensions to estimate precision-recall on coordination detection.

5.2 Baselines and Model Variant

We compare against existing approaches that utilize account activities to identify coordinated accounts. The baselines - extract features of coordination from account activities and use them for supervised or unsupervised detection of coordination, based on both individual or collective behaviours.

- (1) Unsupervised Baselines: Co-activity clustering [18] and Clickstream clustering [24] are based on pre-defined activity features. Co-activity models joint activity in content sharing, and Clickstream clustering models patterns in post, retweet and reply actions. The SOTA approach is IRL [12] based on inverse reinforcement learning to extract features from activity traces, for clustering coordinated accounts.
- (2) Supervised Baselines: IRL(S) [12] is sota supervised variant of inverse reinforcement approach which trains a supervised classifier on extracted features from activity traces (with labeled subset of accounts).

We add another baseline using HP (Hawkes Process) [32] to learn account features unsupervised from activity traces. HP(S) is the supervised variant. HP models the influence between accounts with an additive function. This baseline serves as a ablation of the proposed model to show that latent influence and interaction patterns for coordinated group is more complex and neural point process can better extract these coordination features.

For ablation of different components of proposed model AMDN-HAGE, we also compare with (i) AMDN (without hidden group estimation), which only learns the activity trace model; we can use it to extract account embeddings and cluster with GMM or KMeans to identify the coordinated group as anomalous group. (ii) AMDN-HAGE directly uses the jointly learnt GMM to output the group membership. AMDN-HAGE + Kmeans instead uses account embeddings from AMDN-HAGE with KMeans clustering to find the anomalous coordinated group. (iii) To compare with supervised setting (IRL (S) [12]), we similarly train a classifier on extracted features i.e., learned account embeddings to detect coordinated accounts (assuming subset of labeled coordinated and normal accounts are available for training). The variants are AMDN + NN and AMDN-HAGE + NN which use a two-layer MLP classifier on extracted embeddings from AMDN and AMDN-HAGE resp.

5.3 Results of Coordination Detection

Detection results on IRA dataset: We evaluate on two settings - unsupervised and supervised (as in earlier work [12]). In both, the proposed model is trained as unsupervised from activity traces to obtain the group membership and account embeddings.

¹Our code is available here: <https://github.com/USC-Melady/AMDN-HAGE-KDD21>

²<https://www.recode.net/2017/11/2/16598312/russia-twittertrump-twitter-deactivated-handle-list>

³Active tweet is where an account posts (tweet, retweet, reply) and passive is where the account is mentioned, retweeted, or replied [12]

Table 2: Results on detection of Russian coordinated disinformation campaign (IRA dataset) on Twitter in 2016 U.S. Election

Method (Unsupervised)	AP	AUC	F1@TH=0.5	Prec@TH=0.5	Rec@TH=0.5	MaxF1	MacroF1@TH=0.5
Co-activity	0.208 $\pm$ 0.01	0.592 $\pm$ 0.03	0.292 $\pm$ 0.02	0.206 $\pm$ 0.02	0.510 $\pm$ 0.04	0.331 $\pm$ 0.03	0.515 $\pm$ 0.02
Clickstream	0.169 $\pm$ 0.02	0.535 $\pm$ 0.04	0.215 $\pm$ 0.06	0.205 $\pm$ 0.05	0.228 $\pm$ 0.08	0.215 $\pm$ 0.06	0.532 $\pm$ 0.03
IRL	0.200 $\pm$ 0.00	0.610 $\pm$ 0.02	0.265 $\pm$ 0.02	0.219 $\pm$ 0.02	0.336 $\pm$ 0.03	0.340 $\pm$ 0.02	0.543 $\pm$ 0.01
HP	0.337 $\pm$ 0.04	0.694 $\pm$ 0.05	0.376 $\pm$ 0.05	0.387 $\pm$ 0.06	0.365 $\pm$ 0.05	0.545 $\pm$ 0.03	0.633 $\pm$ 0.03
AMDN + GMM	0.787 $\pm$ 0.05	0.894 $\pm$ 0.03	0.631 $\pm$ 0.06	0.965 $\pm$ 0.03	0.472 $\pm$ 0.07	0.738 $\pm$ 0.05	0.792 $\pm$ 0.03
AMDN + Kmeans	0.731 $\pm$ 0.08	0.901 $\pm$ 0.02	0.727 $\pm$ 0.06	0.806 $\pm$ 0.07	0.663 $\pm$ 0.06	0.752 $\pm$ 0.05	0.841 $\pm$ 0.03
AMDN-HAGE	0.804 $\pm$ 0.03	0.898 $\pm$ 0.02	0.699 $\pm$ 0.05	0.941 $\pm$ 0.04	0.558 $\pm$ 0.06	0.758 $\pm$ 0.04	0.828 $\pm$ 0.03
AMDN-HAGE + Kmeans	0.818 $\pm$ 0.04	0.935 $\pm$ 0.02	0.731 $\pm$ 0.04	0.913 $\pm$ 0.03	0.611 $\pm$ 0.05	0.776 $\pm$ 0.03	0.846 $\pm$ 0.02
Method (Supervised)	AP	AUC	F1@TH=0.5	Prec@TH=0.5	Rec@TH=0.5	MaxF1	MacroF1@TH=0.5
IRL (S)	0.672 $\pm$ 0.08	0.896 $\pm$ 0.03	0.557 $\pm$ 0.06	0.781 $\pm$ 0.06	0.436 $\pm$ 0.06	0.633 $\pm$ 0.07	0.749 $\pm$ 0.03
HP (S)	0.760 $\pm$ 0.04	0.925 $\pm$ 0.02	0.753 $\pm$ 0.02	0.743 $\pm$ 0.04	0.769 $\pm$ 0.06	0.782 $\pm$ 0.03	0.853 $\pm$ 0.01
AMDN + NN	0.814 $\pm$ 0.04	0.918 $\pm$ 0.02	0.733 $\pm$ 0.04	0.710 $\pm$ 0.05	0.761 $\pm$ 0.05	0.763 $\pm$ 0.04	0.841 $\pm$ 0.02
AMDN-HAGE + NN	0.838 $\pm$ 0.04	0.926 $\pm$ 0.03	0.769 $\pm$ 0.04	0.752 $\pm$ 0.05	0.789 $\pm$ 0.05	0.799 $\pm$ 0.04	0.862 $\pm$ 0.02

In the unsupervised, the group membership is directly used to report anomalous coordinated group. In supervised, the learned embeddings are used as features to train a classifier (to predict coordinated from normal accounts). The classifier is trained on subset of labeled coordinated and normal accounts in IRA dataset, with rest (stratified 20% in 5-folds) held-out for evaluation.

Table 2 provides results of model evaluation against the baselines averaged in 5-fold stratified cross-validation on the labeled normal and coordinated accounts in the IRA dataset over five random seeds. We compare the Average Precision (AP), area under the ROC curve (AUC), and F1, Precision, Recall, and MacroF1 at 0.5 threshold, and maxF1 at threshold that maximizes it. AMDN-HAGE outperforms other methods on both unsupervised and supervised settings, due to its ability to capture coordination characteristics with diverse account behaviours by learning latent influence and hidden group behaviours, without pre-specifying features relied on by other baselines.

Moreover, the coordination features learned with the proposed method are robust to unsupervised or supervised setting, unlike IRL and IRL(S) [12] (where even though IRL(S) can learn useful features, it performs poorly in unsupervised setting). In comparison, because AMDN-HAGE models the more *intrinsic behaviours* of coordination, it can extract patterns that can effectively identify anomalous coordinated groups in an unsupervised manner. The margin is larger on unsupervised than supervised setting, where group behaviours are more important, since there is no known set of coordinated accounts to train classifiers from extracted features.

Ablation of proposed model and training: Besides baselines, we also compare AMDN-HAGE with its variants to verify the importance of the joint learning and optimization algorithm.

To verify the importance of joint learning, in Table 2, AMDN-HAGE is compared with AMDN, which only learns the activity trace model, without hidden group estimation. Proposed model AMDN-HAGE captures consistently better coordination behaviours, indicating that modeling group or collective behaviour jointly is useful over only modeling the latent influence between account pairs through observed activity trace modeling.

Figure 3: Comparison of bilevel optimization and Adam.

To demonstrate the effectiveness of the bilevel training algorithm, we present the validation loss (negative log-likelihood on held-out 15% of activity traces) in the training process, comparing direct optimization of the joint log-likelihood using Adam (variant of SGD) and our proposed bilevel algorithm in Fig. 3. As we can see, for the proposed optimization, the loss on the validation set in both the pre-training and joint training stage decline and finally converge. However, in direct optimization with Adam the validation loss decreases to a point but breaks as it reaches an invalid parameter point. Without constraints, Adam reaches a covariance matrix that is not positive definite, and an invalid log-likelihood (NaN).

5.4 Analysis of Coordination Detection

5.4.1 Uncovering characteristic behaviours from influence structure. In this section, we examine the latent influence structure and account interactions learned by AMDN-HAGE on the IRA data. The latent influence is captured by the interpretable attention weights of the model, between account activities in observed traces. Higher attention paid by an event (account activity) on a history event (earlier activity from any account), indicates that the history event has a stronger triggering influence on the future event in consideration.

Figure 4: Analysis of learned influence strength between Coordinated ("trolls" T) and Normal (non-trolls NT) in IRA.

Figure 5: Analysis on how influence weights of account pairs vary with time difference. Green points: Normal (NT) account pairs. Blue: Coordinated ("trolls" T) account pairs.

In Fig4a, we compute the aggregate influence between account pairs learned with AMDN-HAGE over 5 random seeds (as average attention weight from account interactions over all activity traces). The strongest influence ties are between coordinated ("trolls") (T) accounts, and the least influence is between normal ("non-trolls (NT)") and coordinated (T) accounts and their activities. In Figure 4b, the learned account embeddings of coordinated and normal accounts with AMDN-HAGE in the IRA data is visualized. As we see, coordinated accounts form an anomalous cluster distinct from normal accounts, reflecting our model captures coordinated behaviours.

In Fig. 5a and 5b, we examine how influence weights of account pairs vary with time difference. In the two figures, each point represent an account pair appearing in the same activity sequence. Blue points refer to coordinated (T) pairs and green points refer to normal account (NT) pairs. The influence weights (y-axis) correspond to the time difference between activities of two accounts appearing in the sequence (x-axis) is shown for all points in Fig. 5a to reflect the overall trend. In Fig. 5b, we only plot points with highest influence weight in each time difference range of 24 hrs to reflect the trend on strongly interacting pairs in each window. From both figures, we can see that influence weights are higher for shorter time differences between account activities. However, the influence decreases faster for coordinated pairs than normal account pairs with the passage of time between their activities.

Table 3: Overlap between suspended Twitter accounts, and identified coordinated groups/ overall accounts in collected COVID-19 data.

Twitter Suspensions	Overall Overlap	Cluster 1 Anomaly (3.7k)	Cluster 2 Anomaly (5.5k)
Suspended (9k)	7.544 %	12.19 %	11.94 %
State-backed (81)	0.067 %	0.13 %	0.09 %
Coupled (602)	0.504 %	0.72 %	1.33 %
Targets (18.5k)	15.507 %	14.98 %	17.11 %

Table 4: Representative tweets in disinformation topic clusters in identified COVID-19 coordinated accounts groups.

Did coronavirus leak from a research lab in Wuhan? Startling new theory is 'no longer being discounted' amid claims staff 'got infected after being sprayed with blood' #WWG1WGA #QAnon #MAGA #Trump2020 #COVID19
#scamdemic Since China owns WHO, China must be in on the scam The WHO Lied and Created a Global Panic: Second Extensive Study Finds Coronavirus Mortality Rate Is 0.4% Not 3.4% - Similar to Seasonal Flu
VIRUS FRAUD - FAKE NEWS In Spite of Leftist Media Hysteria at the Time, SD Governor Noem Confirms there were ZERO New Cases or 'Outbreaks' over Trump's Rushmore Event, Trump Supporters are Clean, Healthy People...
BREAKING: #BillGates Foundation And The #Covid19 VACCINE NETWORK SCANDAL The British People Are Going To Get Very Angry Very Soon And Will Want ANSWERS.#GlaxoSmithKline #BillGates #Rothschild #WellcomeTrust #COVID19 #CivilService #NoMasks
#BillGates Negotiated \$100 Billion #ContactTracing Deal With #Democratic Congressman in 8/2019 well before #Coronavirus #Pandemic starts. #Gates holds pandemic drill 10/2019 Harvard finds #SARSCOV2 started in 8/2019 in #wuhan. Virus in USA by 1/2020.

This phenomenon suggests that compared to normal social influence, the coordination is more temporal. Meanwhile, it is noticeable that a small fraction of coordination last even longer than 2000 hours (nearly 3 months), reflecting that the assumption on synchronization may fail in detecting such long-term coordination.

5.4.2 Uncovering coordinated groups in COVID-19 data. As mentioned earlier, the data collected on COVID-19 does not contain a ground-truth set of labeled coordinated accounts. But, we can use the proposed method AMDN-HAGE to uncover any suspicious coordinated accounts with analysis of features. AMDN-HAGE is trained on observed account activities in the COVID-19 data of 119k accounts. The method identifies two anomalous clusters of accounts (based on clustering silhouette scores, provided in appendix). We inspect the feature distribution in each account group.

In Table 3, we compare the distribution of suspended Twitter accounts. Twitter additionally labels some suspended accounts as state-backed i.e, accounts Twitter finds as linked to state-sponsored operations (such as from Russia, etc. that tried to interfere with politics in other countries) [8]. In addition, we consider the "coupled" accounts in the collected data that bidirectionally engage with Twitter's state-backed accounts, and "targets" as accounts mentioned (or targeted) by state-backed accounts in their tweets. Amongst all the 119k accounts, the distribution of Twitter accounts (Suspended,

Figure 6: Top-35 (most frequent) unique hashtags in tweets of identified coordinated group and normal accounts.

State-backed, Coupled) was between 1.5-2 times higher than by random chance in the identified anomalous clusters, even when the number of such accounts to be found from the large set of collected accounts is small. For Targets, the distribution is more uniform, since targeted accounts unlike Coupled only capture a unidirectional engagement attempt from state-backed accounts, which is expected because state-backed accounts attempt to manipulate and thus mention or engage with other normal accounts.

In Fig 6, we find most frequent hashtags in tweets posted by accounts in the groups, and plot the top hashtags unique to each group (hashtags of the smaller anomalous cluster is provided in the appendix). We find that the prominent top hundred hashtags in the coordinated group promote anti-mask and anti-vaccine (“No-Masks”, “NoVaccine”, “NoAlaVacuna”), and anti-science theories (“Plandemic”, “Covid-Hoax”), and contain hashtags associated with “QAnon” (“WWG1WGA”), a notorious far-right conspiracy group.

In Table 4 we use topic modeling to find the most representative disinformation tweets posted by the anomalous accounts. We find 4 topic clusters within tweets that are linked to low-credibility (disinformation) news sources. The table shows tweets closest to the topic cluster centers. The narrative in the hashtags and topics suggests strong presence of tweets in Spanish and English (NoAlaMascarillaObligatoria, NoAlaVacuna, NoAlNuevoOrdenMundia, NoVaccine) about no new world order, no masks, no vaccine, and QAnon, all opposing Bill Gates, and suggesting that COVID-19 is a hoax and deep state or political scam to monetize vaccines.

6 CONCLUSIONS

In this work, we proposed a technique to detect coordinated accounts based on their collective behaviours, inferred directly from their activities on social media. The proposed method is independent of linguistic, metadata or platform specific features, and hence can be generalized across platforms and languages or countries from where disinformation campaigns originate. Such features can also be easily incorporated in the proposed model.

With analysis on Russian Interference and COVID-19 datasets, we investigated the behaviours of identified coordinated accounts, finding that influence between coordinated accounts is higher, and influence between coordinated pair of accounts decreases faster than non-coordinated pairs over time. In the COVID-19 data, we identified coordinated groups, and analysis suggests that the main narrative in the coordinated group is that COVID-19 is a hoax and political scam, with anti-vaccine, anti-mask social media posts.

ACKNOWLEDGMENTS

This work is supported by NSF Research Grant (IIS-1254206). Views and conclusions are of the authors and should not be interpreted as representing the social policies of the funding agency, or U.S. Government. Emilio Ferrara acknowledges support by the Air Force Office for Scientific Research (grant no. FA9550-20-1-0224). We are also thankful to Jihao Long for his help in verifying the correctness of the proof of Theorem 1.

REFERENCES

- [1] Aseel Addawood, Adam Badawy, Kristina Lerman, and Emilio Ferrara. 2019. Linguistic cues to deception: Identifying political trolls on social media. In *ICWSM*.
- [2] Adam Badawy, Aseel Addawood, Kristina Lerman, and Emilio Ferrara. 2019. Characterizing the 2016 Russian IRA influence campaign. *SNAM* 9, 1 (2019), 31.
- [3] Jeff Baumes, Mark Goldberg, Malik Magdon-Ismael, and William Al Wallace. 2004. Discovering hidden groups in communication networks. In *ICIS*. Springer.
- [4] Qiang Cao, Xiaowei Yang, Jieqi Yu, and Christopher Palow. 2014. Uncovering large groups of active malicious accounts in online social networks. In *Proceedings of the 2014 ACM Conference on Computer and Communications Security*. 477–488.
- [5] Daryl J Daley and David Vere-Jones. 2007. *An introduction to the theory of point processes: volume II: general theory and structure*. Springer.
- [6] Nan Du, Hanjun Dai, Rakshit Trivedi, Utkarsh Upadhyay, Manuel Gomez-Rodriguez, and Le Song. 2016. Recurrent marked temporal point processes: Embedding event history to vector. In *ACM SIGKDD*. 1555–1564.
- [7] Emilio Ferrara, Onur Varol, Clayton Davis, Filippo Menczer, and Alessandro Flammini. 2016. The rise of social bots. *Commun. ACM* 59, 7 (2016), 96–104.
- [8] Vijaya Gadde and Yoel Roth. 2018. Enabling further research of information operations on Twitter. *Twitter Blog* 17 (2018).
- [9] Sonu Gupta, Ponnurangam Kumaraguru, and Tanmoy Chakraborty. 2019. Malreg: Detecting and analyzing malicious retweeter groups. In *Proceedings of the India Joint International Conference on Data Science and Management of Data*. 61–69.
- [10] Douglas M Hawkins. 1980. *Identification of outliers*. Vol. 11. Springer.
- [11] Jane Im, Eshwar Chandrasekharan, Jackson Sargent, Paige Lighthammer, Taylor Denby, Ankit Bhargava, Libby Hemphill, David Jurgens, and Eric Gilbert. 2020. Still out there: Modeling and identifying russian troll accounts on twitter. In *ACM Web Science*. 1–10.
- [12] Luca Luceri, Silvia Giordano, and Emilio Ferrara. 2020. Detecting troll behavior via inverse reinforcement learning: A case study of Russian trolls in the 2016 US election. In *ICWSM*, Vol. 14. 417–427.
- [13] Diego A Martin, Jacob N Shapiro, and Michelle Nedashkovskaya. 2019. Recent trends in online foreign influence efforts. *Journal of Information Warfare* (2019).
- [14] Hongyuan Mei and Jason M Eisner. 2017. The neural hawkes process: A neurally self-modulating multivariate point process. In *NIPS*. 6754–6764.
- [15] Yurii Nesterov. 1998. Introductory lectures on convex programming volume i: Basic course. *Lecture notes* 3, 4 (1998), 5.
- [16] Takahiro Omi, Kazuyuki Aihara, et al. 2019. Fully neural network based model for general temporal point processes. In *NIPS*. 2122–2132.
- [17] Diogo Pacheco, Pik-Mai Hui, Christopher Torres-Lugo, Bao Tran Truong, Alessandro Flammini, and Filippo Menczer. 2021. Uncovering Coordinated Networks on Social Media. *ICWSM*.

- [18] Natali Ruchansky, Sungyong Seo, and Yan Liu. 2017. CSI: A Hybrid Deep Model for Fake News Detection. In *CIKM*. ACM, 797–806.
- [19] Karishma Sharma, Xinran He, Sungyong Seo, and Yan Liu. 2021. Network Inference from a Mixture of Diffusion Models for Fake News Mitigation. In *ICWSM*.
- [20] Karishma Sharma, Feng Qian, He Jiang, Natali Ruchansky, Ming Zhang, and Yan Liu. 2019. Combating Fake News: A Survey on Identification and Mitigation Techniques. *ACM TIST* (2019).
- [21] Karishma Sharma, Sungyong Seo, Chuizheng Meng, Sirisha Rambhatla, and Yan Liu. 2020. Coronavirus on social media: Analyzing misinformation in Twitter conversations. *arXiv preprint arXiv:2003.12309* (2020).
- [22] Oleksandr Shchur, Marin Bilos, and Stephan Günnemann. 2020. Intensity-Free Learning of Temporal Point Processes. *ICLR*.
- [23] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NeurIPS*. 5998–6008.
- [24] Gang Wang, Xinyi Zhang, Shiliang Tang, Haitao Zheng, and Ben Y Zhao. 2016. Unsupervised clickstream clustering for user behavior analysis. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 225–236.
- [25] Samuel C Woolley and Philip N Howard. 2018. *Computational propaganda: political parties, politicians, and political manipulation on social media*. Oxford University Press.
- [26] Da Xu, Chuanwei Ruan, Evren Korpeoglu, Sushant Kumar, and Kannan Achan. 2019. Self-attention with functional time representation learning. In *NeurIPS*.
- [27] Hongteng Xu. 2018. PoPPy: A Point Process Toolbox Based on PyTorch. *arXiv preprint arXiv:1810.10122* (2018).
- [28] Rose Yu, Xinran He, and Yan Liu. 2015. Glad: group anomaly detection in social media analysis. *ACM Transactions on Knowledge Discovery from Data* 10, 2 (2015).
- [29] Savvas Zannettou, Tristan Caulfield, Emiliano De Cristofaro, Michael Sirivianos, Gianluca Stringhini, and Jeremy Blackburn. 2019. Disinformation warfare: Understanding state-sponsored trolls on Twitter and their influence on the web. In *Companion proceedings of WebConf*. 218–226.
- [30] Savvas Zannettou, Tristan Caulfield, William Setzer, Michael Sirivianos, Gianluca Stringhini, and Jeremy Blackburn. 2019. Who let the trolls out? towards understanding state-sponsored trolls. In *ACM WebSci*. 353–362.
- [31] Qiang Zhang, Aldo Lipani, Omer Kirnap, and Emine Yilmaz. 2020. Self-Attentive Hawkes Process. *ICML*.
- [32] Ke Zhou, Hongyuan Zha, and Le Song. 2013. Learning social infectivity in sparse low-rank networks using multi-dimensional Hawkes processes. In *AISTATS*.
- [33] Simiao Zuo, Haoming Jiang, Zichong Li, Tuo Zhao, and Hongyuan Zha. 2020. Transformer Hawkes Process. *NeurIPS*.

A APPENDIX

A.1 Details of the Baselines

Here we provide the details of our baselines:

- (1) Co-activity clustering [18]. Co-activity features to identify coordinated groups as accounts that repeatedly share the same contents, was proposed in Ruchansky et al. Account features are extracted using SVD on a binary event-participation matrix (of accounts and its posts i.e., tweets, retweets, replies). Clustering on the features are used to detect coordination.
- (2) Clickstream clustering [17, 24]. It is proposed by Wang et al. to analyze account behaviours, and is based on hierarchical clustering of accounts with similar activity patterns to identify coordinated groups. Similarity between activities is represented based on post, reply and re-share patterns.
- (3) IRL (S) [12]. SOTA approach is based on inverse reinforcement learning to discover motives of coordinated accounts from rewards estimated from activity traces. Estimated rewards are used as features for detection. IRL(S) trains AdaBoost on these features on a labeled subset of coordinated and normal accounts. Luceri et al. reported different classifiers in their paper, including two-layer MLP and AdaBoost, with the latter outperforming others on their features.
- (4) IRL [12]. We also compare an unsupervised variant of the SOTA IRL approach based on clustering with GMM or K-means clustering to detect coordinated groups.

- (5) HP and HP (S). In addition, we compare our method to modeling activities using the Hawkes Process (HP)⁴ with influence α_{ij} factorized by learnable embeddings of accounts i and j , with both clustering and supervised detection.

A.2 Details of Experiments

Implementation details We use activity sequences of maximum length 128, splitting longer sequences, batch size of 256 on 4 NVIDIA-2080Ti, embedding dimension in {32, 64}, number of mixture components for the PDF in the AMDN part between {8,16,32}, single head and single layer attention module. As for the component number in the HAGE part, it is set as 2 for IRA dataset and 3 for COVID dataset based on silhouette scores on learned embeddings from the pre-train step of training algorithm (Fig. 7). For the shared covariance matrix of HAGE part, we constrain it to be a diagonal matrix. The second term of the loss is scaled for embedding dimension and size of user set. We implement the model and training algorithm entirely in PyTorch and use Adam with 1e-3 learning rate and 1e-5 regularization to optimize. We train for max 1000 epochs with early stopping based on validation likelihood of sequences (75/15/10 splits). The link⁵ to anonymized code and collected COVID-19 data is in the footnote, and will be made non-anonymous if accepted.

Figure 7: Selection of number of clusters based on silhouette scores in COVID-19 and IRA datasets.

Also, we present some numeric results and statistic properties we found on the COVID-19 dataset. We compare the distribution of account creation years in COVID-19 dataset for each identified accounts group or cluster. As we can see, the ratio of accounts created in recent years is higher in the two detected anomalous clusters, which is consistent to the reality that coordinated campaigns are

⁴Hawkes process code [27] to extract account embeddings from activity traces

⁵Anonymized AMDN-HAGE code and COVID-19 data are available here (clickable).

Figure 8: Distribution of account creation years in COVID-19 dataset for each identified accounts group or cluster.

Figure 9: Hashtag distribution for cluster C1 (COVID-19 dataset).

raised in recent years. Also, we present the hashtag distribution in the C1 cluster in Figure 9. Surprisingly, although the ratio of suspended coordinated users in this cluster is significantly higher than normal, the hashtags look not like disinformation contents or topics. Upon further inspection, we find the political hashtags like “Resist” or “Resistance” in the top-50 unique hashtags of this cluster, refer to U.S. liberal political movement against the reigning president, and have not been linked to disinformation or conspiracies. This phenomenon suggests coordination may not be specific to disinformation spreading. In normal information spreading, there could be the influence from coordinated campaigns.

A.3 Details of Neural Point Process Models

Methods SAHP [31] and THP [33] (Table 1 require Monte Carlo sampling to estimate an integral in the likelihood function used in MLE, as their intensity formulation does not have closed form solutions. This leads to more noisy approximation of gradients in training the model. On the other hand, HP [32], RMTTP [6] define specific functional forms, e.g. exponential for intensity parameterization to have closed form likelihoods but are limited in the flexibility of the selected functional form of the intensity. LogNormMix [22] and FullyNN [16] alleviate these issues by modeling the conditional density and cumulative density getting flexibility and closed-form likelihoods. However, the neural network parameterizations in RMTTP, LogNormMix, FullyNN are based on recurrent neural networks and cannot be directly used to examine the influence between events and event types, which is possible in the other methods. We address this by modeling the conditional density function with interpretable neural network parameterizations, so that the influence structure can be learned and examined to understand coordinated behaviours of malicious accounts, and still have the same form for the point process conditional PDF as LogNormMix.

A.4 Proof of Convergence

Here we provide the proof to Theorem 1.

PROOF. We first prove that the training algorithm converges. In EM, each step increases the likelihood function of a mixture model.

We have:

$$\log P(E^{(i-1)}|\theta_g^{(i)}) \geq \log P(E^{(i-1)}|\theta_g^{(i-1)}) \quad (14)$$

Thus, we obtain

$$L(\theta_a^{(i-1)}, E^{(i-1)}, \theta_g^{(i)}) \leq L(\theta_a^{(i-1)}, E^{(i-1)}, \theta_g^{(i-1)}). \quad (15)$$

Then, from the second condition, we know that:

$$L(\theta_a^{(i)}, E^{(i)}, \theta_g^{(i)}) \leq L(\theta_a^{(i-1)}, E^{(i-1)}, \theta_g^{(i)}) \quad (16)$$

Therefore, we have:

$$L(\theta_a^{(i)}, E^{(i)}, \theta_g^{(i)}) \leq L(\theta_a^{(i-1)}, E^{(i-1)}, \theta_g^{(i-1)}) \quad (17)$$

which means the loss function monotonically decreases. Since we constraint the variance of the mixture model in both point processing model and social group model larger than a constant ϵ , the loss function is bounded by a constant C on a given activity trace set:

$$L(\theta_a^{(i)}, E^{(i)}, \theta_g^{(i)}) \geq C \quad (18)$$

Thus the loss function converges when i increases. Then we prove that the loss function converges at a local minimum or a saddle point. First when the parameters converges, we have:

$$L(\theta_a^{(i)}, E^{(i)}, \theta_g^{(i)}) = L(\theta_a^{(i+1)}, E^{(i+1)}, \theta_g^{(i+1)}). \quad (19)$$

Thus:

$$L(\theta_a^{(i+1)}, E^{(i+1)}, \theta_g^{(i)}) = L(\theta_a^{(i+1)}, E^{(i+1)}, \theta_g^{(i+1)}) \quad (20)$$

$$\log P(E^{(i)}|\theta_g^{(i+1)}) = \log P(E^{(i)}|\theta_g^{(i)}). \quad (21)$$

In EM, if $\log P(E^{(i)}|\theta_g^{(i+1)}) = \log P(E^{(i)}|\theta_g^{(i)})$ then $\theta_g^{(i)} = \theta_g^{(i+1)}$. Since $(\theta_a^{(i)}, E^{(i)})$ is a local minimum or a saddle point, we have:

$$\frac{\partial L(\theta_a^{(i)}, E^{(i)}, \theta_g^{(i)})}{\partial \theta_a^{(i)}} = \frac{\partial L(\theta_a^{(i)}, E^{(i)}, \theta_g^{(i)})}{\partial E^{(i)}} = 0 \quad (22)$$

Since EM is known to converge to a local minimum, we have:

$$\frac{\partial \log P(E^{(i)}|\theta_g^{(i+1)})}{\partial \theta_g^{(i+1)}} = \frac{\partial \log P(E^{(i)}|\theta_g^{(i)})}{\partial \theta_g^{(i)}} = 0 \quad (23)$$

Therefore, $(\theta_a^{(i)}, E^{(i)}, \theta_g^{(i)})$ is a local minimum or a saddle point. $\square$