



HAL
open science

Uncovering Social Network Activity Using Joint User and Topic Interaction

Gaspard Abel, Argyris Kalogeratos, Jean-Pierre Nadal, Julien Randon-Furling

► **To cite this version:**

Gaspard Abel, Argyris Kalogeratos, Jean-Pierre Nadal, Julien Randon-Furling. Uncovering Social Network Activity Using Joint User and Topic Interaction. IEEE Transactions on Computational Social Systems, 2026, <10.1109/TCSS.2026.3664984>. <hal-05411622>

HAL Id: hal-05411622

<https://hal.science/hal-05411622v1>

Submitted on 12 Dec 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Uncovering Social Network Activity Using Joint User and Topic Interaction

Gaspard Abel^{*×}, Argyris Kalogeratos^{*}, Jean-Pierre Nadal^{×†}, Julien Randon-Furling^{*+}

Abstract—The emergence of online social platforms, such as social networks and social media, has drastically affected the way people apprehend the information flows to which they are exposed. In such platforms, various information cascades spreading among users is the main force creating complex dynamics of opinion formation, each user being characterized by their own behavior adoption mechanism. Moreover, the spread of multiple pieces of information or beliefs in a networked population is rarely uncorrelated. In this paper, we introduce the *Mixture of Interacting Cascades* (MIC), a model of marked multidimensional Hawkes processes with the capacity to model jointly non-trivial interaction between cascades and users. We emphasize on the interplay between information cascades and user activity, and use a mixture of temporal point processes to build a coupled user/cascade point process model. Experiments on synthetic and real data highlight the benefits of this approach and demonstrate that MIC achieves superior performance to existing methods in modeling the spread of information cascades. Finally, we demonstrate how MIC can provide, through its learned parameters, insightful bi-layered visualizations of real social network activity data.

Index Terms—Hawkes processes, information diffusion, opinion formation, topic interaction, online social networks.



1 Introduction

THE advent of wireless communication technologies and online social platforms, coupled with the miniaturization and accessibility of mobile devices, has transformed the way individuals are exposed to information, as well as how they interact with each other in digital diffusion networks. The growing role of online social platforms, namely online social media and even more typically online social networks (OSNs) that we henceforth use as reference, in public debate and opinion formation has attracted research interest in the last decades, trying to understand how the information diffusion shapes the beliefs of a population. The way a diffusion process propagates in a population and impacts its agents is investigated in the field of collective phenomena, among other fields. The associated class of problems is relevant in various domains, such as epidemiology [1], [2], econometrics [3], viral diffusion [4], [5], or product recommendation [6], [7]. Additionally, studies on social network activity have considered an information cascades approach. This formulation is supported by the temporal property of information propagation [8] and the universal properties of information sharing [9], [10], [11].

When users share content in a social network, they express themselves on a specific topic and at the same time assist in spreading the associated information cascade to their vicinity.

Despite being initiated locally, this phenomenon may extend to a non-negligible fraction of the population, up to a point where consensus is reached. On another level, OSNs are also characterized by the simultaneous spread of a multitude of topics, which may be related in a way that affects user behavior as well. However, notably limited attention has been put to the coupling between information diffusion, involving multiple correlated cascades, and opinion formation. This brings us to the main motivation of this work, where we argue that considering non-trivial interactions between the topics of the information cascades is necessary for encompassing refined social network activity in real life. To this end, this work aims at providing conceptual tools and results that bridge the gap between theoretical and data-driven analysis.

A summary of the contribution of this article is as follows:

- We present the *Mixture of Interacting Cascades* (MIC) model, a two-layer temporal point process, namely multidimensional marked Hawkes processes, with the capacity to model complex patterns of social network activity generated by the interplay of users and cascades. More specifically, MIC models jointly cascade-to-cascade, cascade-to-user, and finally user-to-user interactions, where the interactions between cascades are implicit in the event generation process driven by users (see Fig. 1). To achieve that, MIC encodes the pairwise influence between cascades in a *cascade interaction matrix*. Moreover, MIC encompasses existing models in the literature, namely the *Independent Cascades* (IC) [12], [13] and *Correlated Cascades* (CC) [14] models, and its gains come with a reasonable complexity overhead, given that the cascading topics are usually substantially less in number than the users.
- We derive closed-form expressions for the conditional intensity and the number of events, which we validate with numerical simulations. Providing this type of theoretical results is not typical in the related literature.
- We provide a simple yet effective Maximum Likelihood Estimation scheme for learning the model parameters based

* *Université Paris Saclay, Université Paris Cité, ENS Paris Saclay, CNRS, SSA, INSERM, Centre Borelli, F-91190, Gif-sur-Yvette, France*

× *Centre d'Analyse et de Mathématique Sociales, EHESS, CNRS, 75006 Paris, France.*

† *Laboratoire de Physique de l'École Normale Supérieure, ENS, Université PSL, CNRS, Sorbonne Université, Université Paris Cité, F-75005 Paris, France*

+ *College of Computing, UM6P, Ben Guerir, Morocco.*

✉ *Corresponding authors with equal contribution: Gaspard Abel and Argyris Kalogeratos. E-mails: name.surname@ens-paris-saclay.fr.*

Manuscript received December XX, XXXX; revised XXXX XX, XXXX.

on input data, and we validate MIC's advantage in capturing the dynamics of multiple interacting cascades in synthetic experiments comparing to existing Hawkes-process-based models. By evaluating further the model with a series of experiments on real data, we show that MIC not only manages to capture more accurately information spreading dynamics, but is more adaptable to various types of datasets and evaluation measures. Moreover, results on real data highlight the inherent limitation of all the compared temporal point processes methods in handling the heterogeneity of volume of user (and cascade) activity. This panel of evaluation measures suggest that MIC is the best candidate for encompassing such properties.

■ We demonstrate that MIC can also offer insights on the structure of information pathways by visualizing social network activity, as shown in the schematic view of Fig. 1.

The rest is organized as follows: Sec. 2 gives an overview of the literature on the subject of information diffusion and on opinion formation in diffusion networks. Sec. 3 introduces the basic notation required to develop the theoretical framework in which the MIC model is grounded. Sec. 6 presents the derivation of closed-form expressions for the conditional intensity and the number of events for the MIC model, and validates these results with numerical experiments. Finally, in Sec. 7, the MIC model is evaluated with a series of experiments on real and synthetic data. Concluding remarks are issued in Sec. 8.

2 Related work

2.1 Information diffusion in online social networks

The spread of information in OSNs is studied at different granularity levels: from local user-based dynamics, to global network analysis aiming at predicting the popularity of contents. Considering the diffusion of content, *Information Cascades* is a popular formulation [15]. Although the network structure is assumed to be known, the complexity of these spreading processes lies in the heterogeneous social influence between users. In fact, learning the underlying influence weights of the associated graph has attracted significant attention: NetInf [16] and [4] propose optimization methods to learn the local infectivity of a graph that is assumed to be low rank and sparse, while [17] tackles the joint problem of diffusion network inference and meme tracking to highlight the importance of characterizing topics for the detection of viral content.

OSN dynamics are also investigated through the lens of influence estimation and dynamic graph evolution [18]. Unlike traditional broadcast media, the information cascades in OSNs are characterized by user interactions and content sharing. *Independent Cascades* (IC) [12] was among the first models for information spread in networks. More refined models were introduced [6], [7], [19], [20], [21], also the *Correlated Cascades* (CC) model [14] that include correlations between products and user influence, yet without encompassing the complex interplay between users and information cascades.

2.2 Opinion formation in social networks

Another important line of work studies the opinion evolution in social networks. Over the last decades, opinion dynamics and theoretical modeling are being also intensively researched, from the prototypical DeGroot model [22] to more complex

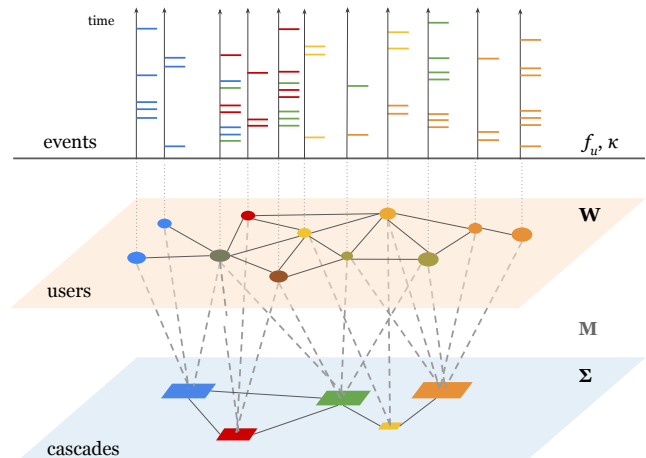


Fig. 1. Bi-layer scheme for the MIC model. Users interact on the top of a layer of interacting cascades. Node size is proportional to the associated volume of events, and node color depicts the mixture of cascades. MIC encompasses complex patterns of social network activity related to both layers and their interplay, namely by modeling jointly cascade-to-cascade (Σ), cascade-to-user (M), and user-to-user (W) interactions. Cascade(-to-cascade) interactions are implicit through the event generation process driven by users. Model parameters Σ, W, M , and the design components $f_u(\cdot), \kappa(\cdot)$ (see Sec. 4) appear on the right.

models that account for the influence of the network structure, the presence of stubborn agents [23], an acceptance threshold [24], [25], multiple topics [26], or the influence of a global steering mechanism [27]. Beyond theoretical modeling that makes the unrealistic assumption of having access to the true opinion of agents, studying opinion formation in OSNs can be enabled by estimating user opinions based on their observed online activity, e.g. via sentiment analysis [28], [29], or by looking at retweet dynamics [30], [31]. Notably, the SLANT model [32] comprises a set of marked jumped stochastic differential equations to capture both asynchronous event dynamics and user opinions, while in [33] an optimal control framework is developed. The limitation of these models is that they only consider that each user has a varying opinion with respect to a single topic, and not on multiple interacting topics that is the case in the content shared in real OSNs. In this context, the implicit interactions between topics take roots into their contextual links, as well as user bias [5].

2.3 Temporal point processes for social network activity

Recurrent temporal point processes have been used for capturing the patterns of information diffusion in social networks. In particular, starting from the original Hawkes processes [34], the extension of multidimensional Hawkes processes [35] has been considered for capturing microscopic information diffusion pathways [6], [7], [18]. This formulation is convenient for accounting for both social and self-exciting collective phenomena, as well for considering the multitude of cascades propagating through an OSN. There has been a growing interest in efficient parameter estimation methods for temporal point processes [4], [17], [21], [36], as well as in hybrid variations that combine deep learning architectures [37]. However, far less effort has been invested in providing comprehensive modeling frameworks to tackle analytically and numerically complex information diffusion patterns at a microscopic scale.

3 Background on Hawkes processes

Preliminaries. Let a directed graph $\mathcal{G} = (\mathcal{V}, E)$ representing a network, with \mathcal{V} being a set of N_u nodes (users) and E a set of edges, together inducing an adjacency matrix \mathbf{A} . Let $\mathbf{W} = \{w_{vu}\}_{u,v}$ be the associated weight matrix expressing with a non-zero value the *social influence* from user v to user u , for each edge $(v, u) \in E$. Let also \mathcal{F}_u be the set of ‘followers’ of u , i.e. nodes with directed edges to u . Suppose a set $\mathcal{C} = \{c_i\}_{i=1}^{N_c}$ of *information cascades* (tweets/posts that represent topics, opinions, products). User activity is a sequence of events, represented by the order set $\mathcal{H} = \{e_i = (u_i, t_i, c_i)\}_{i=1}^{N_c}$, and each event refers to a user u_i , a time of occurrence t_i (e.g. time of a tweet), and a cascade c_i to which it contributes. The event history $\mathcal{H}(t) = \{e_i = (u_i, t_i, c_i) | t_i < t\}$ is the set of events occurring in the time interval $[0, t]$ related to any cascade; $\mathcal{H}_u(t)$ refers to the events of user u ; $\mathcal{H}^{(c)}(t)$ refers to cascade c , and $\mathcal{H}_u^{(c)}(t)$ is only about u and c . We omit the time index when referring to the total history of length T , i.e. $\mathcal{H} := \mathcal{H}(T)$, and similarly for other time-dependent notation. Tab. A provides a detailed symbol list.

Multidimensional Hawkes processes (MHPs). Hawkes processes [34] is a class of self-exciting temporal point processes suitable for describing recurrent discrete asynchronous events. They are widely used in modeling collective phenomena characterized by activity bursts, like those observed in OSNs. MHPs can be formulated via a counting process vector $\mathbf{n}(t) \in \mathbb{N}^{N_u}$, which counts the amount of events up to time t for each user, and $d\mathbf{n}(t)$ denotes the infinitesimal increase of $\mathbf{n}(t)$ at t . Given a history $\mathcal{H}(t)$, the occurrence rate for an event created by any user in the interval $[t, t+dt]$ is expressed by the conditional *intensity function* $\boldsymbol{\lambda}(t)$:

$$\boldsymbol{\lambda}(t)dt := \mathbb{E}[d\mathbf{n}(t) | \mathcal{H}(t)]. \quad (1)$$

The intensity for a user u is defined as (see [34]):

$$\lambda_u(t) = \mu_u(t) + \sum_{v \in \mathcal{F}_u} w_{vu}(t) \star dn_v(t), \quad (2)$$

where $\mu_u(t)$ is the *background intensity*, \mathcal{F}_u is the set of users that u follows, $w_{vu}(t)$ is a temporal kernel for the *social influence* exerted by v on u , and \star denotes the convolution product. Note that in many existing works on social network analysis with practical considerations [6], [14] the background intensity μ_u is set to be stationary, and the social influence of v to u induced by an event of v that occurred at time t' is expressed as $w_{vu}(t) := w_{vu}\kappa(t-t')$, where $\kappa(\cdot)$ is a time-decaying kernel. Thus, the simplified intensity function is:

$$\lambda_u(t) = \mu_u + \sum_{v \in \mathcal{F}_u} w_{vu} \sum_{e_i \in \mathcal{H}_v(t)} \kappa(t-t_i). \quad (3)$$

Unlike synchronous and social averaging models, a time-decaying kernel allows for more refined user attention mechanisms. In this work, we use a typical exponential kernel: $\kappa(t) := \mathbb{1}\{t > 0\} \exp(-t/\tau)$, where smaller values of $\tau \in \mathbb{R}_+$ accelerate the fade-out effect, though our analysis remains valid for other types of time-decaying kernels. As a literature note, MHPs have been generalized to multiple cascades or event types [38], [39], and learning non-parametric triggering kernels for Hawkes processes is of broad interest [4], [36].

Multidimensional marked Hawkes processes (MMHPs). Marked Hawkes processes offer one way to

generalize MHPs for a network of users by accommodating multiple cascading topics. Each mark c corresponds to one cascade and is assigned to a user u via the user’s *conditional mark probability density function* $f_u(c|t)$. The marked intensity of user u for cascade c is written as (see [40]):

$$\lambda_u^{(c)}(t) = \lambda_u(t)f_u(c|t), \quad (4)$$

which is time-dependent, hence depends on the event history. The total marked intensity of cascade c is: $\lambda^{(c)}(t) = \sum_{u \in \mathcal{V}} \lambda_u(t)f_u(c|t)$. An interesting link with opinion formation is that f_u can be seen as a representation of u ’s relative opinion on the N_c topics, and $\lambda_u^{(c)}(t)$ as the expressed opinion through the marked user intensity. Most importantly, f_u allows to formalize joint user-cascade dynamics. Probabilistic models have been defined for f_u (e.g. in [38], a Dirichlet distribution has been used for cascade dynamics, yet user dynamics were not considered), while the CC model [14] proposed the use of the *Boltzmann distribution* (i.e. the *soft-max* function) for reweighting f_u ’s components and this way allowing for cascade competition. Finally, note that the IC model can be expressed as an MMHP by considering f_u to be a simple linear normalization of the user intensities $f_u(\cdot|t)$ [14]. More light will be shed on this discussion in Sec. 4.

4 Mixture of Interacting Cascades (MIC)

The MIC model is defined over the MMHPs framework and builds upon the idea that user interactions take place over a landscape of contextually interacting topics (see the conceptual illustration in Fig. 1). More specifically, the cascade interactions are represented as *positive pairwise contextual links*, which are taken into account for defining a context-sensitive independent user intensity for each cascade. These are then fed to a mixing function that shapes the final conditional mark probability density f_u . The main formulas of the model follow.

The MIC model

- Independent user intensity for a cascade:

$$\nu_u^{(c)}(t) = \mu_u^{(c)} + \sum_{v \in \mathcal{F}_u} w_{vu} \sum_{e_i \in \mathcal{H}_v^{(c)}(t)} \kappa(t-t_i) \quad (5)$$

- Context-sensitive indep. user intensity for a cascade:

$$\nu_u^{*(c)}(t) = \mu_u^{(c)} + \sum_s \sigma_{sc} \sum_{v \in \mathcal{F}_u} w_{vu} \sum_{e_i \in \mathcal{H}_v^{(s)}(t)} \kappa(t-t_i) \quad (6)$$

- Cascade mixing function:

$$f_u(c|t; \phi) = \frac{\phi(\nu_u^{*(c)}(t))}{\sum_k \phi(\sum_s \nu_u^{*(s)}(t))} \quad (7)$$

- Global user intensity:

$$\lambda_u(t) = \sum_{c \in \mathcal{C}} \nu_u^{(c)}(t) = \mu_u + \sum_{v \in \mathcal{F}_u} w_{vu} \sum_{e_i \in \mathcal{H}_v(t)} \kappa(t-t_i) \quad (8)$$

- Marked user intensity for a cascade:

$$\lambda_u^{(c)}(t) = \lambda_u(t)f_u(c|t) \quad (9)$$

Above, $\mu_u^{(c)}$ represents the baseline user interest for cascade c , where $\forall u, \mu_u = \sum_{c \in \mathcal{C}} \mu_u^{(c)}$; the context-sensitive independent user intensity for a cascade, $\nu_u^{*(c)}$, combines all the associated contextual links of the *cascade interaction matrix* $\boldsymbol{\Sigma} = [\sigma_{sc}]_{sc}$

TABLE 1

Regimes of pairwise cascade interaction for two cascades, s and c , induced by the cascade mixing function f_u (Eq. 7) of the MIC model.

The characterization depends on the involved map ϕ , the scaling hyperparameter β of the Boltzmann distribution, and the influence σ_{sc} . MIC encompasses its linear variant **linMIC**, **IC** [12], and **CC** [14].

	MIXING FUNCTION f_u	NEUTRALITY $\sigma_{sc} = 0$	REINFORCEMENT $0 < \sigma_{sc} \leq 1$
UNIFORMITY	$\phi(x) := \exp(\beta x), \beta \rightarrow 0$	CC MIC	MIC
LINEARITY	$\phi(x) := x$	IC linMIC	linMIC
COMPETITION	$\phi(x) := \exp(\beta x), \beta > 0$	CC MIC	MIC

where $\sigma_{sc} \in [0, 1]$ is the contextual influence that a cascade s exerts on cascade c , and by convention $\forall c, \sum_s \sigma_{sc} = 1$. The contextual links are generally asymmetric, i.e. $\sigma_{sc} \neq \sigma_{cs}$; also, being positive implies interactions between two cascades ranging from being neutral to being reinforcing.

The *cascade mixing function* (Eq. 7) maps each $\nu_u^{*(s)}$ through ϕ ; for instance, with $\phi(x) := x$, f_u performs a simple normalization, while for $\phi(x) := \exp(\beta x)$ (with $\beta \geq 0$ being a hyperparameter), f_u becomes the Boltzmann distribution that reshapes the density. Regardless the choice of ϕ , f_u introduces one-vs-all cascade competition. Especially for the case of the Boltzmann distribution, the competition is modulated by β : As $\beta \rightarrow \infty$, f_u assigns all the activity to the most popular cascade, i.e. for $c = \arg \max_s f_u(s|t)$: $f_u(c|t) \approx 1$, otherwise $f_u(c|t) \approx 0$. For $\beta \rightarrow 0$, f_u converges to a uniform distribution, i.e. $f_u(c|t) \approx 1/N_c$. Finally, note that from a statistical mechanics point-of-view, the denominator of the Boltzmann distribution is the *partition function* of the conditional mark probability f_u , denoted by $Z_u(t) = \sum_k \exp(\beta \nu_u^{*(s)}(t))$.

The complete MIC model is defined by the parameters set $\Theta = ([\mu_u^{(c)}]_{uc}, [\sigma_{sc}]_{sc}, [w_{vu}]_{vu}) = (\mathbf{M}, \mathbf{\Sigma}, \mathbf{W})$, which gives a total of $N_u \times N_c + N_c^2 + N_u^2$ parameters to be inferred, and β can be estimated by cross-validation. Given that usually $N_c \ll N_u$, inferring the additional N_c^2 parameters of the model does not change the order of the computational cost.

Comparison with other closely related models. For $\mathbf{\Sigma} = \mathbf{I}$ and linear f_u mixing ($\phi(x) := x$), MIC recovers the IC model from [13] (hence, it holds $\lambda_u^{(c)}(t) = \nu_u^{(c)}(t)$). Otherwise, using the f_u with exponential mixing (Boltzmann distribution), MIC recovers the CC model from [14]. As mentioned earlier, f_u introduces mere one-vs-all competition, which is included in all the IC, CC, and MIC models. However, our formulation is more flexible since it allows pairwise asymmetric contextual links between cascades, which in turn help us define contextual user intensities, and this way extend the notion of cascade competition. Tab. 1 summarizes how the compared models encompass the different regimes of cascade interactions.

5 Learning a MIC model

MIC proposes the addition of N_c^2 parameters to capture pairwise cascade interactions (see Sec. 4). Here, our intention is show that those parameters can be efficiently learned by a comprehensive Maximum Likelihood Estimation (MLE) approach, which at the same time is not complex. Moreover, when the likelihood can be factorized into a product of user likelihoods, then thanks to the partial log-likelihood formulas, these optimization sub-problems can be solved in parallel. Although developing a highly competitive optimization scheme

is not in the scope of this paper, in the literature there are works dedicated to parameter inference for Hawkes processes, which use methods such as kernel decomposition [4], [21], [36] or Bayesian approaches [41], [42].

5.1 Computation of the log-likelihood

First, using the formulation introduced by [43], for a MMHP and given the model parameters Θ , the conditional density function for the event $e_i = (u_i, c_i, t_i)$ to occur at t_i , combined with no other event to occur before t_{i+1} , is written as:

$$l(e_i, t_{i+1}; \Theta) = \lambda_{u_i}^{(c_i)}(t_i) \mathbb{P}(\text{no event in } (t_i, t_{i+1}]), \quad (10)$$

where $\lambda_{u_i}^{(c_i)}(t_i)$ is the rate of occurrence of event e_i . The probability of Eq. 10 is known as *exponential survival function*, which we write here for an arbitrary time interval $(t, t']$ as:

$$\begin{aligned} S(t, t') &:= \mathbb{P}(\text{no event in } (t, t']) = \exp\left(-\int_t^{t'} \sum_{u \in \mathcal{V}} \lambda_u(\tau) d\tau\right) \\ &= \exp\left(-\sum_{u \in \mathcal{V}} \left[(t' - t)\mu_u + \sum_{v \in \mathcal{F}_u} \sum_{e_j \in \mathcal{H}_v(t, t')} w_{vu} K(t, t' - t_j)\right]\right), \end{aligned} \quad (11)$$

where $K(t, t' - t_j) = \int_t^{t-t_j} \kappa(\tau) d\tau$, with $t < t_j \leq t'$ and $\kappa(\cdot)$ the exponential kernel.

Next, we express MIC's full observation likelihood of the whole event history $\mathcal{H} = \{e_i = (u_i, t_i, c_i)\}_{i=1}^{N_e}$ as:

$$\begin{aligned} L(\mathcal{H}; \Theta) &= \prod_{e_i \in \mathcal{H}} l(e_i, t_{i+1}; \Theta) = \prod_{e_i \in \mathcal{H}} [\lambda_{u_i}^{(c_i)}(t_i) S(t_i, t_{i+1})] \\ &= \left[\prod_{e_i \in \mathcal{H}} \lambda_{u_i}^{(c_i)}(t_i) \right] S(0, T), \end{aligned} \quad (12)$$

where $\prod_{e_i \in \mathcal{H}} S(t_i, t_{i+1}) = S(0, T)$, given Eq. 11. As mentioned earlier, the full likelihood can be factorized by the partial likelihoods for each user u as: $L(\mathcal{H}; \Theta) = \prod_{u \in \mathcal{V}} L_u(\mathcal{H}; \Theta)$, where, using the marked user intensity of Eq. 4, we can write:

$$\begin{aligned} L_u(\mathcal{H}; \Theta) &= \left[\prod_{e_i \in \mathcal{H}_u} \lambda_{u_i}^{(c_i)}(t_i) \right] S(0, T) \\ &= \left[\prod_{e_i \in \mathcal{H}_u} \lambda_u(t_i) f_u(c_i | t_i) \right] S(0, T). \end{aligned} \quad (13)$$

Next step is to give the expression for the log-likelihood:

$$\begin{aligned} \mathcal{L}(\mathcal{H}; \Theta) &= \left[\sum_{e_i \in \mathcal{H}} \log(\lambda_{u_i}(t_i) f_{u_i}(c_i | t_i)) \right] + S(0, T) \\ &= \left[\sum_{e_i \in \mathcal{H}} \log(\lambda_{u_i}(t_i)) + \beta \nu_{u_i}^{*(c_i)}(t_i) \right. \\ &\quad \left. - \sum_{e_i \in \mathcal{H}} \log\left(\sum_{s \in \mathcal{C}} \exp(\beta \nu_{u_i}^{*(s)}(t_i))\right) \right] \\ &\quad + S(0, T) \\ &= \left[\sum_{e_i \in \mathcal{H}} \log(\lambda_{u_i}(t_i)) + \beta \nu_{u_i}^{*(c_i)}(t_i) \right. \\ &\quad \left. - \sum_{e_i \in \mathcal{H}} \log(Z_{u_i}(t_i)) \right] \\ &\quad - \sum_{u \in \mathcal{V}} \left[T\mu_u + \sum_{v \in \mathcal{F}_u} \sum_{e_j \in \mathcal{H}_v} w_{vu} K(0, T - t_j) \right], \end{aligned} \quad (14)$$

where $Z_{u_i}(t_i)$ is the partition function defined in Sec. 4. The partial log-likelihood is therefore expressed as follows:

$$\begin{aligned} \mathcal{L}_u(\mathcal{H}_u; \Theta) = & \left[\sum_{e_i \in \mathcal{H}_u} \log(\lambda_u(t_i)) + \beta \nu_u^{*(c_i)}(t_i) \right] \\ & - \sum_{e_i \in \mathcal{H}_u} \log(Z_{u_i}(t_i)) \\ & - T\mu_u - \sum_{v \in \mathcal{F}_u} \sum_{e_j \in \mathcal{H}_u} w_{vu} K(0, T - t_j). \end{aligned} \quad (15)$$

Proposition 1: The negative log-likelihood \mathcal{L} is a convex function of the parameters Θ .

Proof. The convexity of the negative partial log-likelihood is a direct consequence of the convexity of the LogSumExp function $\varphi: \mathbb{R}^{N_c} \rightarrow \mathbb{R}$ [44], defined as:

$$\varphi(\mathbf{z}_i) = \log \left(\sum_{s=1}^{N_c} \exp(z_i^{(s)}) \right), \quad (16)$$

where $\forall s \in \mathcal{C}$, $z_i^{(s)} = \beta \nu_u^{*(c_i)}(t_i)$. For MIC, the set of parameters $\Theta = (\mathbf{M}, \Sigma, \mathbf{W})$ are positive, ensuring the preservation of convexity. Similarly, the other terms of the negative log-likelihood are convex functions of the parameters, which proves the overall convexity of \mathcal{L}_u , and subsequently of \mathcal{L} . In the case of linear mixing, in Eq. 12 we replace $\lambda_{u_i}^{(c_i)}(t_i)$ by $\nu_{u_i}^{(c_i)}(t_i)$ and we perform the same reasoning, which also leads to a convex formulation of the log-likelihood. \square

5.2 Parameter inference algorithm

Relying on the expression of the partial and total log-likelihood (Eq. 14 and Eq. 15) we decompose the parameter inference scheme into two optimization steps, presented in Alg. 1. Problem 17 uses Maximum Likelihood Estimation (MLE) to find the cascade interaction matrix $\Sigma = [\sigma_{cs}]_{cs}$ by minimizing $-\mathcal{L}(\mathcal{H}; \Theta)$, with (\mathbf{M}, \mathbf{W}) fixed. For this, we use Newton's method to approximate Σ for when the respective gradient is zero, $\nabla_{\Sigma} \mathcal{L} = 0$. Then, Problem 18 uses Maximum Partial Likelihood Estimation (MPLE) to find the user parameters $\mathbf{M}_u = (\mu_u^{(c)})_c$ and $\mathbf{W}_u = (w_{vu})_v$, with Σ fixed. Due to the per-user factorization of the log-likelihood, the N_u sub-problems can be solved numerically in parallel.

The two problems are treated by the alternating optimization scheme of Alg. 1 until convergence. Note that the hyperparameter β is estimated by cross-validation. As in this work focuses on modeling the interplay of user and cascade interactions, we are mainly interested in showing that learning is feasible with a typical optimization scheme. In that sense, although our proposition employs alternating optimization, the rest of the inference scheme comes from the prior line of related work and up to [14].

6 Analytic derivation of characteristic quantities

In this section, we derive closed-form expressions for the conditional intensity and the number of events for the MIC model. The derivation follows a rather simple method to obtain moments of Hawkes processes, based on [45]. These analytical results are compared and validated in Sec. 7 with experiments both on synthetic (Fig. 3) and real data (Fig. 5).

Algorithm 1 Inference of MIC parameters

Input: history of train events \mathcal{H} , termination threshold ϵ
Output: the set of inferred parameters $\hat{\Theta} = (\hat{\mathbf{M}}, \hat{\mathbf{W}}, \hat{\Sigma})$

- 1: Initialize $\Theta = (\mathbf{M}, \mathbf{W}, \Sigma)$
- 2: Let $\hat{\Theta} = \Theta$
- 3: **repeat**
- 4: ■ Step 1 – Underlying cascade interaction
- 5: $\hat{\Sigma} = \arg \min_{\Sigma} -\mathcal{L}(\mathcal{H}; \mathbf{M}, \mathbf{W})$ (17)
 s.t. $\forall c, s \in \mathcal{C}, \sigma_{sc} \in [0, 1]$ and $\sum_c \sigma_{sc} = 1$
- 6: ■ Step 2 – Social influence and background intensity
- 7: **for** $u \in \mathcal{V}$ **do in parallel**
- 8: $\hat{\mathbf{M}}_u, \hat{\mathbf{W}}_u = \arg \min_{\mathbf{M}_u, \mathbf{W}_u} -\mathcal{L}_u(\mathcal{H}; \hat{\Sigma})$ (18)
 s.t. $\mathbf{M}_u, \mathbf{W}_u \geq 0$
- 9: **end for**
- 10: $\hat{\Theta} = \hat{\Theta}$
- 11: **until** $\mathcal{L}(\hat{\Theta}) - \mathcal{L}(\hat{\Theta}) < \epsilon$
- 12: **return** $\hat{\Theta} = (\hat{\mathbf{M}}, \hat{\mathbf{W}}, \hat{\Sigma})$

6.1 Generic derivation

Step 1: Let us consider the quantity $\mathbb{E}[n_v^{(s)m}(t) \nu_u^{(c)q}(t)]$, where $n_v^{(s)}(t)$ is the number of events of user v for cascade s and $\nu_u^{(c)}(t)$ is the independent intensity of user u for cascade c , both at time t and exponentiated by $m, q \in \mathbb{N}$.

Step 2: We calculate the probabilities of the absence or occurrence of an arbitrary event $e_i = (u_i, c_i, t_i)$ that is or is not associated to user v and cascade s , at order $o(\Delta t)$:

$$\begin{aligned} \mathbb{P}(\neg e_i | \mathcal{H}_t) &= 1 - \Delta t \sum_{v \in \mathcal{V}} \sum_{p \in \mathcal{C}} \lambda_v^{(p)}(t) + o(\Delta t) \\ \mathbb{P}(e_i \neq (v, s, t_i) | \mathcal{H}_t) &= \Delta t \sum_{v \in \mathcal{V}} \sum_{\substack{p \in \mathcal{C} \\ (v, p) \neq (v, s)}} \lambda_v^{(p)}(t) + o(\Delta t) \quad (19) \\ \mathbb{P}(e_i = (v, s, t_i) | \mathcal{H}_t) &= \Delta t \lambda_v^{(s)}(t) + o(\Delta t). \end{aligned}$$

Step 3: We express the evolution of the independent cascade intensity at time $t + \Delta t$ given the absence and the occurrence of an event $e_i = (u_i, c_i, t_i)$: When no event occurs:

$$\begin{aligned} \nu_{u,0}^{(c)}(t + \Delta t) &= \mu_u^{(c)} + \sum_{j \in \mathcal{F}_u} w_{vu} \sum_{e_j \in \mathcal{H}_u(t)} e^{-(t-t_j+\Delta t)/\tau} \\ &= \nu_u^{(c)}(t) - \Delta t (\nu_u^{(c)}(t) - \mu_u^{(c)})/\tau + o(\Delta t), \end{aligned} \quad (20)$$

where τ controls the scale of the temporal kernel κ (see Sec. 3). When there is an event (u_i, c_i, t_i) with $t < t_i < t + \Delta t$:

$$\nu_{u,e_i}^{(c)}(t + \Delta t) = \nu_{u,0}^{(c)}(t + \Delta t) + w_{u_i u} \delta_{c_i, c} (1 - \Delta t') + o(\Delta t')/\tau, \quad (21)$$

where $\delta_{\cdot, \cdot}$ is the Kronecker delta function and $\Delta t' = t_i - t$ is the time passed after the occurrence of event e_i (i.e. $0 < \Delta t' < \Delta t$). Here we extend the notation of influence: $w_{u_i u} = 0$ if $u_i \notin \mathcal{F}_u$.

Step 4: We calculate and simplify the conditional expectation $h(t + \Delta t) = \mathbb{E}[n_v^{(s)}(t + \Delta t) \nu_u^{(c)}(t + \Delta t) | \mathcal{H}_t]$ given the two events:

$$\begin{aligned} h(t + \Delta t) &= n_v^{(s)m}(t) (\nu_{u,0}^{(c)})^q(t) \mathbb{P}(\neg e_i | \mathcal{H}_t) \\ &\quad + n_v^{(s)}(t) (\nu_{u,e_i}^{(c)})^q(t) \mathbb{P}(e_i \neq (w, s, t_i) | \mathcal{H}_t) \\ &\quad + (n_v^{(s)}(t) + 1)^m (\nu_{u,e_i}^{(c)})^q(t) \mathbb{P}(e_i = (w, s, t_i) | \mathcal{H}_t). \end{aligned} \quad (22)$$

With $\lim_{t \rightarrow 0} \frac{\mathbb{E}[h(t+\Delta t) - h(t)]}{\Delta t} = \frac{d\mathbb{E}[n_v^{(s)m}(t)\nu_u^{(c)q}(t)]}{dt}$, using Eq. 19, Eq. 20 and Eq. 21, and taking the full expectation, we obtain the following differential equation for $\mathbb{E}[n_v^{(s)m}(t)\nu_u^{(c)q}(t)]$:

$$\begin{aligned} \frac{d}{dt} \mathbb{E}[n_v^{(s)m}(t)\nu_u^{(c)q}(t)] = & \\ & - n \mathbb{E}[n_v^{(s)m}(t)\nu_u^{(c)q-1}(t)(\nu_u^{(c)}(t) - \mu_u^{(c)})]/\tau \\ & + \sum_{j \in \mathcal{V}} \sum_{k=0}^{q-1} \binom{n}{k} w_{ju}^{q-k} \mathbb{E}[n_v^{(s)m}(t)\lambda_j^{(c)}(t)\nu_u^{(c)k}(t)] \\ & + \sum_{l=0}^{m-1} \sum_{k=0}^q \binom{m}{l} \binom{n}{k} w_{vu}^{q-k} \delta_{s,c}^{q-k} \mathbb{E}[n_v^{(s)l}(t)\lambda_v^{(s)}(t)\nu_u^{(c)k}(t)]. \end{aligned} \quad (23)$$

Note that the derivation relies on the same assumptions of regularity of Hawkes processes as in [45]. The moments of the conditional intensity and the number of events $\mathbb{E}[n_v^{(s)m}(t)\lambda_u^q(t)]$ can be derived in a similar way for the MIC model:

$$\begin{aligned} \frac{d}{dt} \mathbb{E}[n_v^{(s)m}(t)\lambda_u^q(t)] = & - n \mathbb{E}[n_v^{(s)m}(t)\lambda_u^{q-1}(t)(\lambda_u(t) - \mu_u)]/\tau \\ & + \sum_{j \in \mathcal{V}} \sum_{k=0}^{q-1} \binom{n}{k} w_{ju}^{q-k} \mathbb{E}[n_v^{(s)m}(t)\lambda_j(t)\lambda_u^k(t)] \\ & + \sum_{l=0}^{m-1} \sum_{k=0}^q \binom{m}{l} \binom{n}{k} w_{vu}^{q-k} \mathbb{E}[n_v^{(s)l}(t)\lambda_v(t)\lambda_u^k(t)]. \end{aligned} \quad (24)$$

6.2 Application: computation of the global user intensity and the number of events

The differential Eq. 23 can be solved for specific values of m, n .

Expectation of the cascade intensity. For $m=0$ and $q=1$, we obtain the differential equation for the expectation of the independent cascade intensity:

$$\begin{aligned} \frac{d}{dt} \mathbb{E}[\nu_u^{(c)}(t)] = & - (\mathbb{E}[\nu_u^{(c)}(t)] - \mu_u^{(c)})/\tau + \sum_{j \in \mathcal{V}} w_{ju} \mathbb{E}[\lambda_j(t) f_j(c|t)] \\ = & - (\mathbb{E}[\nu_u^{(c)}(t)] - \mu_u^{(c)})/\tau + \sum_{j \in \mathcal{V}} w_{ju} \mathbb{E}[\lambda_j^{(c)}(t)]. \end{aligned} \quad (25)$$

In this form, the first term corresponds to the exponential decay that accounts for the recency of past events, while the second term is the contribution from every neighboring activity at time t for cascade c .

Summing Eq. 25 over c (or equivalently taking $(m, n) = (0, 1)$ in Eq. 24), we obtain the differential equation for the expectation of user intensity $\mathbb{E}[\lambda_u(t)]$:

$$\frac{d}{dt} \mathbb{E}[\lambda_u(t)] = - (\mathbb{E}[\lambda_u(t)] - \mu_u)/\tau + \sum_{v \in \mathcal{V}} w_{vu} \mathbb{E}[\lambda_v(t)]. \quad (26)$$

Eq. 26 can be expressed in vectorial form for all users u :

$$\frac{d}{dt} \mathbb{E}[\boldsymbol{\lambda}(t)] = - (\mathbb{E}[\boldsymbol{\lambda}(t)] - \boldsymbol{\mu})/\tau + \mathbf{W}^T \mathbb{E}[\boldsymbol{\lambda}(t)]. \quad (27)$$

This equation is analogous to the one derived in [18], [32] and can be analytically solved for the MIC model, if we consider the initial condition $\mathbb{E}[\boldsymbol{\lambda}(0)] = \boldsymbol{\mu}$:

$$\mathbb{E}[\boldsymbol{\lambda}(t)] = \left[\mathbf{B} + (\mathbf{I} - \mathbf{B})e^{-\mathbf{B}^{-1}t/\tau} \right] \boldsymbol{\mu}, \quad (28)$$

where $\mathbf{B} = (\mathbf{I} - \mathbf{W}^T \tau)^{-1}$. The expected intensity converges to the stationary value $\boldsymbol{\lambda}_\infty = \mathbf{B}\boldsymbol{\mu}$ if it holds for the spectral radius: $\rho(\mathbf{W}^T) < 1/\tau$. This condition can be interpreted as the balance between the amplitude of event social influence and its temporal decay. Then, after a negative exponential increase, the expected intensity tends to the stationary value $\mathbf{B}\boldsymbol{\mu}$.

Expected number of events. Next, the definition of the expectation of the number of events can be recovered in Eq. 23 for $m=1, q=0$: $\frac{d\mathbb{E}[n_v^{(s)}(t)]}{dt} = \mathbb{E}[\lambda_v^{(s)}(t)]$. The expectation of the number of events can then be expressed as:

$$\mathbb{E}[\mathbf{n}(t)] = \left[\mathbf{B}t + (\mathbf{I} - \mathbf{B})\mathbf{B}\tau \left(\mathbf{I} - e^{-\mathbf{B}^{-1}t/\tau} \right) \right] \boldsymbol{\mu}, \quad (29)$$

with the initial condition $\mathbb{E}[\mathbf{N}(0)] = 0$. Therefore, the expected number of events tends to a linear regime of slope $\mathbf{B}\boldsymbol{\mu}$. Given the non-independence between the intensity λ_u and the cascade mixing function f_u , the analytical expression of the expected cascade-related intensity is not straightforward and goes beyond our scope. Instead, we use the following *stationary state* approximation:

$$\mathbb{E}[\lambda_u^{(c)}(t)] \approx \mathbb{E}[\lambda_u(t)] \frac{\exp(\beta(\mathbf{B}\boldsymbol{\mu}\boldsymbol{\Sigma})_{u,c})}{\sum_{s \in \mathcal{C}} \exp(\beta(\mathbf{B}\boldsymbol{\mu}\boldsymbol{\Sigma})_{u,s})}, \quad (30)$$

where recall that $\boldsymbol{\Sigma}$ is the cascade interaction matrix.

7 Experiments

This section evaluates both qualitatively and quantitatively the performance of the MIC model. The compared methods and the evaluation measures used are in Sec. 7.1.

In Sec. 7.2, we start by designing synthetic experiments with data generated by simulating the MIC model. The aim is to evaluate MIC's capacity to encompass different information cascade dynamics, and therefore to show that MIC generalizes over the existing models discussed earlier. In Sec. 7.3, we test how the model handles the complexity of real data. We are also interested in providing an ablation study of the features of MIC, so as to assess the impact of the cascade mixing function (Boltzmann or linear).

Next, in Sec. 7.4 we analyze the models' behavior from a user-centric point of view. This is a frequently neglected side in similar studies that focus on the global behavior of the agents within a social network. Although this latter is of high significance [9], [10], neglecting the user-wise evaluation of user-based models may limit our understanding of what happens at the level of individual users.

Finally, in Sec. 7.5 we demonstrate MIC's potential for producing comprehensive data visualizations of the social network activity landscape, which provide a qualitative insight into the model's behavior and the underlying user-to-user, user-to-cascade, and cascade-to-cascade interactions.

7.1 Competitors and evaluation measures

We consider as competitors the Independent Cascades (IC) model [19], [20], [32], and the Correlated Cascades (CC) model [14]. We also include the linear variant presented earlier, **linMIC**. As discussed in Sec. 4 (see also Tab. 1), all compared models can be mapped to special cases of MIC, therefore, the results that follow can be also analyzed from an ablation study viewpoint for MIC's features. Moreover, treating them as MIC




Fig. 2. Heatmaps of the test log-likelihood ratios between the competitors and MIC, with varying β (y-axis) and the cascade interaction σ_{21} (x-axis), on synthetic event logs generated by MIC. A ratio value larger than 1 means that MIC performs better. The color scale is the same for all the plots.

cases conveniently allowed us to perform the simulation of all models using the same code implementation¹².

Each experiment takes as input the adjacency matrix of a user graph and an activity log containing events of the form: $\mathcal{H} = \{e_i = (\text{user: } u_i, \text{cascade: } c_i, \text{timestamp: } t_i)\}_{i=1}^{N_e}$. This can be a real event log, or a synthetic one where the events are generated by a source model, e.g. using Ogata’s thinning algorithm [46]. Then, the objective is to infer the model parameters, e.g. for MIC, these are $\Theta = (\Sigma, \mathbf{M}, \mathbf{W})$. We use a train-test split where the first 80% of the events is used as training set, and the last 20% is used for testing the models. We use two ways to evaluate and compare the models:

- *Comparing the goodness-of-fit*: we compare how well a probabilistic model fits to an input event log, namely, the *test log-likelihoods* of the models computed on the test events of a given dataset $\mathcal{H}^{\text{test}}$. A high test log-likelihood, $\mathcal{L}^{\text{test}} < 0$ and as close as possible to zero is desired. The test log-likelihood can be also computed for a varying fraction of the train dataset $\mathcal{H}^{\text{train}}$ (e.g. the last 20% or 60% train events before the test dataset $\mathcal{H}^{\text{test}}$) to evaluate model efficiency as a function of the data availability.
- *Comparing the realized event activity*: we compare the empirical event intensity (rate of events over time) of the events generated by a model against the activity that comes either from a reference model (when this is available), or by an input event log. Similarly, we compare the activity in terms of number of events generated. We use as evaluation measures to compare such quantities: i) *Inverse l_1 -distance*; ii) *Pearson correlation*; iii) *Distribution of ranked user (resp. cascade) activity*: this refers to the comparison of the distribution of the user (resp. cascade) activity with respect to their ranking according to their total volume of activity.

These measures can be used at a global level (for all users and all cascades), user-wise and/or cascade-wise, to evaluate model performance from different perspectives. We repeat

1. The implementation for MIC and all the competitors will be made publicly available.

2. Although relevant, we do not include the Competing Products model [6] as it has been shown that it is clearly outperformed by the former competitors. Note also that our results can be put in perspective to the slightly different context of opinion formation, as for $N_c = 2$ cascades, the SLANT model [32] with binary opinion values is analogous to the IC model.

each experiment 10 times to get a model’s empirical average behaviors.

7.2 Experiments on synthetic data

7.2.1 Data generation

Synthetic data are generated by the MIC model, in a random Erdős-Rényi network of $N_u = 50$ users and for $N_c = 3$ cascades. For each edge (v, u) between users, we set the social influence at $w_{vu} \sim \mathcal{U}([0, 1])$. For each user u we set the baseline intensities for the three cascades at $\mu_u^{(c)} \sim \mathcal{U}([0, 0.2])$. We explore the cascade dynamics along two dimensions of parameterization:

- by varying the cascade competitiveness via $\beta \in [0.01, 100]$,
- by varying the cascade interaction via Σ : the reinforcement from c_2 to c_1 ($\sigma_{21} \in [0, 1]$) is gradually increasing as $\Sigma = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \rightarrow \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & \sigma_{21} \\ 0 & 0 & 1 \end{bmatrix}$.

This explores a wide range of interactions between cascades, gradually going from pure independence to complex and asymmetric competition or cooperation. To get statistically sounder results, synthetic event logs are generated 10 times for each parameterization, and for $T = 500$ time units and $\tau = 3$ times units.

7.2.2 Goodness-of-fit evaluation

The scenario challenges the compared models to learn from data generated by the MIC model, as described in Sec. 7.2.1. Note that CC and MIC are learned taking as input the true β and τ values used to generate the train data. The models are then compared on how well they generalize to unseen synthetic event logs. Fig. 2 shows heatmaps of the test log-likelihood ratios $\mathcal{L}_{\langle \text{model} \rangle}^{\text{test}} / \mathcal{L}_{\text{MIC}}^{\text{test}}$. Starting from the uniform and independent cascades regime (low β), MIC has on average a similar or slightly better performance to the three other methods. As β increases and brings nonlinear competitiveness, MIC becomes gradually more accurate than IC (Fig. 2a) and linMIC (Fig. 2b). Looking at those two first subfigures, we can also indirectly compare IC and linMIC: the heatmap patterns are rather similar except from a difference in the scale of the ratio. It can also be observed that the performance of IC decreases as σ_{21} grows (along the horizontal axis) whereas linMIC is not affected by this variation, as expected by its design, which suggests that overall linMIC behaves similarly

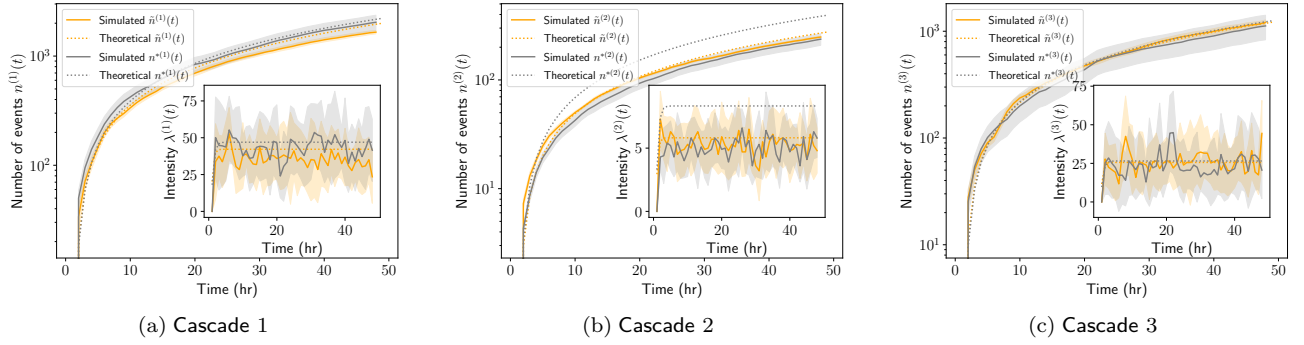


Fig. 3. Number of events and intensity of the three cascades over time, for the true $(n^{*(\cdot)}(t), \lambda^{*(\cdot)}(t))$, and the generated events given the learned MIC model $(\tilde{n}^{(\cdot)}(t), \tilde{\lambda}^{(\cdot)}(t))$. Error bars correspond to the simulation on 10 event generations. Dotted lines are the theoretical expected quantities for the MIC model, computed using both the true and the inferred parameter values. The event dataset has been generated with the following parameterization: $\beta = 33.37$ and $\Sigma = \begin{bmatrix} 1 & 0 & 0 \\ .71 & .29 & 0 \\ 0 & 0 & 1 \end{bmatrix}$.

to IC, i.e. it is insensitive to both dimensions of setting variation, while having a bit better performance. Additionally, as the interaction between cascades becomes more intricate, MIC starts to outperform CC (Fig. 2c). The results suggest that MIC is more flexible than the compared special cases, which is consistent with the regime summary given in Tab. 1.

7.2.3 Comparing empirical intensities and number of events

In this part we focus on the evaluation of the MIC inference scheme (Sec. 5.2), as well as the precision of the theoretical results (Sec. 6). We choose an instance of the previous experiment, namely the synthetic dataset generated by MIC with $\beta = 33.37$, $\Sigma = \begin{bmatrix} 1 & 0 & 0 \\ .71 & .29 & 0 \\ 0 & 0 & 1 \end{bmatrix}$. We treat β and τ as known so as to assess independently the inference of the cascade interaction matrix (Alg. 1), which in this case is recovered accurately: $\tilde{\Sigma} = \begin{bmatrix} 1 & 0 & 0 \\ .73 & .25 & 0 \\ 0 & 0 & 1 \end{bmatrix}$. Next, to evaluate the precision of the other inferred parameters $(\tilde{\mathbf{M}}, \tilde{\mathbf{W}})$, we compare the number of events and intensity of the true $(n^{*(\cdot)}(t), \lambda^{*(\cdot)}(t))$, with the ones generated by simulations of the learned MIC model $(\tilde{n}^{(\cdot)}(t), \tilde{\lambda}^{(\cdot)}(t))$. To validate our theoretical results, we also compute the analytical expectation of number of events and intensity for both ground truths and learned model.

Fig. 3 shows the temporal evolution of both simulated and theoretical conditional event intensities and the realized number of events for each cascade in the generated data. It can be noticed that the inferred model is able to replicate the true event generating process, both in terms of average intensities and variance. In addition to that, the theoretical and simulated conditional intensities agree well for each cascade: the expected intensity converges to the stationary value after a negative exponential increase. The second cascade, less active, is however underestimated by simulations, which probably originates from the stationary state approximation made in Eq. 30. Nevertheless, this experiment indicates that MIC is not only statistically more adaptive than its competitors (as also shown earlier in Fig. 2), but also that the accurate parameter inference enables the model to provide accurate predictions for new events over time, both numerically and theoretically.

7.3 Experiments on real data

7.3.1 Data description and metrics

We use two datasets from Twitter and a one from a music platform, containing user activity logs while interacting with different content types. Note that social networks are characterized by user and cascade activity that are heterogeneous and non-stationary, i.e. large differences in overall volume of events per user and per cascade, as well as temporal variations and bursts of the intensity of events. For the real datasets, the hyperparameters β and τ are cross-validated (see Tab. 3).

music2 dataset [14]. It is a registry of Twitter users sharing URLs to Google Play and Youtube (taken as two cascades) of songs during one month of 2015. The dataset is filtered so to keep the $N_u = 93$ users with the highest adoption of these $N_c = 2$ cascades (more than 50 tweets) which amounts to a total of $N_e = 79065$ events. It is noted that the Youtube and Google Play related events of this dataset are neither uncorrelated nor in pure competition. As a matter of fact, beyond the market competition between streaming services, album releases cause a burst of interest on every platform.

url dataset [13]. It contains tweets with URLs transformed by shortening services, commonly used when users share links on Twitter. We selected a (rather) stationary segment of the stream, and then we filtered the event log so as to keep $N_u = 637$ users related to the $N_c = 6$ cascades: `bit.ly`, `migre.me`, `tinyurl.com`, `tiny.ly`, `j.mp`, and `is.gd`, during 10 days of 2010, which amounts to a total of $N_e = 65321$ events.

lastFM dataset [47]. This is a registry of user listenings to music from the LastFM platform and spans over $T = 6$ months. For the sake of visualization, we filter this dataset so as to keep the $N_c = 50$ most popular artists (taken as cascades), as well as $N_u = 559$ users with at least 50 events each in those cascades. The final dataset is then composed of a total of $N_e = 56296$ listening events.

7.3.2 Results

For the music2 dataset (Fig. 4a), we first notice that both linear models, IC and linMIC, have the same learning behavior over different fractions of the test set, though linMIC has a better test-likelihood score. Similarly for CC and MIC, which are equipped with Boltzmann mixing function, where

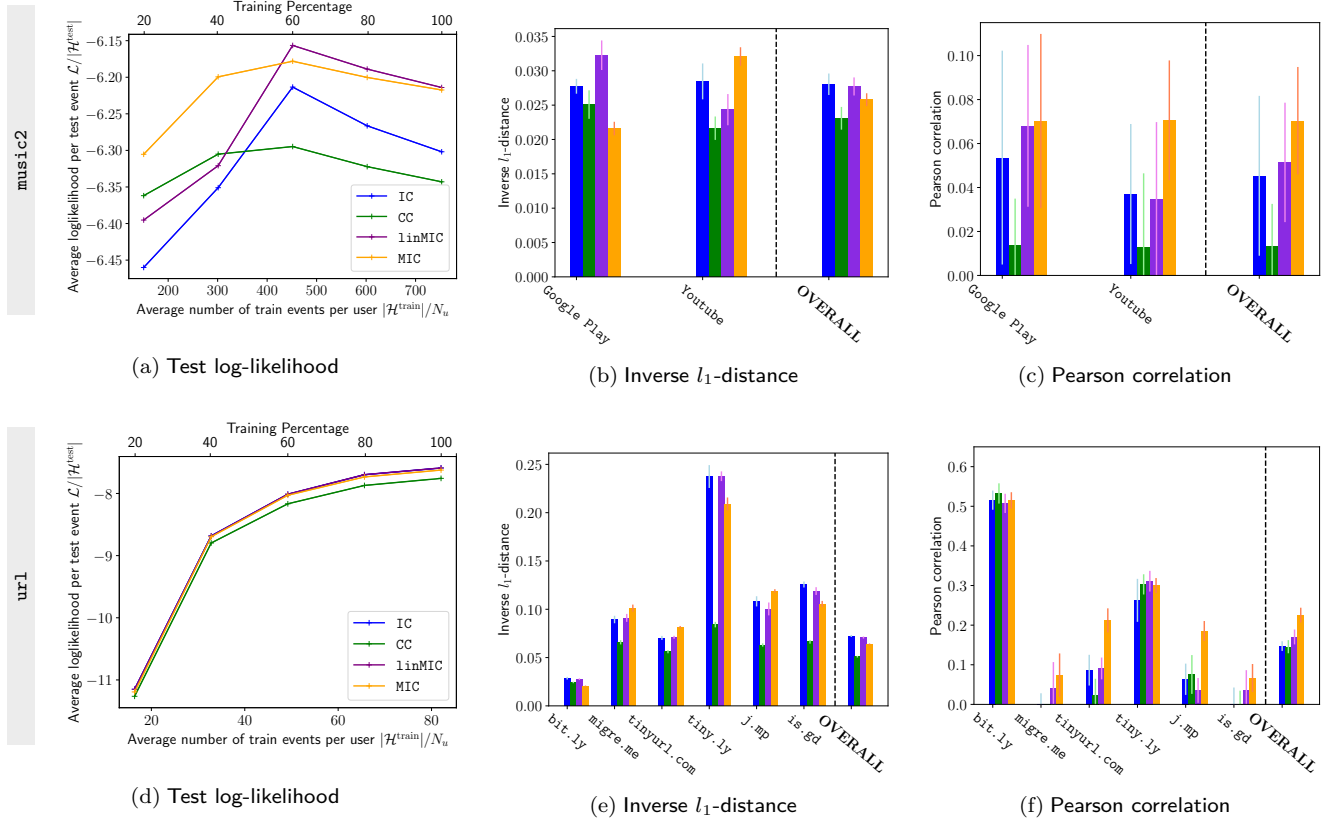


Fig. 4. Evaluation of the compared methods when applied on the **music2** dataset (top row, 2 cascades) and the **url** dataset (bottom row, 6 cascades), using three measures: (a,d) Test log-likelihood for a varying fraction of the initial train dataset (the bottom x-axis shows number of training events per user, and the top x-axis shows the correspondence to percentages). (b,e) Inverse l_1 -distance for the adoption of each cascade, and for the overall intensity. (c,f) Pearson correlation between the real and the simulated intensity of each cascade.

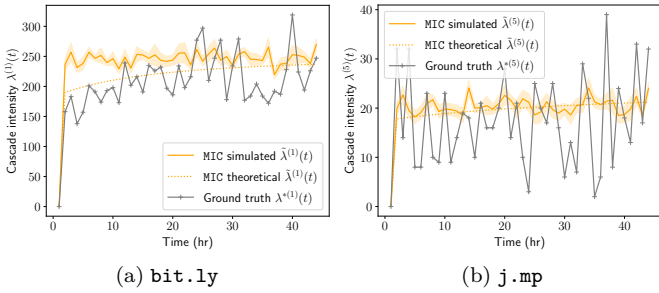


Fig. 5. Temporal evolution of the conditional event intensities of each model when applied on the **url** dataset. For each cascade, a signal is plotted for MIC and is compared to the theoretical expected intensity of the MIC model (Eq. 28) and to test data.

MIC performs better. Next, using the models learned on the training data, we generate event logs and compare them with the real test event logs. In Fig. 4b, we observe that the inverse l_1 -distance is slightly better for linear models, however the Pearson correlation in Fig. 4c suggests that MIC replicates best the real event generation process.

Regarding the **url** dataset, the test log-likelihoods plotted in Fig. 4d indicate that most of the models present the same learning tendency, with CC being slightly inferior. Fig. 4d shows that all models benefit from as the training set grows (note that IC is plotted exactly below linMIC); note that the number of events per user are here smaller than in Fig. 4a,

also this dataset has greater complexity. For the generation of events, Fig. 4e indicates that even if the scores vary notably across the cascades, MIC achieves a better approximation for 3 out of the 6 cascades, and is overall a little below linear models. Nevertheless, the Pearson correlation in Fig. 4f shows that, cascade-wise, MIC is better or comparable to the other models, and overall exhibits a higher correlation to the real data. It is worth noting that IC and CC obtain negative scores for the 2nd and the 6th cascade. To qualitatively assess the performance of the model, we also compare the test intensity to the one sampled from the learned model, as well as the theoretical expected intensity. Fig. 5 shows that the empirical conditional intensities over time of MIC model are in agreement with the real test data, indicating an agreement between our numerical and analytical results in the context of the **url** dataset.

The results for the **lastfm** dataset are displayed in Fig. 6. The test log-likelihood in Nonlinear models appear to have a superior test log-likelihood score than linear ones, with MIC being the model with the best performance. Given the large number of cascades, we omit plotting for this dataset the inverse l_1 -distance and Pearson correlation measures. Nevertheless, the larger diversity of this dataset can provide insights regarding the distribution of cascade popularity shown in Fig. 6b in log-scale. There, MIC exhibits a more accurate distribution of events per cascade compared to the other models, both in terms of number of events generated and in its ordered shape. This figure suggests that none of the models gets sufficiently close to the test data, especially for the most

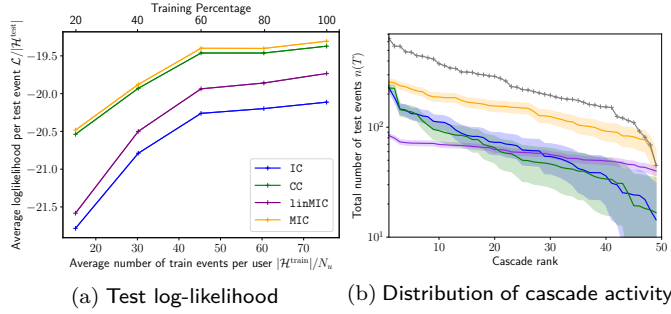


Fig. 6. Evaluation of the compared methods when applied on the `lastfm` dataset (50 cascades) using two measures: (a) Test log-likelihood for a varying size of the initial train dataset (the bottom x-axis shows number of training events per user, and the top x-axis shows the correspondence to percentages). (b) Ranked number of events for each cascade. Real data is compared to the generated events by each of the models.

popular cascades, hence they seem unable to capture the full complexity of this dataset. In fact, for the `lastfm` dataset, the largest event counts mostly originate from exogenous factors, such as the release of a new album or events like concerts.

7.4 Heterogenous user activity in OSNs

We shift our focus to the behavior of the models at the level of individual users. We first compare in Fig. 7 user activity in the real test data to the activity generated by each trained model. The results visually validate the truncated power law behavior of social network activity, which is similarly to what was earlier observed cascade-wise in Fig. 6. For completeness we provide all the plots despite visible differences appear in some of the cases. For instance, there is little to comment about the `music2` dataset where all methods generate roughly the same distribution. However, when we look at cascade-specific distributions from `url` and `lastfm`, MIC manages to replicate much better the distribution of user activity.

Interestingly, when comparing the test log-likelihoods at different quantiles of most active users for the three datasets, we can see in the results reported in Tab. 2 that MIC achieves a better score than the other models. Moreover, the distance to the other models seems to grow as we focus on more active users, which are notably the users for which more can be learned by analyzing their higher activity. These results emphasize the importance of user- and cascade-wise evaluation measures, and gives evidence on MIC’s superior performance in capturing user heterogeneity in social networks. This kind of evaluation constitutes an additional viewpoint on the performance of such models, typically overlooked in the literature. Nevertheless, evaluating further on this direction would require working with more diverse datasets, like `lastfm`.

7.5 Bi-layered visualization of social network activity

This section meets the promise made by the schematic view in Fig. 1, by demonstrating that MIC model can indeed produce such insightful visualizations for real data. Before discussing the results displayed in Fig. 8, we should note that representing social network activity data as a bi-layered network, one layer for the users and another for the cascades, is intuitive, natural, and has been used in prior exploratory studies [48], [49]. The contribution of our work in this aspect stems from the fact that MIC parameters $\Theta = (\Sigma, \mathbf{M}, \mathbf{W})$ can describe




Fig. 7. Ranked number of events per user of the compared methods when applied on the `music2`, `url`, and `lastfm` datasets. The first row shows the event distribution over all the cascades, which are however representative of the datasets. The error bars correspond to the standard deviation of the number of events per user.

TABLE 2
Test log-likelihoods for the top 5%, 10%, and 25% most active users for the three datasets. The best scores are highlighted in gray (smaller values closer to zero are better; draws imply differences $< 10^{-3}$).

		DATASET		
$\mathcal{L}_{\langle \text{model} \rangle}^{\text{test}}$	MODEL	music2	url	lastfm
top 5% users	IC	-3045	-836	-3457
	linMIC	-2967	-836	-3376
	CC	-3073	-838	-3388
	MIC	-2923	-838	-3371
top 10% users	IC	-1841	-640	-2573
	linMIC	-1787	-640	-2511
	CC	-1840	-645	-2505
	MIC	-1771	-643	-2492
top 25% users	IC	-1192	-423	-1636
	linMIC	-1169	-423	-1604
	CC	-1208	-426	-1584
	MIC	-1167	-424	-1578

this structure, therefore such a bi-layered visualization can be computed numerically within our modeling framework by inferring the model parameters for real data.

In the first row of Fig. 8, the `url` and `lastfm` datasets are visualized using a typical bi-layer spring layout based of the real data. The adjacency matrix ($\mathbf{W} = \mathbf{A}$) gives the user graph at top layer, and the cascade activity is at the bottom layer, though lacking structure as this is not provided by the data ($\Sigma = \mathbf{I}$). Node size is proportional to the associated volume of events, by users or cascades, and the node color of a user illustrates the mixture of his activity regarding the cascades.

In the second row, in Fig. 8c and 8d, we provide the bi-layered visualizations computed based on the MIC model, specifically using the respective inferred parameters $\Theta = (\tilde{\Sigma}, \tilde{\mathbf{M}}, \tilde{\mathbf{W}})$. Edges within the cascade (resp. user) layer exist only if they are inferred in $\tilde{\Sigma}$ (resp. $\tilde{\mathbf{W}}$), also are drawn as

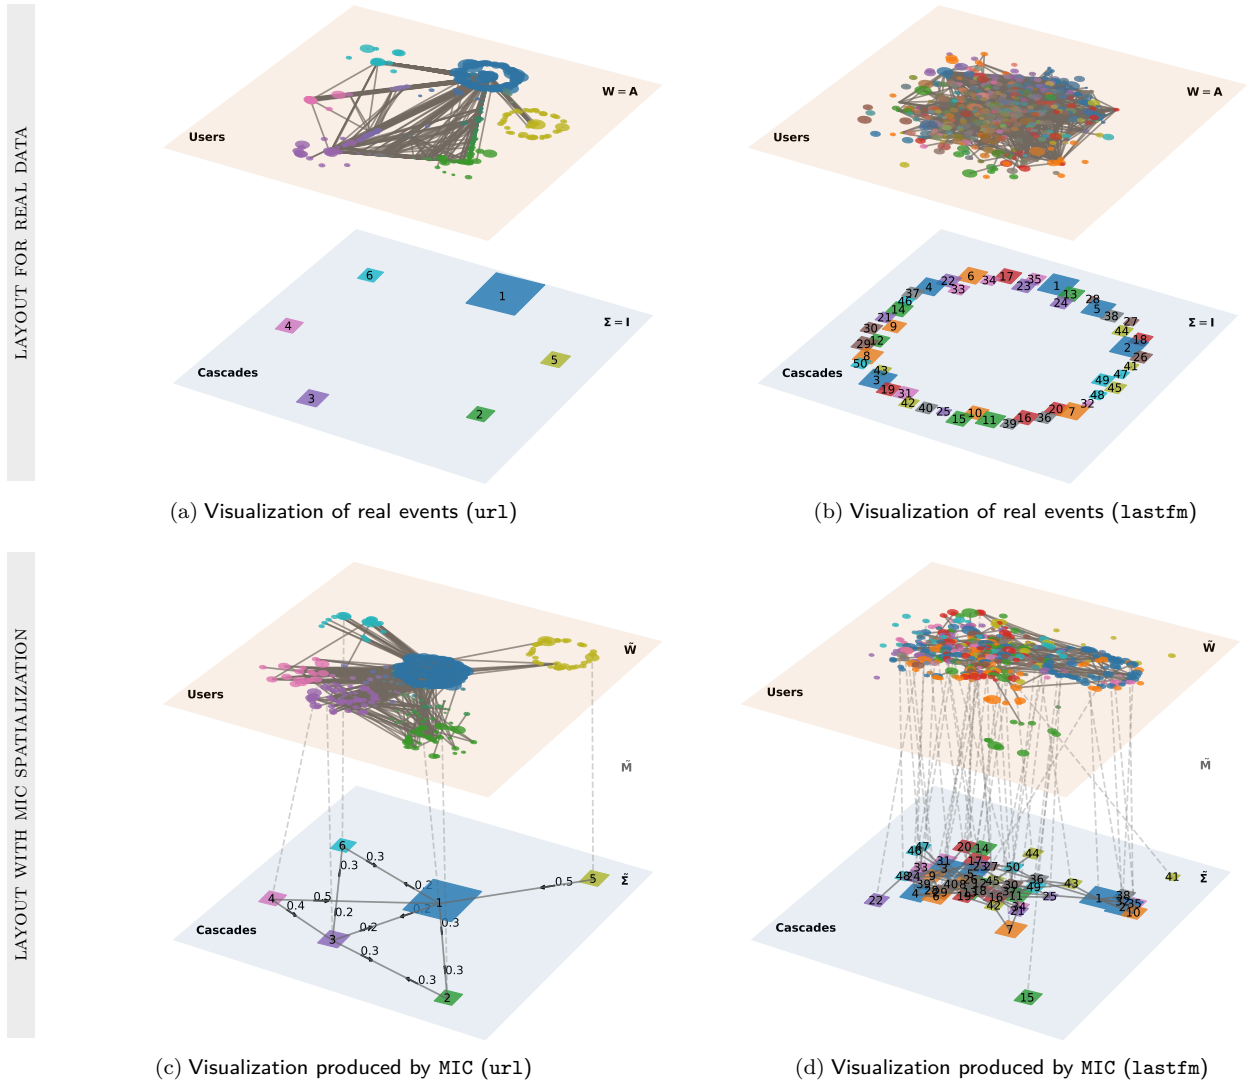


Fig. 8. Bi-layered network visualization of the activity of the `ur1` (left) and `lastfm` (right) datasets. First row: the real data are visualized using a classic two-layer spring layout. Node size is proportional to the associated volume of events, and edges indicate the presence of influence/interactions. User node color illustrates the mixture of activity regarding the cascades. The unweighted adjacency matrix of the user graph is used for the top layers, while in the bottom layers there is no actual graph, so nodes are arranged in a circular layout aligned with the upper layer. Second row: for each visualization, a MIC model is trained on the event log of a dataset, and the layout is based on the inferred parameters $\hat{\Theta} = (\hat{\Sigma}, \hat{M}, \hat{W})$. Edges within the cascade (resp. user) graph are drawn given $\hat{\Sigma}$ (resp. \hat{W}) and edges between layers are drawn according to the highest values of \hat{M} .

being undirected for simplicity. Dashed edges between the two layers correspond to the highest value of the cascade-to-user coupling intensity $\tilde{M}_{(\cdot)}^{(c)}$ for each cascade c . The position of cascades is computed according to their inferred interactions $\tilde{\Sigma}$. The position of each user u is computed based on his activity with regards to the cascades. Specifically, the learned cascade mixing function $f_u(\cdot|t)$ maps user u 's activity to a probability distribution over the cascades, which is then transformed into coordinates in the bottom layer. Fig. 8c visualizes the `ur1` dataset where specific cascades (`bit.ly` and `tinyurl`) appear to be central, as also noted in previous works for this dataset [6]. Links between shortening services are uncovered: for instance, the weighted cascades interactions, depicted with arrows at the bottom layer, can illustrate the varying popularity of URL shortening services shared on Twitter. Additionally, this graphical view highlights the interplay between users and cascades, as it can shed light on

the coupling between network structure and topic alignment, e.g. to reveal clusters of users and cascades that are closely related in terms of content [50]. In fact, for the `ur1` dataset we indeed observe communities of users to be well-organized around specific URL shortening services, and migrating users to be spatially located in transition zones between cascades.

Fig. 8d visualizes a different landscape for the `lastfm` dataset. The higher number of cascades leads to a more dense network structure, with highly central artists (e.g. c_1 and c_3). Nevertheless, the inferred network presents a typical center-periphery-isolated structure, as well as hubs connecting with distinct communities (e.g. c_1). The stronger mixing of cascades through the user social network is most possibly due to the diversity of music tastes and cross-genre listening among users.

8 Conclusion

In this paper, we have presented the *Mixture of Interacting Cascades* (MIC) model, aiming at providing a more accurate modeling tool for the underlying information diffusion mechanisms in OSNs. The use of a marked point process framework ensures an easily interpretable and more refined description of the coupling between information cascades and user activity. Benefiting from the convex and factorization properties of the likelihood, parameter learning can be performed in parallel for each user via Maximum Likelihood Estimation methods. The model is also supported by a theoretical framework that allows the derivation of differential equations of the moments of the number of events and the conditional intensity, which provide analytical insights for the understanding of the underlying information diffusion dynamics.

The presented experiments provided a comprehensive evaluation of the MIC model, using synthetic and real data, several evaluation measures, and from various viewpoints. The design of an ablation study, as well as the comparison with its competitors quantitatively supports the different contributions of MIC. The learned model is also compared to the event generating process from a numerical and analytical perspective. The results suggests that MIC manages to bridge the gap between theoretical and numerical approaches to the problem of modeling intricate patterns of information diffusion. We finally demonstrated that MIC can be used to visualize insightfully social network activity in a bi-layered network. Such representation highlights the interplay between users and cascades, with a focus on cascade interactions that shape the underlying structure of information pathways.

Further investigation can be concentrated on improving the accuracy of the model by refining some modeling choices and the parameter estimation method. For instance, one could consider introducing time-dependent parameters (Θ_t) and non-parametric triggering kernels [4], [36]. Data preprocessing can be beneficial for treating non-stationarity and/or heterogeneity, e.g. using existing methods to detect change-points in Hawkes processes [51], or deal with parameter learning given unbalanced data availability [52]. Additional work could elaborate further the theoretical properties of the model.

Acknowledgments

G.A. was funded by CDSN scholarship from ENS Paris-Saclay. A.K. was supported by the Industrial Data Analytics and Machine Learning Chair hosted at ENS Paris-Saclay. J.R.-F. was funded by CPJ ModSHS at ENS Paris-Saclay.

Appendix

TABLE 3

The cross-validated values for the hyperparameters β and τ , for each model when used on the different real datasets.

HYPERPARAMETER	MODEL	DATASET		
		music2	url	lastfm
Time-decay parameter τ	<i>all</i>	6 min	6 min	10h
competitiveness β	CC	10	5×10^3	10^4
	MIC	5×10^3	10^4	10^4

TABLE 4
Table of symbols.

Symbol	Description
$\mathcal{G}(\mathcal{V}, E)$	Directed network representing the OSN
\mathcal{V}	Set of N_u nodes (users)
\mathbf{A}	Graph adjacency matrix (with binary edge weights)
E	Set of edges
\mathcal{F}_u	Set of followers of user u
\mathcal{C}	Set of N_c information cascades
$\mathbf{n}(t)$	Counting process vector
\mathcal{H}	Sequence of N_e events
$e_j = (u_i, t_i, c_i)$	Event j by user u_i , at time t_i , for cascade c_i
$\mathcal{H}(t)$	Set of events occurring in $[0, t]$
$\mathcal{H}_u(t)$	Subset of events of user u
$\mathcal{H}^{(c)}(t)$	Subset of events of cascade c
$\mathcal{H}_u^{(c)}(t)$	Subset of events of user u and cascade c
$\mathcal{H}_u(t, t')$	Subset of events of user u in the interval $(t, t']$
$\lambda(t)$	Conditional intensity vector
\mathbf{M}	Set of baseline intensities
μ_u	Background intensity of user u
\mathbf{W}	Weight matrix of social influence
w_{vu}	Social influence kernel from v to u
$\kappa(t)$	Temporal kernel
$\nu_u^{(c)}(t)$	Independent intensity of user u for cascade c
$\lambda_u^{(c)}(t)$	Conditional marked intensity of user u for cascade c
$f_u(c t)$	Conditional mark probability function
$n_u^{(c)}(t)$	Number of events of user v for cascade c at time t
Σ	Cascade interaction matrix
σ_{sc}	Influence of cascade s on cascade c
Θ	Set of MIC parameters, $\Theta = (\mathbf{M}, \Sigma, \mathbf{W})$
$\tilde{\Theta}$	Tilde for any inferred model parameters
Θ^*	Star for any true model parameters used to generate synthetic data
$S(t', t)$	Survival function between t and t'
$Z_u(t)$	Partition function of the conditional mark probability
β	Hyperparameter modulating cascade competition
\mathcal{L}	Log-likelihood
\mathcal{L}_u	Partial log-likelihood for user u
$K(T, t_i)$	Integral of the temporal kernel
\mathbf{I}	Identity matrix
$\mathbb{E}[\cdot]$	Expectation operator
$\mathbb{P}(\cdot)$	Probability of an event

References

- [1] L.-X. Yang, X. Yang, and Y. Y. Tang, "A bi-virus competing spreading model with generic infection rates," *IEEE Trans. on Network Science and Engineering*, vol. 5, no. 1, pp. 2–13, 2018.
- [2] H. Hethcote, "The mathematics of infectious diseases," pp. 599–653, 2000.
- [3] E. Bacry, S. Delattre, M. Hoffmann, and J. F. Muzy, "Modelling microstructure noise with mutually exciting point processes," *Quantitative Finance*, vol. 13, no. 1, pp. 65–77, 2013.
- [4] K. Zhou, H. Zha, and L. Song, "Learning triggering kernels for multi-dimensional Hawkes processes," in *Intern. Conf. on Machine Learning*, 2013, pp. 1301–1309.
- [5] L. Weng, A. Flammini, A. Vespignani, and F. Menczer, "Competition among memes in a world with limited attention," *Scientific Reports*, vol. 2, no. 1, p. 335, 2012.
- [6] I. Valera and M. Gomez-Rodriguez, "Modeling adoption and usage of competing products," in *IEEE Intern. Conf. on Data Mining*, 2015, pp. 409–418.
- [7] S. A. Hosseini, K. Alizadeh, A. Khodadadi, A. Arabzadeh, M. Farajtabar, H. Zha, and H. R. Rabiee, "Recurrent Poisson factorization for temporal recommendation," in *ACM SIGKDD Intern. Conf. on Knowledge Discovery and Data Mining*, 2017, pp. 847–855.
- [8] S. A. Myers and J. Leskovec, "The bursty dynamics of the Twitter information network," in *Intern. Conf. on World Wide Web*, 2014, pp. 913–924.
- [9] L. Muchnik, S. Pei, L. C. Parra, S. D. S. Reis, J. S. Andrade Jr, S. Havlin, and H. A. Makse, "Origins of power-law degree

- distribution in the heterogeneity of human activity in social networks,” *Scientific Reports*, vol. 3, no. 1, p. 1783, 2013.
- [10] A. M. Sadri, S. Hasan, S. V. Ukkusuri, and M. Cebrian, “Exploring network properties of social media interactions and activities during Hurricane Sandy,” *Transportation Research Interdisciplinary Perspectives*, vol. 6, p. 100143, 2020.
- [11] D. Notarmuzi, C. Castellano, A. Flammini, D. Mazzilli, and F. Radicchi, “Universality, criticality and complexity of information propagation in social media,” *Nature Communications*, vol. 13, no. 1, p. 1308, 2022.
- [12] J. Goldenberg, B. Libai, and E. Muller, “Talk of the Network: A complex systems look at the underlying process of word-of-mouth,” *Marketing Letters*, vol. 12, pp. 211–223, 2001.
- [13] N. O. Hodas and K. Lerman, “The Simple Rules of Social Contagion,” *Scientific Reports*, vol. 4, no. 1, p. 4343, 2014.
- [14] A. Zarezade, A. Khodadadi, M. Farajtabar, H. Rabiee, and H. Zha, “Correlated Cascades: Compete or cooperate,” *AAAI Conf. on Artificial Intelligence*, pp. 238–244, 2017.
- [15] F. Zhou, X. Xu, G. Trajcevski, and K. Zhang, “A survey of information cascade analysis: Models, predictions, and recent advances,” *ACM Computing Surveys*, vol. 54, no. 2, pp. 1–36, 2022.
- [16] M. G. Rodriguez, D. Balduzzi, and B. Schölkopf, “Uncovering the temporal dynamics of diffusion networks,” in *Intern. Conf. on Intern. Conf. on Machine Learning*, 2011, pp. 561–568.
- [17] S.-H. Yang and H. Zha, “Mixture of mutually exciting processes for viral diffusion,” in *Intern. Conf. on Machine Learning*, 2013, pp. 1–9.
- [18] M. Farajtabar, Y. Wang, M. Gomez-Rodriguez, S. Li, H. Zha, and L. Song, “COEVOLVE: A joint point process model for information diffusion and network evolution,” *Journal of Machine Learning Research*, vol. 18, no. 41, pp. 1–49, 2017.
- [19] T. Iwata, A. Shah, and Z. Ghahramani, “Discovering latent influence in online social activities via shared cascade poisson processes,” in *ACM SIGKDD Intern. Conf. on Knowledge Discovery and Data Mining*, 2013, pp. 266–274.
- [20] S. Linderman and R. Adams, “Discovering latent network structure in point process data,” in *Intern. Conf. on Machine Learning*, 2014, pp. 1413–1421.
- [21] Y. A. Isik, C. Davis, P. Chapfuwa, and R. Henao, “Flexible triggering kernels for Hawkes process modeling,” 2022.
- [22] M. H. Degroot, “Reaching a consensus,” *Journal of the American Statistical Association*, vol. 69, no. 345, pp. 118–121, 1974.
- [23] G. Verma, A. Swami, and K. Chan, “The impact of competing zealots on opinion dynamics,” *Physica A: Statistical Mechanics and its Applications*, vol. 395, pp. 310–331, 2014.
- [24] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch, “Mixing beliefs among interacting agents,” *Advances in Complex Systems*, vol. 03, no. 01n04, pp. 87–98, 2000.
- [25] H. Rainer and U. Krause, “Opinion Dynamics and Bounded Confidence: Models, Analysis and Simulation,” *Journal of Artificial Societies and Social Simulation*, vol. 5, no. 3, 2002.
- [26] F. Baumann, P. Lorenz-Spreen, I. M. Sokolov, and M. Starnini, “Emergence of polarized ideological opinions in multidimensional topic spaces,” *Physical Review X*, vol. 11, no. 1, p. 011012, 2021.
- [27] I. Conjeaud, P. Lorenz-Spreen, and A. Kalogeratos, “Degroot-based opinion formation under a global steering mechanism,” *IEEE Trans. on Computational Social Systems*, vol. 11, no. 3, pp. 4040–4057, 2024.
- [28] C. Tan, L. Lee, J. Tang, L. Jiang, M. Zhou, and P. Li, “User-level sentiment analysis incorporating social networks,” in *ACM SIGKDD Intern. Conf. on Knowledge Discovery and Data Mining*, 2011, pp. 1397–1405.
- [29] R. West, H. S. Paskov, J. Leskovec, and C. Potts, “Exploiting social network structure for person-to-person sentiment analysis,” *Trans. of the Association for Computational Linguistics*, vol. 2, pp. 297–310, 2014.
- [30] S. Liu, H. Shen, H. Zheng, X. Cheng, and X. Liao, “CT LIS: Learning influences and susceptibilities through temporal behaviors,” *ACM Trans. on Knowledge Discovery from Data*, vol. 13, no. 6, pp. 1–21, 2019.
- [31] D. Chavalarias, P. Bouchaud, and M. Panahi, “Can a Single Line of Code Change Society? The Systemic Risks of Optimizing Engagement in Recommender Systems on Global Information Flow, Opinion Dynamics and Social Structures,” *Journal of Artificial Societies and Social Simulation*, vol. 27, no. 1, p. 9, 2024.
- [32] A. De, I. Valera, N. Ganguly, S. Bhattacharya, and M. Gomez Rodriguez, “Learning and forecasting opinion dynamics in social networks,” in *Advances in Neural Information Processing Systems*, vol. 29, 2016.
- [33] Y. Wang, E. Theodorou, A. Verma, and L. Song, “A stochastic differential equation framework for guiding online user activities in closed loop,” in *Intern. Conf. on Artificial Intelligence and Statistics*, vol. 84, 2018, pp. 1077–1086.
- [34] A. G. Hawkes, “Spectra of some self-exciting and mutually exciting point processes,” *Biometrika*, vol. 58, no. 1, pp. 83–90, 1971.
- [35] Thomas Liniger, “Multivariate Hawkes processes,” Ph.D. dissertation, ETH Zurich, 2009.
- [36] R. Lemonnier, K. Scaman, and A. Kalogeratos, “Multivariate Hawkes processes for large-scale inference,” *AAAI Conf. on Artificial Intelligence*, vol. 31, no. 1, 2017.
- [37] Q. Cao, H. Shen, K. Cen, W. Ouyang, and X. Cheng, “Deep-Hawkes: Bridging the gap between prediction and understanding of information cascades,” in *ACM Conf. on Information and Knowledge Management*, 2017, pp. 1149–1158.
- [38] G. Poux-Médard, J. Velcin, and S. Loudcher, “Multivariate powered Dirichlet-Hawkes process,” in *Advances in Information Retrieval*, J. Kamps, L. Goeriot, F. Crestani, M. Maistro, H. Joho, B. Davis, C. Gurrin, U. Kruschwitz, and A. Caputo, Eds. Springer Nature Switzerland, 2023, pp. 47–61.
- [39] S. Zhu and Y. Xie, “Spatiotemporal-textual point processes for crime linkage detection,” *The Annals of Applied Statistics*, vol. 16, no. 2, 2022.
- [40] D. Daley and D. Vere-Jones, *An Introduction to the Theory of Point Processes*, ser. Probability and Its Applications. New York: Springer-Verlag, 2003.
- [41] R. Zhang, C. Walder, M.-A. Rizozi, and L. Xie, “Efficient non-parametric Bayesian Hawkes processes,” *Preprint arXiv:1810.03730*, 2022.
- [42] A. Nayak, S. Hosseinalipour, and H. Dai, “Smart information spreading for opinion maximization in social networks,” in *IEEE Conf. on Computer Communications*, 2019, pp. 2251–2259.
- [43] O. Aalen, O. Borgan, and H. Gjessing, *Survival and Event History Analysis: A Process Point of View*. Springer, 2008.
- [44] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [45] L. Cui, A. Hawkes, and H. Yi, “An elementary derivation of moments of Hawkes processes,” *Advances in Applied Probability*, vol. 52, pp. 102–137, 2020.
- [46] Y. Ogata, “On Lewis’ simulation method for point processes,” *IEEE Trans. on Information Theory*, vol. 27, no. 1, pp. 23–31, 1981.
- [47] O. C. Herrada, *Music Recommendation and Discovery in the Long Tail*. Springer, 2010.
- [48] M. De Domenico, M. A. Porter, and A. Arenas, “MuxViz: A tool for multilayer analysis and visualization of networks,” *Journal of Complex Networks*, vol. 3, no. 2, pp. 159–176, 2015.
- [49] A. P. Logan, P. M. LaCasse, and B. J. Lunday, “Social network analysis of Twitter interactions: A directed multilayer network approach,” *Social Network Analysis and Mining*, vol. 13, no. 1, p. 65, 2023.
- [50] L. M. Aiello, A. Barrat, R. Schifanella, C. Cattuto, B. Markines, and F. Menczer, “Friendship prediction and homophily in social media,” *ACM Trans. on the Web*, vol. 6, no. 2, pp. 1–33, 2012.
- [51] Z. Zhang, X. LU, and F. Zhou, “Conjugate Bayesian two-step change point detection for hawkes process,” in *Advances in Neural Information Processing Systems*, 2024.
- [52] F. Salehi, W. Trouleau, M. Grossglauser, and P. Thiran, “Learning Hawkes processes from a handful of events,” in *Advances in Neural Information Processing Systems*, vol. 32, 2019.