

"© Quyu Kong, Pio Calderon, Rohit Ram, Olga Boichak, Marian-Andrei Rizoiu | ACM {2023}. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in WWW '23: Proceedings of the ACM Web Conference 2023, <http://doi.org/10.1145/3543507.3583481>

Interval-censored Transformer Hawkes: Detecting Information Operations using the Reaction of Social Systems

Quyu Kong
Alibaba Group &
University of Technology Sydney
Hangzhou, China
kongquyu.kqy@alibaba-inc.com

Pio Calderon
University of Technology Sydney
Sydney, Australia
pio.calderon@student.uts.edu.au

Rohit Ram
University of Technology Sydney
Sydney, Australia
rohit.ram@student.uts.edu.au

Olga Boichak
University of Sydney
Sydney, Australia
olga.boichak@sydney.edu.au

Marian-Andrei RizoIU
University of Technology Sydney
Sydney, Australia
marian-andrei.rizoIU@uts.edu.au

ABSTRACT

Social media is being increasingly weaponized by state-backed actors to elicit reactions, push narratives and sway public opinion. These are known as Information Operations (IO). The covert nature of IO makes their detection difficult. This is further amplified by missing data due to the user and content removal and privacy requirements. This work advances the hypothesis that the very reactions that Information Operations seek to elicit within the target social systems can be used to detect them. We propose an Interval-censored Transformer Hawkes (IC-TH) architecture and a novel data encoding scheme to account for both observed and missing data. We derive a novel log-likelihood function that we deploy together with a contrastive learning procedure. We showcase the performance of IC-TH on three real-world Twitter datasets and two learning tasks: future popularity prediction and item category prediction. The latter is particularly significant. Using the retweeting timing and patterns solely, we can predict the category of YouTube videos, guess whether news publishers are reputable or controversial and, most importantly, identify state-backed IO agent accounts. Additional qualitative investigations uncover that the automatically discovered clusters of Russian-backed agents appear to coordinate their behavior, activating simultaneously to push specific narratives.

1 INTRODUCTION

Online social platforms are explicitly designed to boost user interaction and engagement [31]. Recently, there has been an increase in false and misleading claims being deliberately propagated and legitimized as tools of foreign interference [30] — also known as computational propaganda or Information Operations (IO). IO actors leverage a wide spectrum of problematic online content (from satire and parody to manipulated and outright fabricated content) and behavior (inauthentic coordination) to elicit reactions and sway public opinion [29]. Content- and user-based detection tools are notoriously difficult to build due to language nuances and suffer from language drift and adversarial attacks. Here, we start from the hypothesis that the actions of IO actors are designed to elicit particular reactions from the target audience. This makes analyzing the posting behavior of users and the reaction of the social systems surrounding them a helpful starting point to map actors and agendas

behind those posts. This paper presents a novel approach to model the reactions of social systems, to disentangle the structural and functional types of users from the content posted by them. Based solely on the timing of retweet cascades, we can separate YouTube videos' categories, quantify news publishers' trustworthiness, and even identify users involved in state-linked coordinated operations.

The ubiquitous diffusion cascades generated in the online environment, such as reshare events on Twitter and replies on Reddit, unveil the temporal dynamics of online items [20, 38] and have been shown to be a proxy for characterizing them [15]. A typical difficulty when analyzing reshare cascades is partial data loss — i.e., when some of the events in the cascade are missing. For instance, the IO dataset made available by the Twitter Moderation Research Consortium [1] contains tweets from users relating to state-linked coordinated operations; however, for privacy reasons, it does not contain any of the activities of organic, non-malicious users. Besides, a recent study by Wu et al. [33] comprehensively examined the sampling effect of retweet cascades collected from Twitter, where the loss of tweets may lead to modeling bias. We ask the question: **can we model reshare cascades with partial data loss, containing both event times and missing event counts?** This work develops the Interval-censored Transformer Hawkes (IC-TH) model, an enhancement of the current state-of-the-art temporal point process model, the Transformer Hawkes model [36]. IC-TH leverages two sources of information — the event timestamps and the counts of missing events. We propose a unified data encoding scheme that jointly accounts for the two data sources, and we derive a novel log-likelihood function that augments the canonical point process likelihood function [9]. Using synthetic experiments, we evaluate IC-TH robustness concerning data loss. We find that cascades generated from distinct parameter combinations can be near-perfectly separated even when 90% of the events in each cascade are missing.

Characterizing problematic content (such as mis- and dis-information) or identifying malicious users typically demands significant manual labeling efforts. In practice, only limited amounts of labeled data are available, raising the question: **how do we leverage unlabeled cascade data for learning the representation of online items and users?** Inspired by recent progress in pre-training deep neural nets [11, 34], we propose a contrastive learning procedure to train IC-TH on unlabeled cascade groups. We build positive and

negative pairs by sampling from cascade groups — a cascade group contains all the cascades pertaining to a given online item or in which a given user participates. We train the model to differentiate cascades from different groups and maximize the similarities among cascades within the same group. Contrastive learning enables IC-TH to be pre-trained on large, unrelated datasets and later fine-tuned on the available labeled data. We show that pre-training improves performance when only small amounts of labeled data are available.

The final question is **can we employ IC-TH to disentangle between content types and detect malicious online users?** We test IC-TH on several tasks and case studies. We study two tasks — popularity prediction and online item categorization — on three real-world cascade datasets (ACTIVERT2017, RNCNIX and IO). We show that IC-TH consistently outperforms the current state-of-the-art models, such as Mean Behavior Poisson [25], HawkesN [16] and Transformer Hawkes [41] for popularity prediction, and Dual Mixture Model [15] and Transformer Hawkes [41] for item classification. Specifically, we augment the IO dataset with organic Twitter users and the prediction result on the IO dataset shows a macro-F1 of 98.7% at separating users relating to Iran-, Russia- and Saudi Arabia-backed information operations and organic users. Furthermore, embeddings generated by IC-TH allow for analyzing the strategies of groups of users. We identify three clusters of Russia-backed users; the users in each cluster appear to be coordinated in their behavior, activating simultaneously to push specific narratives such as the *#columbiachemicals* hoax in 2014.

The main contributions of this work are:

- The interval-censored Transformer Hawkes for modeling the mixed temporal data formats of events and event counts, along with a unified temporal data coding scheme and a log-likelihood function.
- A pre-training procedure based on contrastive learning to leverage the large number of unlabeled cascade groups relating to different online items/users.
- A set of experimental results on large real-world Twitter cascade datasets where we observe improved performances on item/user category and final popularity predictions from the proposed model and the pre-training procedure.
- Qualitative and quantitative analysis of user embeddings generated by our model on the IO dataset.

Related work. Diffusion cascades are commonly observed in the forms of individual events and event counts in between observation times. Event-based point process models employ different intensity functions (e.g., predefined parametric forms [13, 20, 24, 38, 39], deep neural nets [12, 19, 21, 36, 41] and non-parametric variants [37, 40]) to characterize the likelihood of event emergence, and the model parameters are typically estimated from the general point process likelihood function [9]. On the other hand, count-based point process models generally acknowledge the loss in the interval-censored raw data and attempt to uncover the underlying parameters of the corresponding event-based models. For instance, Kirchner [14] shows that the Hawkes process parameters can be obtained from event-count data using an integer-valued auto-regressive (INAR) model, and Rizoïu et al. [25] connect the Hawkes processes with a non-homogeneous Poisson process that can be fitted on event

counts. This work extends both classes of models by proposing a Hawkes process variant based on the Transformer architecture with a novel likelihood function mixing both events and count data.

On the application front, prior works use the models above, including the Hawkes processes [15] and deep learning models [23, 26] to encode online items and users based on their temporal activities for detecting coordinated accounts from Twitter, while earlier works apply pre-defined features to summarize online user behavior [2, 35]. Our work leverages a model pre-training procedure that further learns from a large unlabeled dataset to improve the derived embeddings of online items and users.

2 PRELIMINARIES

This section covers the required preliminaries concerning modeling event cascades and series of interval-censored event counts. We first mathematically define cascades in two distinct forms: event timestamps and event counts. We then introduce the Hawkes process [13] — alongside its extensions HawkesN [24], NeuralHawkes [19] and Transformer Hawkes [41] — and Mean Behavior Poisson (MBP) process [25] for modeling event times and event counts, respectively. As far as we are aware, no existing model can handle mixtures of event times and counts of missing events; therefore, a gap exists between models for interval-censored data and models for event cascades.

Diffusion cascades. On online social media platforms, such as Twitter, users can post/tweet content, and others may reshare or retweet it, resulting in cascades of retweet events [3]. On other platforms, the sharing events may be only available as aggregated counts with selected granularity, such as Youtube, where only daily view counts are available [32]. Formally, an event cascade, $\mathcal{H}_T = \{t_0, t_1, \dots\}$, is a set of timestamps of individual sharing events observed until time T . In the interval-censored setting, a set of observation times, $O_T = \{o_0, o_1, \dots\}$, partitions an event cascade into time segments, and only the number of events $C(o_i, o_{i+1}]$ in each segment is observed. Prior works successfully model diffusion cascades with both data forms [25, 38], as we discuss next.

2.1 Models for event data

Hawkes process is a particular type of point process model with the self-exciting property — i.e., the future event intensity depends on all past events [13]. Its event intensity at time t is given as

$$\lambda(t) = \mu + \sum_{t_i \in \mathcal{H}_T} \phi(t - t_i), \quad (1)$$

where $\phi : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is known as the decay function capturing the intensity excited by past events and μ is the background intensity which is normally considered 0 for reshare cascades [38]. Common choices for the decay function include the exponential function, $\phi_{EXP}(\tau) = \kappa\theta e^{-\theta\tau}$, and the power-law function, $\phi_{PL}(\tau) = \kappa(\tau + c)^{-(1+\theta)}$ [16]. A finite-population extension of Hawkes processes, dubbed **HawkesN processes** [24], has shown superior modeling performance on online social media data [16]. The intensity function of HawkesN processes is modulated by the proportion of remaining unaffected individuals in the population, i.e.,

$$\lambda(t) = \frac{N - N_t}{N} \sum_{t_i \in \mathcal{H}_T} \phi(t - t_i). \quad (2)$$

where N is the population size and N_t is the number of events up to time t .

Neural Hawkes process generalizes the intensity function of vanilla Hawkes process with deep recurrent neural networks (RNNs) [19, 21], i.e.,

$$\lambda^N(t) = f(\mathbf{w}^\top \mathbf{h}(t)), \quad (3)$$

where \mathbf{w} is a weight matrix and $\mathbf{h}(t)$ is a hidden state that encodes event history information produced by a continuous-time LSTM. $f(x) = \beta \log(1 + e^{\frac{x}{\beta}})$ is a softplus function parameterized by β .

Transformer Hawkes (TH) process [36, 41] is the natural extension to the Neural Hawkes processes, leveraging an attention-based model [27] that showed superior performance in sequence modeling compared to RNNs. The hidden state of Transformer Hawkes is obtained via a self-attention module. Specifically,

$$\mathbf{h}(t_j) = \mathbf{H}(\mathbf{j}, :), \quad (4)$$

$$\mathbf{H}(\mathbf{j}, :) = \text{ReLU}(\mathbf{S}\mathbf{W}_1 + \mathbf{b}_1)\mathbf{W}_2 + \mathbf{b}_2, \quad (5)$$

$$\mathbf{S} = \text{Concat}(\text{head}_1, \text{head}_2, \dots)\mathbf{W}^O, \quad (6)$$

$$\text{head}_i = \text{Softmax}\left(\frac{\mathbf{X}\mathbf{W}_i^Q(\mathbf{X}\mathbf{W}_i^K)^\top}{\sqrt{d_k}}\right)\mathbf{X}\mathbf{W}_i^V, \quad (7)$$

where $\mathbf{W}_i^Q, \mathbf{W}_i^K \in \mathbb{R}^{d_m \times d_k}$, $\mathbf{W}_i^V \in \mathbb{R}^{d_m \times d_v}$, $\mathbf{W}^O \in \mathbb{R}^{hd_v \times d_m}$, $\mathbf{W}_1 \in \mathbb{R}^{d_m \times h}$, $\mathbf{b}_1 \in \mathbb{R}^h$, $\mathbf{W}_2 \in \mathbb{R}^{h \times d_m}$ and $\mathbf{b}_2 \in \mathbb{R}^{d_m}$ are learned weights, d_m is the embedding dimension, d_k, d_v are the hidden dimensions of the projection subspace, and h is the number of heads. The input \mathbf{X} is obtained via the temporal encoding, i.e.,

$$\mathbf{X}(\mathbf{j}, i) = \begin{cases} \cos(t_j/1000^{\frac{i-1}{d_m}}) & \text{if } i \text{ is odd,} \\ \sin(t_j/1000^{\frac{i}{d_m}}) & \text{if } i \text{ is even.} \end{cases} \quad (8)$$

where trigonometric functions are used to encode each event timestamp t_j into a d_m -dimensional feature vector, $\mathbf{X}(\mathbf{j}, :)$.

Parameter estimation. The parameters of the models mentioned above are estimated via maximizing the general point process log-likelihood [9], i.e.,

$$\mathcal{L}_{\text{LL}} = \sum_{t_i \in \mathcal{H}_T} \lambda(t_i) - \int_0^T \lambda(\tau) d\tau. \quad (9)$$

2.2 Models for interval-censored data

Mean Behavior Poisson (MBP) processes are a particular variant of Poisson processes that fit on interval-censored data. MBP is defined as the point process whose event intensity is the mean Hawkes process intensity over all the possible realizations of a parameter set. As a result, MBP’s parameters directly correspond to their Hawkes process counterparts [7, 25]. Specifically, given a Hawkes process in Eq. (1), its MBP equivalence is defined by a deterministic intensity function

$$\xi(t) = \mu + \int_0^t \xi(\tau)\phi(t-\tau)d\tau, \quad (10)$$

where $\xi(t) = \mathbb{E}_{\mathcal{H}_t}[\lambda(t)]$. The MBP log-likelihood is

$$\mathcal{L}_{\text{IC-LL}}(\theta) = \sum_{i=1}^m C(o_{i-1}, o_i) \log \Xi(o_{i-1}, o_i; \theta) - \sum_{i=1}^m \Xi(o_{i-1}, o_i; \theta) \quad (11)$$




Figure 1: 4 observed events (lollipops with solid lines) of an example retweet cascade with the presence of unobserved retweet events (lollipops with dotted lines). Only retweet counts are given (i.e., $rtc_0, rtc_1, rtc_2, \dots$) by Twitter API and the exact timestamps of unobserved events are missing.

where $C(o_{i-1}, o_i)$ is the number of events recorded in the interval (o_{i-1}, o_i) , and we define the compensator

$$\Xi(o_{i-1}, o_i; \theta) = \int_{o_{i-1}}^{o_i} \xi(z) dz \quad (12)$$

3 INTERVAL-CENSORED TRANSFORMER HAWKES

This section first explores where the mixing of event counts and event times is applicable (Section 3.1). We then propose to encode such data in a flexible format, extend the Transformer Hawkes to model it, and propose a novel log-likelihood function (Section 3.2). Finally, we discuss a contrastive learning and pre-training procedure for training IC-TH without large volumes of labeled data (Section 3.3).

3.1 Missing events on Twitter

Twitter sampling in data collection. While existing studies modeling information diffusions on Twitter generally analyze diffusion cascades with individual retweet events crawled via the official API endpoint [16, 20, 38], the sampled-down effect from the API [33] has been usually overlooked and may lead to modeling bias. For instance, the sampled streaming API returns roughly 1% of all public tweets in real-time and the widely adopted filtered streaming API also rate limits on popular search queries.

Piecing together the missing events count. The missing event counts between two consecutive observed events can be derived by leveraging the ‘retweeted_count’ in tweet metadata. This field is included in every retweet and indicates how many times the original tweet has been retweeted (including by the current retweet). Fig. 1 depicts an example Twitter cascade with missing tweets as dashed lollipops. Four tweets are observed at times t_0, t_1, t_2 and t_3 . rtc_i indicates the ‘retweeted_count’ value in the i -th observed tweet. We compute c_0, c_1, c_2, \dots as the number of missing retweets between two observed tweets. This leads to a cascade that mixes both event counts and individual events.

A mixed data format. We propose the triple representation (o_i, d_i, c_i) for the i^{th} observation, where o_i is the starting observation time, d_i is the observation duration and c_i is the number of events (observed or missing) within the time segment $(o_i, o_i + d_i)$ – note, excluding interval ends. Using this notation, we denote an observed event at t_i as $(t_i - dt, 2dt, 1)$ where dt is an infinitesimal time period. We denote the missing events between t_i and t_{i+1} as

$(t_i, t_{i+1} - t_i, c_i)$. Note that c_i would correspond to $C(t_i, t_{i+1})$ in MBP’s interval-censored notation in Section 2.2. Finally, we can denote the reshare cascade as $\mathcal{H}^* = \{(t_i, d_i, c_i) \mid i = 0, \dots, m, o_i + d_i \leq T\}$ given a maximum observation time T and m the number of observed timestamps. For instance, one can denote the cascade in Fig. 1 as $= \{(t_0 - dt, 2dt, 1), (t_0, t_1 - t_0, 0), (t_1 - dt, 2dt, 1), (t_1, t_2 - t_1, 1), \dots\}$.

3.2 Interval-censored Transformer Hawkes

In this section, we propose the Interval-censored Transformer Hawkes (IC-TH) that generalizes the Transformer Hawkes accounting for event times and missing event counts. IC-TH takes data in the interval-censored event sequences format proposed in Section 3.1, and produces the hidden states using the self-attention mechanism proposed by Zuo et al. [41].

Model event intensity. Due to missing event timestamps (only their counts are available), IC-TH models reshare cascades using the expected conditional event intensity. We adopt the notations of Rizoiu et al. [25] and denote the expected event intensity function as

$$\xi(t \mid \mathcal{H}^*) = \mathbb{E}_{\mathcal{H}_u^*} [\lambda(t)] = f(\mathbf{w}^\top \mathbf{h}(t)), \quad (13)$$

where $\mathcal{H}_u^* = \{(t_i, d_i, c_i) \in \mathcal{H}^* \mid d_i > 2dt, c_i > 0\} \subseteq \mathcal{H}^*$ is the subset of triples denoting missing event counts. The intensity function is modeled from the hidden state $\mathbf{h}(t)$ via a softplus function $f(\cdot)$ and a weight matrix \mathbf{w} .

We note that when $\mathcal{H}_u^* = \emptyset$ or $\mathcal{H}^* = \mathcal{H}_u^*$, $\xi(t \mid \mathcal{H}^*)$ degrades to $\lambda(t \mid \mathcal{H}^*)$ or $\xi(t)$ discussed in Section 2, respectively.

Encoding \mathcal{H}^* . IC-TH employs the same temporal encoding procedure as TH to convert the timestamp t_i into an input vector X_i' . However, unlike the timestamps, the duration and event count information needs extra encoding layers before being fed to the self-attention module. Inspired by Li et al. [17], we encode the duration d_i and the event count c_i as masks to X_i' . Specifically, for d_i , we compute the context vector c_i^d using a fully connected layer f_θ , i.e., $c_i^d = f_\theta(\log(d_i))$. The duration mask m_i^d is then obtained via a linear transformation parameterized by \mathbf{W}_d , \mathbf{b}_d and a sigmoid function σ , i.e., $m_i^d = \sigma(c_i^d \mathbf{W}_d + \mathbf{b}_d)$. Similarly, the event count mask m_i^c is computed following the same layers with a different parameter set. Finally, $X_i = X_i' \odot m_i^d \odot m_i^c$, where \odot is an element-wise multiplication.

Log-likelihood function. Starting from Eq. (11), we compute the log-likelihood function for the mixings of event counts and

individual events.

$$\begin{aligned} \mathcal{L}_{\text{IC-TH-LL}}(\theta) &= \sum_{i \in \mathcal{H}^*} c_i \log \Xi(o_i, o_i + d_i) - \sum_{i \in \mathcal{H}^*} \Xi(o_i, o_i + d_i) \\ &= \underbrace{\sum_{i \in \mathcal{H}_u^*} c_i \log \Xi(t_i, t_{i+1})}_{\text{missing event counts}} + \underbrace{\sum_{i \in \mathcal{H}_c^*} \log \Xi(t_i - dt, t_i + dt)}_{\text{observed event times}} \\ &\quad - \sum_{i \in \mathcal{H}^*} \Xi(t_i, t_{i+1}) \\ &= \sum_{i \in \mathcal{H}_u^*} c_i \log \Xi(t_i, t_{i+1}) + \sum_{i \in \mathcal{H}_c^*} \log \xi(t_i) 2dt - \sum_{i \in \mathcal{H}^*} \Xi(t_i, t_{i+1}) \\ &= \sum_{i \in \mathcal{H}_u^*} c_i \log \Xi(t_i, t_{i+1}) + \sum_{i \in \mathcal{H}_c^*} \log \xi(t_i) - \sum_{i \in \mathcal{H}^*} \Xi(t_i, t_{i+1}) \\ &\quad + \sum_{i \in \mathcal{H}_c^*} \log 2dt, \end{aligned} \quad (14)$$

where $\mathcal{H}_c^* = \{(t_i, d_i, c_i) \in \mathcal{H}^* \mid d_i = 2dt, c_i = 1\}$ is the subset of triples denoting observed event times. This is equivalent to

$$\begin{aligned} \mathcal{L}_{\text{IC-TH-LL}}(\theta) &= \sum_{i \in \mathcal{H}_u^*} c_i \log \Xi(t_i, t_{i+1}) + \sum_{i \in \mathcal{H}_c^*} \log \xi(t_i) - \sum_{i \in \mathcal{H}^*} \Xi(t_i, t_{i+1}). \end{aligned} \quad (15)$$

The Linformer trick. The quadratic complexity of the transformer architecture in memory restricts the encoding of long cascades and large cascade groups. Therefore, we employ the Linformer trick, which optimizes the model to linear complexity [28]. Specifically, the *heads* in Linformer are computed as

$$\text{head}_i = \text{Softmax} \left(\frac{\mathbf{XW}_i^Q (\mathbf{E}_i \mathbf{XW}_i^K)^\top}{\sqrt{d_k}} \right) \mathbf{F}_i \mathbf{XW}_i^V \quad (16)$$

where $\mathbf{E}_i, \mathbf{F}_i \in \mathbb{R}^{k \times n}$ are two projection matrices that simplify the computation via $k \ll n$.

Implementation details. Throughout the paper, we conduct our experiments on a machine with an Intel Xeon CPU @ 2.20GHz processor and with Nvidia Tesla Volta V100 GPUs (32 GB memory). The proposed model is implemented in PyTorch 1.9.0 [22].

3.3 Model pre-training

Intuition. Here, we build a contrastive representation learning pipeline to obtain informative embeddings from groups of unlabeled diffusion cascades. The intuition is to leverage large amounts of unlabeled (and possibly unrelated to the learning task) diffusion cascades from the Twitter API.

We group the cascades into discrete groups based on a specific criterion — such as all cascades in which a given user participates or cascades related to a given online item (such as a YouTube video). We build the representation for a group by aggregating the cascade embeddings for cascades of that group. The contrastive representation learning algorithm aims to maximize latent distances between a group representation and other groups and minimize its latent distance to its own cascades.

Technical detail. We first augment the dataset by building positive and negative pairs of cascade groups via downsampling. Specifically, given a group of cascades, $\mathcal{C}_i = \{\mathcal{H}_1, \mathcal{H}_2, \dots\}$ posted

Table 1: Statistics of the datasets.

	ACTIVERT2016	ACTIVERT2017	RNCNIX	IO
#items	155, 105, 987 videos	75, 717 videos	102, 429 articles	32, 486 users
#cascades	881, 587, 021	30, 535, 891	8, 129, 126	19, 476, 766
#tweets	1, 212, 945, 195	85, 334, 424	56, 397, 252	22, 845, 053

by the same user or related to the same online item denoted as i , we randomly split the cascade group into two, $C_{i,1}$ and $C_{i,2}$, and consider them as a positive pair. Naturally, examples formed by pairing cascade groups related to different items or users are deemed negative.

We then obtain the cascade embeddings $h_{i,1}$ and $h_{i,2}$ with the proposed model. Following prior works [34], we further introduce a MLP-based projection head to a new representation, i.e., $z_{i,j} = MLP(h_{i,j})$. Finally, we optimize the contrastive loss function [8]:

$$\mathcal{L}_i^{contrastive} = -\log \frac{\exp(f_{sim}(z_{i,1}, z_{i,2})/\tau)}{\sum_{k=1}^N \sum_{l=1}^2 \sum_{j=1}^2 \exp(f_{sim}(z_{i,l}, z_{k,j})/\tau)}, \quad (17)$$

where $f_{sim}(z_1, z_2)$ is the cosine similarity between z_1 and z_2 .

4 EXPERIMENTS AND RESULTS

In this section, we present the quantitative performance results for IC-TH. First, we present the four datasets and learning tasks (Section 4.1). Next, we study the robustness of IC-TH concerning data loss (Section 4.2). Finally, we evaluate IC-TH against the state-of-the-art in two predictive tasks: popularity prediction (Section 4.3) and online item categorization (Section 4.4).

4.1 Datasets

This work employs four real-world Twitter datasets. Two datasets originate from prior works (ACTIVERT2017 and RNCNIX). We contribute two new datasets: one for model pre-training (ACTIVERT2016) and an Information Operations (IO) dataset based on the unhashed datasets provided by Twitter Moderation Research Consortium [1].

ACTIVERT2017 [15] was constructed by Kong et al. [15] by collecting tweets mentioning popular Youtube videos. Cascade groups are constructed based on the YouTube video ID: all the cascades mentioning the same video will be placed in the same cascade group. The categorization task for this dataset is predicting the video’s category (Entertainment, Gaming, Music and News&Politics, etc.) based on retweet cascade dynamics. This dataset suffers from Twitter API’s down-sampling, given the popularity of YouTube videos on Twitter (see Wu et al. [33]). We augment each retweet cascade with the missing event counts between consecutively recorded retweets (as described in Section 3.1).

ACTIVERT2016. We contribute a dataset by following the same setup as ACTIVERT2017 using tweets from 2016. We employ ACTIVERT2016 to pretrain the IC-TH to obtain a better representation of cascade groups. As we use the Twitter API v2 for collecting this dataset, the cascades are complete. To make ACTIVERT2016 comparable with ACTIVERT2017, we sample cascades at the same level as ACTIVERT2017.

RNCNIX [6]. The tweets in this dataset were collected by querying the Twitter search endpoint for tweets mentioning articles from

a list of controversial news publishers and a list of leading Australian news outlets [4–6]. The cascade groups were constructed based on the news publishers, i.e., all cascades relating to news articles from the same publisher are in the same group. The categorization task for this dataset is predicting whether a news publisher is reputable or controversial based on retweet cascade dynamics. The data is obtained retrospectively using a paid service, and the cascades are complete without the sampling-down effect.

Information Operations (IO) dataset contains the users (also known as IO operatives) and their tweets linked to state-backed information operations identified and released officially by Twitter Moderation Research Consortium [1]. We obtained the unhashed version of this dataset in Nov 2021; the dataset spans from Nov 4, 2010, to Aug 21, 2020. We select IO operatives from three countries (Russia, Iran and Saudi Arabia) with the most significant number of users and highest activity.

We augment the dataset with a matched set of organic Twitter users, i.e., users who were involved in the discussions with the IO operatives but were not identified by Twitter as operatives. Building the organic set presents challenges, as when Twitter suspends an IO operative account, it also deletes all its activity, including the retweets emitted by organic users. We cannot use the Twitter API to search for the activity of IO operatives (as all of their traces have been removed), and the dataset does not contain organic activity. Instead, we cross-link with a 1% sampled tweet stream. We follow the following steps:

- (1) We first collect the IDs of tweets posted or retweeted by removed users.
- (2) We augment the tweets by searching for the retweet cascades using these tweet IDs in a complete archive of 1% sampled tweet stream¹ hosted at archive.org².
- (3) We consider users who participated in these retweet cascades but were not removed by Twitter as organic users. To obtain users with similar popularity levels, we sample 10,000 organic users by matching the follower count distribution of removed users.
- (4) Last, we repeat the first two steps for organic users to augment the dataset with all retweets cascades engaged by them.

We construct cascade groups based on the users, i.e., all the cascades in which a given user participates are placed in the same group. The categorization task is predicting whether a user is an IO operative from one of the three countries or an organic user (4 classes). Table 1 summarizes the basic statistics of the four datasets.

4.2 Effect of data loss on IC-TH

In this section, we evaluate the robustness of IC-TH’s learned group embeddings concerning missing data. We deploy a synthetic setup in which cascades are synthetically generated and grouped into cascade groups. Events are randomly removed with given probabilities. Finally, we generate group embeddings using the contrastive learning loss and evaluate the separability based on the group type.

¹<https://developer.twitter.com/en/docs/twitter-api/tweets/volume-streams/introduction>

²<https://archive.org/details/twitterstream>

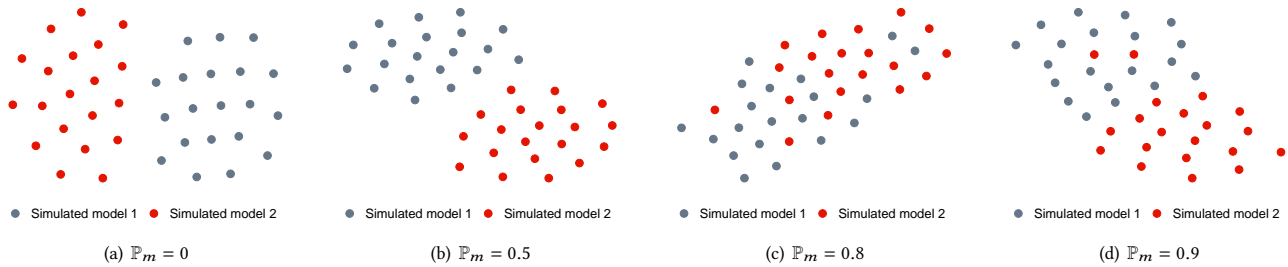


Figure 2: t-SNE [18] visualizations of the group embeddings of synthetic cascades obtained using IC-TH. Each circle represents a group of cascades. The events in each cascade are removed with the probability \mathbb{P}_m , varying from no missing data ($\mathbb{P}_m = 0$, (a)) to 90% of the events missing ($\mathbb{P}_m = 0.9$, (d)). Apart from down-sampling, the groups and the cascades are identical across (a)-(d).




Figure 3: Performance of predicting final popularity measured using the absolute percentage error (APE) on ACTIVER2017—lower is better. We compare IC-TH (with and without pre-training, in bold font) against three state-of-the-art baselines: HawkesN [16], MBP [25] and Transformer Hawkes (TH) [41]. The boxplots summarize all the cascades in the dataset. The red diamond shows the mean APE.

Setup. We sample the cascades from two Hawkes process models — a power-law decay kernel and an exponential decay kernel, with randomly selected parameters. We sample 10,000 cascades for each model, split into 20 groups of 500 cascades. This results in 40 groups of cascades. We synthetically emulate Twitter’s sample-down mechanism for each cascade (see Wu et al. [33] for an in-depth discussion of Twitter’s sample-down mechanism). We choose a sampling-down probability \mathbb{P}_m , and we remove the tweets in each cascade independently and randomly with the probability \mathbb{P}_m . Once an event is marked as missing, we remove it from its cascade and update the event count in the time interval. We conduct contrastive learning on the simulated cascade groups and apply a widely employed dimension reduction tool, t-SNE [18], on the group embeddings for visualization. We repeat the experiments with varying \mathbb{P}_m missing probability to examine the effect of different levels of data loss.

Fig. 2 depicts the embeddings of cascade groups produced from the trained IC-TH, for the various \mathbb{P}_m . When there is no missing data ($\mathbb{P}_m = 0$) or half of the tweets are missing ($\mathbb{P}_m = 0.5$), the color groups are perfectly separated. As the data loss increases ($\mathbb{P}_m > 0.5$), the color groups start mixing together. However, we achieve near-perfect separability even at very large data sampling rates ($\mathbb{P}_m = 0.9$ in Fig. 2(d)). This shows that IC-TH captures the differences in cascade groups even as most events are missing and represented as missing event counts.

Table 2: Macro-F1 of predicting the category of YouTube videos (on ACTIVER2017), controversial news publishers (RNCNIX) and state-backed Information Operations operatives (IO). We compare three flavors of IC-TH against two state-of-the-art models. Higher is better.

Models	ACTIVER2017	RNCNIX	IO
DMM [15]	0.488	0.675	0.968
TH [41]	0.469	0.823	0.983
IC-TH w/o missing counts	0.495	0.840	0.985
IC-TH	0.499	-	0.987
Pre-trained IC-TH	0.503	0.853	0.987

4.3 Future popularity prediction

In this section, we evaluate IC-TH for the cascade future popularity prediction task on the ACTIVER2017 dataset. We evaluate the impact of pre-training for this task by testing IC-TH with and without pre-trained weights on the ACTIVER2016 dataset.

Setup and baselines. We compare IC-TH against the baseline models discussed in Section 2. This includes an event count model (MBP [25]) and two event time models (the generalized HawkesN [16] and Transformer Hawkes [41]). As the ACTIVER2017 is a sampled-down dataset, it contains both retweet counts and retweet events. We process the data to adapt it to the requirements of each baseline as follows. For event time models, we model only the observed event times in a cascade. For the event count model (i.e., MBP), we consider the observed event timestamps as observation times and compute the total event counts for each observation period — e.g., $c_1 + 1$ for the observation period t_0 to t_1 . These operations introduce information loss; however, we highlight that we are unaware of any other model that can leverage event count and individual event data simultaneously.

Results. Fig. 3 shows the results of the popularity prediction. Visibly, all transformer-based approaches (IC-TH and TH) provide a significant performance boost over the generative approaches (HawkesN and MBP). Furthermore, our proposed IC-TH outperforms TH in terms of both mean and median Absolute Percentage Error (APE). Surprisingly, unlike the categorization prediction (see next section), pre-training does not boost performance. The pre-trained IC-TH shows performances on par with TH.




Figure 4: Comparison of prediction performances on ACTIVEVERT2017 of IC-TH with and without pre-trained weights after being trained on different proportions of training data. Mean F1 scores and standard deviations among 10 runs are shown as colored lines and areas, respectively.

4.4 Categorical prediction of online items

Here, we evaluate the categorization of online items outside Twitter (video categories, publisher reputation, IO state-backed agents) using solely the retweet cascades temporal information.

Setup. We follow the experimental setup outlined by Kong et al. [15]. For ACTIVEVERT2017 we predict the category of YouTube videos among four video categories (Entertainment, Gaming, Music and News&Politics). For RNCNIX, we predict whether a publisher is controversial or reputable. For IO, we distinguish among four user state-backed agent identities (Russia-, Iran-, Saudi Arabia-backed and Organic User). We evaluate the prediction performance using the macro-F1 with the same 50%-50% train-test split setup as [15] and employing 5% of the training data for validation. We compare our proposed IC-TH (with and without pre-trained weights) to state-of-the-art baselines TH [41] and DMM [15]. We also evaluate a version of IC-TH where retweet count data is excluded from the input – i.e., we remove the missing event counts and show only the event times – to evaluate the contribution of knowing how many events are missing. Note the dash for IC-TH on RNCNIX, as this dataset does not have any missing data; therefore IC-TH is identical to IC-TH w/o missing counts.

Results. We show the category prediction results in Table 2. We make three observations. First, the classification task is more straightforward on the datasets RNCNIX and IO than on ACTIVEVERT2017, as all algorithms obtain significantly higher macro-F1 scores. Second, our pre-trained IC-TH outperforms all other baselines and IC-TH flavors. Third, we estimate the performance impact of the three modeling contributions of pre-trained IC-TH over TH. The new data representation introduced in Section 3.1 appears to have the highest impact (IC-TH w/o missing counts compared to TH). Introducing missing counts and pre-training both lead to moderate performance increases.

Effect of pre-training weights. To explore the benefits of pre-trained weights, we test the IC-TH model (with and without pre-training) on training sets of ACTIVEVERT2017 of varying sizes (in percentage). We repeat the train-test procedure 10 times for each training set size and report the mean and standard deviation of Macro-F1. All pre-training is performed on ACTIVEVERT2016. Fig. 4 depicts the performance increase with the size of the training dataset. We see that the improvements from the pre-trained

Table 3: Mean \pm standard deviation of Jaccard similarity in hashtag usage between pairs of users in different clusters (inter-cluster line) and within the same cluster (intra-cluster line) for users affiliated with each country. Higher scores indicate more similarity.

	Russia ($k = 3$)	Saudi Arabia ($k = 3$)	Iran ($k = 4$)
Inter-cluster sim.	0.064 \pm 0.0183	0.038 \pm 0.0467	0.034 \pm 0.0171
Intra-cluster sim.	0.133 \pm 0.0289	0.260 \pm 0.2390	0.089 \pm 0.0405

weights become notable and consistent when using more than 40% of training data.

5 ANALYZING STATE-BACKED IO AGENTS

In this section, we perform an in-depth quantitative (Section 5.1) and qualitative analysis (Section 5.2) of the users in the IO dataset, studying both the temporal dynamics and the content of the tweets.

5.1 Identify clusters of state-backed accounts

Method. We leverage IC-TH fine-tuned on ACTIVEVERT2016 to derive cascade group embeddings for users – a cascade group corresponds to each user and contains all the retweet cascades in which the user is involved. We visualize in Fig. 5(a) the user embeddings from IO with t-SNE. The figure shows a clear separation between the removed state-backed agents from the three countries and the organic users who were not removed. This is not surprising given the macro-F1 score of 0.987 obtained by the pre-trained IC-TH on IO (see Table 2). However, clear divisions also exist within each country group – e.g., at least three clusters can be observed for Russia-linked users. Next, we investigate whether the users in the same clusters have more in common than temporal tweeting patterns.

Assess content similarity of clusters. We first quantitatively measure the pairwise similarities of users within clusters. For state-backed removed users, we apply the k-means clustering on the t-SNE space with the k number manually selected to match the cluster numbers in Fig. 5(a): 3 for Russia, 3 for Saudi Arabia and 4 for Iran. We then build the sets of hashtags used by each user and compute the Jaccard similarity scores between all pairs of users. We show in Table 3 the average similarity scores and standard deviations of inter-cluster and intra-cluster user pairs. The scores consistently indicate higher similarity among users within the same cluster for all countries. This indicates that users placed in the same cluster based on their temporal tweeting patterns (captured by IC-TH) exhibit similar hashtag usage. In other words, the identified clusters appear coherent both temporally and content-wise.

5.2 Investigation of Russia-backed Information Operations strategies

Here we analyze the agenda-setting episodes for each cluster and investigate for signs of possible coordination across clusters by checking whether they post similar messages with increased frequency on the same day.

First, we map the temporal posting patterns of the identified Russia-backed clusters (i.e., number of tweets per day) to determine




Figure 5: Understanding the strategies of state-backed agents. (a) T-SNE visualizations of the users in the IO dataset. The colors indicate the countries of origin for state-backed agents, as identified by Twitter. Three clusters of Russian-backed users are highlighted. (b) Temporal patterns of coordinated activity across the clusters of Russian-backed users. Annotations show the most used messages and hashtags on particular days.

the days of increased activity within each cluster. We identify the “spikes” — the days characterized by increased frequency of tweets from the users in each particular cluster — and sample and analyze their tweets’ content at peak times. For example, in Fig. 5(b), we observe a distinct spike in activity among *Cluster 2* and *Cluster 3* on 18-19 July 2014. On 17 July 2014, Malaysia Airlines flight MH17 was shot down using an anti-aircraft system belonging to the Russian military, killing all 298 passengers and crew on board. We see from the IO tweets data that on the next day (18 July), users from both *Clusters 2 and 3* released an avalanche of accusatory tweets in an attempt to shift blame towards Ukraine. The content of the Russian-language tweets included hashtags such as #KievShotDown-theBoeing, #KievProvocation and #KievTelltheTruth. The narratives alleged Ukraine’s responsibility for committing the attack. They also drew conclusions regarding Ukraine’s future as an independent state. These narratives were aligned with Russia’s strategic geopolitical interests — shifting the blame to Ukraine and delegitimizing the Ukrainian government allowed them to justify their military intervention in the region.

Interestingly, while we can observe instances of coordinated activity across clusters like the one above, we can also see that each cluster exhibits unique information-spreading patterns. After the downing of MH17, *Cluster 2* continued tweeting Russian news and messages of support for the Russian government in the Russian language. In contrast, users in *Cluster 3* have switched to English in an apparent effort to target a new audience on the platform. Some of these accounts posted inspirational quotes and engaged in popular culture discussions marked by common hashtags such as #music and #usa. Others engaged in conversations with other users, providing relationship advice and even consolation. Then on 11/09/2014, all of these accounts started pushing information using the hashtag #columbianchemicals — a known hoax claiming an explosion at a chemical plant in Centerville, Louisiana in 2014 [10]. #columbianchemicals was not the only issue discussed in the cluster — other parts were dedicated to spreading false narratives about Ukraine, including negative portrayals of Ukraine’s Revolution of Dignity and snapshots of a happy civilian life in the

Russian-occupied Crimea. Another slightly smaller fraction of these conversations was dedicated to discussing the Russian soul and the “indomitable Russian spirit” that were often linked to Christianity and Russia’s much-celebrated victory in World War II. Thus, *Cluster 3* captures a broad repertoire of Russia’s attempts at foreign and domestic influence perpetrated by seemingly “ordinary” accounts.

Unlike *Cluster 3*, the tweets in *Clusters 1 and 2* were exclusively in the Russian language. *Cluster 1* contains retweets of news reported by the Russian mainstream media. Most of the tweets from this cluster have a notably patriotic framing, praising the Russian president Vladimir Putin or Russia’s achievements in the automotive industry and sports. On certain days, users from this cluster would participate in coordinated campaigns, such as the #MadeInRussia campaign that aimed to justify Russia’s import substitution policy. Following the official announcement, these accounts shared information on various Russian-made goods and praised their advantages compared to their sanctioned analogs.

Cluster 2 is thematically similar to *Cluster 1* but has a distinct prevalence of regional (rather than national) news with a mix of conservative sentiment. It has been activated to promote pro-Russian hashtags that mimic public sentiments, such as #LavrovBeStrong on 21/03/2015 or #TurkeyAggressor on 24/11/2015.

Overall, the temporal analysis of coordinated behavior within clusters suggests they might represent separate organizational units, such as troll farms. This is a plausible scenario in which each organizational unit/troll farm would be engaged in a separate set of activities that constitute Russia’s information operations for domestic or international audiences. However, as we can observe in Fig. 5(b), in the case of significant political events, they can also join their forces and participate in coordinated campaigns across clusters.

6 CONCLUSION

This paper proposes the Interval-censored Transformer Hawkes (IC-IH), an extension to the Transformer Hawkes model that accounts for both observer event times and counts of missing events between observed events. We start by observing that most Twitter

feeds are sampled-down, and tweets are missing at varying rates. We show that the number of missing tweets can be reconstructed using the ‘*retweeted_count*’ information embedded in each retweet. We propose a novel data embedding scheme that accounts for both observed and missing data, a new log-likelihood function for training IC-TH and a contrastive learning approach for training it with large volumes of unlabeled data.

We test IC-TH on three large datasets with two learning tasks: future popularity prediction and online item categorization. We show that IC-TH outperforms the current state-of-the-art baselines on all datasets and tasks: generative and neural network-based, interval-censored, and event time-based. Most importantly, we show that the category of YouTube videos, the reputability of news publishers and the status of state-backed actors for Twitter users can be predicted using solely the timing of the retweet cascades associated with these items.

IC-TH for IO detection. As validation suggests, grouping users solely based on their temporal interaction patterns represented by retweet cascades allows researchers to identify clusters of users engaged in information operations. Moreover, the clusters’ granularity might be sufficient in many cases to identify each distinct information operation and its textual and hypertextual features. Based on these preliminary findings, the interval-censored Transformer Hawkes might be used to identify coordinated inauthentic behavior in other datasets independent of their content, making it a potentially helpful tool for combatting foreign interference on social technology platforms.

Acknowledgments

This work was supported by the Commonwealth of Australia (represented by the Defence Science and Technology Group) through a Defence Science Partnerships Agreement. This research used resources and services from the National Computational Infrastructure (NCI), which is supported by the Australian Government.

REFERENCES

- [1] 2022. Twitter Moderation Research Consortium. <https://transparency.twitter.com/en/reports/moderation-research.html>
- [2] Aseel Addawood, Adam Badawy, Kristina Lerman, and Emilio Ferrara. 2019. Linguistic cues to deception: Identifying political trolls on social media. In *ICWSM*.
- [3] Olga Boichak, Jeff Hemsley, Sam Jackson, Rebekah Tromble, and Sikana Tanupabrungsun. 2021. Not the Bots You Are Looking For: Patterns and Effects of Orchestrated Interventions in the U.S. and German Elections. *International Journal of Communication* (2021).
- [4] Axel Bruns. 2016. Big Data Analysis. *The Sage handbook of digital journalism* (2016).
- [5] Axel Bruns. 2017. Making Audience Engagement Visible: Publics for Journalism on Social Media Platforms. *The Routledge Companion to Digital Journalism Studies* (2017).
- [6] Axel Bruns and Tobias Keller. 2020. News diffusion on Twitter: Comparing the dissemination careers for mainstream and marginal news. In *Social Media & Society 2020 Conference*.
- [7] Pio Calderon, Alexander Soen, and Marian-Andrei Rizoiu. 2021. Linking Across Data Granularity: Fitting Multivariate Hawkes Processes to Partially Interval-censored Data. (2021).
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *ICML*.
- [9] Daryl J Daley and David Vere-Jones. 2008. Conditional Intensities and Likelihoods. In *An introduction to the theory of point processes*. Vol. I. Springer, Chapter 7.2.
- [10] Aaron Delwiche. 2019. Computational Propaganda and the Rise of the Fake Audience. *The SAGE Handbook of Propaganda* (2019), 105.
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv* (2018).
- [12] 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 *KDD*.
- [13] Alan G Hawkes. 1971. Spectra of some self-exciting and mutually exciting point processes. *Biometrika* (1971).
- [14] Matthias Kirchner. 2016. Hawkes and INAR(∞) processes. *Stochastic Processes and their Applications* (2016).
- [15] Quyu Kong, Marian-Andrei Rizoiu, and Lexing Xie. 2020. Describing and predicting online items with reshare cascades via dual mixture self-exciting processes. In *CIKM*.
- [16] Quyu Kong, Marian-Andrei Rizoiu, and Lexing Xie. 2020. Modeling Information Cascades with Self-exciting Processes via Generalized Epidemic Models. In *WSDM*.
- [17] Yang Li, Nan Du, and Samy Bengio. 2018. Time-dependent representation for neural event sequence prediction. *ICLR Workshop track* (2018).
- [18] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *JMLR* (2008).
- [19] Hongyuan Mei and Jason M Eisner. 2017. The neural Hawkes process: A neurally self-modulating multivariate point process. In *NeurIPS*.
- [20] Swapnil Mishra, Marian-Andrei Rizoiu, and Lexing Xie. 2016. Feature Driven and Point Process Approaches for Popularity Prediction. In *CIKM*.
- [21] Swapnil Mishra, Marian-Andrei Rizoiu, and Lexing Xie. 2018. Modeling Popularity in Asynchronous Social Media Streams with Recurrent Neural Networks. In *ICWSM*.
- [22] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *NeurIPS*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett (Eds.). Curran Associates, Inc.
- [23] Rohit Ram and Marian-Andrei Rizoiu. 2022. Data-driven ideology detection: a case study of far-right extremist. In *Defence Human Sciences Symposium*. Sydney, Australia, 1–2.
- [24] Marian-Andrei Rizoiu, Swapnil Mishra, Quyu Kong, Mark Carman, and Lexing Xie. 2018. SIR-Hawkes: on the Relationship Between Epidemic Models and Hawkes Point Processes. In *WWW*.
- [25] Marian-Andrei Rizoiu, Alexander Soen, Shidi Li, Pio Calderon, Leanne Dong, Aditya Krishna Menon, and Lexing Xie. 2022. Interval-censored Hawkes processes. *JMLR* (2022).
- [26] Karishma Sharma, Yizhou Zhang, Emilio Ferrara, and Yan Liu. 2021. Identifying coordinated accounts on social media through hidden influence and group behaviours. In *KDD*.
- [27] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *NeurIPS* (2017).
- [28] Sinong Wang, Belinda Z Li, Madian Khabsa, Han Fang, and Hao Ma. 2020. Linformer: Self-attention with linear complexity. *arXiv* (2020).
- [29] Claire Wardle and Hossein Derakhshan. 2017. Information disorder: Toward an interdisciplinary framework for research and policymaking.
- [30] Samuel C Woolley and Philip N Howard. 2018. Introduction: Computational propaganda worldwide. *Computational propaganda: Political parties, politicians, and political manipulation on social media* (2018), 3–18.
- [31] Siqi Wu, Marian-Andrei Rizoiu, and Lexing Xie. 2018. Beyond views: Measuring and predicting engagement in online videos. In *ICWSM*.
- [32] Siqi Wu, Marian-Andrei Rizoiu, and Lexing Xie. 2019. Estimating Attention Flow in Online Video Networks. *Proceedings of the ACM on Human-Computer Interaction CSCW* (2019).
- [33] Siqi Wu, Marian-Andrei Rizoiu, and Lexing Xie. 2020. Variation across scales: Measurement fidelity under twitter data sampling. In *ICWSM*.
- [34] Xovee Xu, Fan Zhou, Kunpeng Zhang, and Siyuan Liu. 2022. CCGL: Contrastive Cascade Graph Learning. *TKDE* (2022).
- [35] 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 *WWW Companion*.
- [36] Qiang Zhang, Aldo Lipani, Omer Kirnap, and Emine Yilmaz. 2020. Self-attentive Hawkes process. In *ICML*.
- [37] Rui Zhang, Christian Walder, Marian-Andrei Rizoiu, and Lexing Xie. 2019. Efficient non-parametric Bayesian Hawkes processes. *IJCAI* (2019).
- [38] Qingyuan Zhao, Murat A. Erdogdu, Hera Y. He, Anand Rajaraman, and Jure Leskovec. 2015. SEISMIC: A Self-Exciting Point Process Model for Predicting Tweet Popularity. In *KDD*.
- [39] Feng Zhou, Quyu Kong, Zhijie Deng, Jichao Kan, Yixuan Zhang, Cheng Feng, and Jun Zhu. 2022. Efficient Inference for Dynamic Flexible Interactions of Neural Populations. *JMLR* (2022).
- [40] Ke Zhou, Hongyuan Zha, and Le Song. 2013. Learning social infectivity in sparse low-rank networks using multi-dimensional Hawkes processes. In *AISTATS*.
- [41] Simiao Zuo, Haoming Jiang, Zichong Li, Tuo Zhao, and Hongyuan Zha. 2020. Transformer Hawkes process. In *ICML*.