

Multivariate Spatiotemporal Hawkes Processes and Network Reconstruction*

Baichuan Yuan[†], Hao Li[†], Andrea L. Bertozzi[†], P. Jeffrey Brantingham[‡], and
Mason A. Porter[†]

Abstract. There is often latent network structure in spatial and temporal data, and the tools of network analysis can yield fascinating insights into such data. In this paper, we develop a nonparametric method for network reconstruction from spatiotemporal data sets using multivariate Hawkes processes. In contrast to prior work on network reconstruction with point-process models, which has often focused on exclusively temporal information, our approach uses both temporal and spatial information and does not assume a specific parametric form of network dynamics. This leads to an effective way of recovering an underlying network. We illustrate our approach using both synthetic networks and networks that we construct from real-world data sets (a location-based social-media network, a narrative of crime events, and violent gang crimes). Our results demonstrate that, in comparison to using only temporal data, our spatiotemporal approach yields improved network reconstruction, providing a basis for meaningful subsequent analysis—such as examinations of community structure and motifs—of the reconstructed networks.

Key words. multivariate Hawkes processes, community structure, spatiotemporal data, social networks, motifs

AMS subject classifications. 60G55, 62H11, 91D30

DOI. 10.1137/18M1226993

1. Introduction. Digital devices such as smartphones and tablets generate an enormous amount of spatiotemporal data about human activities, providing a wonderful opportunity for researchers to gain insight into human dynamics through our “digital footprints”. A wide variety of human activities are now analyzed using such data, creating new disciplines such as computational social science and digital humanities [38]. Examples of such activities include online check-ins in large cities [15], human mobility [6] and currency flow [11], online communications during Occupy Wall Street [16], crime reports in Los Angeles [33], and many others.

Network analysis is a powerful approach for representing and analyzing complex systems of interacting components [51], and network-based methods can provide important insights into the structure and dynamics of complex spatiotemporal data [8]. It has been valuable for studies of both digital human footprints and human mobility [7]. To give one recent example, Noulas et al. [53] studied geographic online social networks to illustrate similarities and heterogeneities in human mobility patterns.

*Received by the editors November 15, 2018; accepted for publication (in revised form) April 30, 2019; published electronically June 27, 2019.

<http://www.siam.org/journals/simods/1-2/M122699.html>

Funding: This work is supported by NSF grants DMS-1417674 and DMS-1737770, DARPA grant FA8750-18-2-0066, and NIJ fellowship 2018-R2-CX-0013.

[†]Department of Mathematics, University of California, Los Angeles, Los Angeles, CA 90095 (byuan@math.ucla.edu, lihao0809@math.ucla.edu, bertozzi@math.ucla.edu, mason@math.ucla.edu).

[‡]Department of Anthropology, University of California, Los Angeles, Los Angeles, CA 90095 (branting@ucla.edu).

Suppose that each node in a network represents an entity, and that the edges (which can be either undirected or directed, and can be either unweighted or weighted) represent spatiotemporal connections between pairs of entities. For instance, in a data set of check-ins on a social-media platform, one can model each user as a node, which has associated check-in times and locations. In this case, one can suppose that an edge exists between a pair of users if they follow each other on the platform. One can use edge weights to quantify the amount of “influence” between users, where a larger weight signifies a larger impact. In our investigation, we assume that the relationships between nodes are time-independent.¹ In some cases, the entities and relationships are both known, and one can investigate the structure and dynamics of the associated networks. However, in many situations, network data are incomplete—with potentially a large amount of missing data in the form of missing entities, interactions, and/or metadata [69]—and it may not be possible to directly observe the relationships between nodes [63]. For example, social-media companies attempt to infer friendship relationships between their users to provide accurate friendship recommendations for online social networks.

In the last few years, there has been much effort on inferring missing data (both structure and weights) in networks. A basic approach for inferring relationships among entities is to calculate cross-correlations of their associated time series [37]. Another approach is to use coefficients from a generalized linear model (GLM) [49], a generalization of linear regression that allows response variables to have a non-Gaussian error-distribution. Recently, researchers have begun to use point-process methods [66] in network reconstruction. For example, Perry and Wolfe [58] examined the modeling of a network as a multivariate point process and then inferred covariate-based edges (both their existence and their weights) by estimating a point process. One well-studied family of point-process models is Hawkes processes (also known as “self-exciting point processes”),² which have been employed often for studying human dynamics [23, 40]. Hawkes-process models are characterized by mutual “triggering” among events [55], as one event may increase the probability for subsequent events to occur. Such models can capture inhomogeneous interevent times and causal (temporal) correlations, which are important considerations for human dynamics [32]. These properties illustrate the relevance of using Hawkes processes in social-network applications [30]. It thus seems promising to employ such processes for network inference on dynamic human data, such as crime events and online social activity. For example, Linderman and Adams [40] proposed a fully Bayesian Hawkes model that they reported to be more accurate at inferring missing edges for their data than GLMs, cross-correlations, and a simple self-exciting point process with an exponential kernel. Very recently, self-exciting point processes were applied in [71] to reconstruct multilayer networks [31], a generalization of ordinary graphs. However, despite the successes of these studies, the aforementioned temporal point-process models are not without limitations. For example, most of these models do not use spatial information, even when it plays a significant role in a system’s dynamics. Furthermore, many studies assume an a priori model [40] or a specific parametrization [72] for their point processes.

¹Depending on the relative time scales between spatiotemporal processes and network dynamics, it may be important to consider time-dependent edges [28, 59].

²We use these terms interchangeably in this paper.

In the present paper, we propose a nonparametric and multivariate version of a spatiotemporal Hawkes process. Spatiotemporal Hawkes processes have been used previously to study numerous topics, including crime [46], social media [36], and earthquake forecasting [22]. In our model, each node in a network is associated with a spatiotemporal Hawkes process. The nodes can “trigger” each other, so events that are associated with one node increase the probability that there will be events associated with the other nodes. We measure the extent of such mutual-triggering effects using a $U \times U$ “triggering matrix” $\mathbf{K}$, where U is the number of nodes. If one considers an exclusively temporal scenario, a point process u does not “cause” (in the Granger sense [24]) a point process v if and only if $\mathbf{K}(u, v) = 0$ [20]. Because triggering between point processes reflects an underlying connection, one can try to recover latent relationships in a network from $\mathbf{K}$. Such triggering should decrease with both distance and time according to spatial and temporal kernels. In our work, instead of assuming exponential decay [23] or some other distribution [40, 72], we adopt a nonparametric approach [43] to learn both spatial and temporal kernels from data using an expectation-maximization-type (EM-type) algorithm [75]. Recently, Chen et al. [14] studied Hawkes processes with both mutual excitations and inhibitory relationships using a nonparametric approach, although they only considered exclusively temporal kernels.

Our paper helps fill a gap in the literature on incorporating spatial information into multivariate self-exciting point processes [62]. To the best of our knowledge, it is the first paper that uses a multivariate spatiotemporal Hawkes process with a nonparametric method to estimate a triggering kernel. Inspired by the successful employment of spatiotemporal univariate Hawkes processes in earthquake forecasting [22, 43] and predictive policing [47], our work extends these ideas to multivariate Hawkes processes and uses these ideas in an application to network reconstruction. Our approach outperforms other recent point-process network-reconstruction methods [23, 40] on both synthetic and real-world data sets with spatial information. Additionally, our results illustrate the importance both of incorporating spatial information and of using nonparametric kernels. Although we assume that the relationships between nodes are time-independent, our model still recovers a causal structure among events in synthetic data sets. Based on this information, we build event-causality networks on data sets about violent crimes of gangs and examine gang-retaliation patterns using motif analysis.

Our paper proceeds as follows. In section 2, we review self-exciting point processes and recent point-process methods for network reconstruction. In section 3, we introduce our nonparametric spatiotemporal model and our approaches for model estimation and simulation. In section 4, we compare our model with other models using both synthetic and real-world data sets. We construct our two examples of the latter from (1) a location-based social-media platform and (2) crime topics. We conclude in section 5, and we give details about our preprocessing for the social-media data set in Appendix A.

2. Self-exciting point processes. Given a complete, separable metric space, a *point process* S is a random measure that has values in $\{0, 1, 2, \dots\} \cup \{\infty\}$ [64]. We first consider a *temporal point process*, which consists of a list $\{t_1, t_2, \dots, t_N\}$ of N time points, with corresponding events $1, 2, \dots, N$. Let $S[a, b)$ denote the number of points (i.e., events) that occur in a finite time interval $[a, b)$, with $a < b$. One typically models the behavior of a simple temporal point process (multiple events cannot occur at the same time) by specifying its

conditional intensity function $\lambda(t)$, which represents the rate at which events are expected to occur around a particular time t , conditional on the prior history of the point process before time t . Specifically, when $H_t = \{t_i | t_i < t\}$ is the history of the process up to time t , one defines the *conditional intensity function*

$$\lambda(t) = \lim_{\Delta t \downarrow 0} \frac{\mathbb{E}[S[t, t + \Delta t] | H_t]}{\Delta t}.$$

One important point-process model is a *Poisson process*, in which the number of points in any time interval follows a Poisson distribution and the numbers of points in disjoint sets are independent of each other. A Poisson process is called *homogeneous* if $\lambda(t) \equiv \text{constant}$, and it is thus characterized by a constant rate at which events are expected to occur per unit time. It is called *inhomogeneous* if the conditional intensity function $\lambda(t)$ depends on the time t (e.g., $\lambda(t) = e^{-t}$). In both situations, the numbers of points (i.e., events) in disjoint intervals are independent random variables.

We now discuss self-exciting point processes, which allow one to examine a notion of causality in point-process models. If we consider a list $\{t_1, t_2, \dots, t_N\}$ of time stamps, we say that a point process is *self-exciting* if

$$\text{Cov}[S[t_{k-1}, t_k), S[t_k, t_{k+1}]] > 0, \quad \text{with } t_{k-1} < t_k < t_{k+1},$$

where k is a positive integer. In a self-exciting point process, if an event occurs, another event becomes more likely to occur locally in time.

A *univariate temporal Hawkes process*, which we express using the common cluster representation [27], has the following conditional intensity function:

$$(2.1) \quad \lambda(t) = \mu(t) + K \sum_{t_k < t} g(t - t_k),$$

where the background rate $\mu(t)$ can be either a constant or a time-dependent function that describes how the likelihood of some process (crimes, e-mails, tweets, and so on) evolves in time. For example, violent crimes are more likely to occur at night than during the day, and business e-mails are less likely to be sent during the weekend than on a weekday. One can construe the rate $\mu(t)$ as a process that designates the likelihood of an event to occur, independent of the other events. The summation term in (2.1) describes the self-excitation: past events increase the current conditional intensity. The function $g(t) \geq 0$ is called the *triggering kernel*, and the parameter $K \geq 0$ denotes the mean number of events that are triggered by an event. One standard example is a Hawkes process with an exponential kernel $g(t) = \omega e^{-\omega t}$, where the constant decay rate ω for the triggering kernel controls how fast the rate $\lambda(t)$ returns to its baseline level $\mu(t)$ after an event occurs.

2.1. Temporal multivariate models. In network reconstruction, one seeks to infer the relationships (i.e., edges) and the strengths of such relationships (i.e., edge weights) among a set of entities (i.e., nodes). When modeling the relationships in a network, it is more appropriate to use a multivariate point process than a univariate one. In a temporal multivariate point process, there are U different point processes $(S_u)_{u=1, \dots, U}$; and the corresponding conditional intensity functions are $(\lambda_u(t))_{u=1, \dots, U}$. We seek to infer the intensity functions from

observed data $(t_j, u_j)_{j=1,\dots,N}$ in a time window $[0, T]$, where t_j and u_j , respectively, are the time and point-process indices of event j . There are numerous applications of temporal multivariate point processes; they include financial markets [5], real-time crime forecasting [76], and neuronal spike trains [12]. In the present paper, we focus on the application of network reconstruction.

Let's first consider two examples of multivariate processes that are not self-exciting. A trivial example of a multivariate point process is the multivariate Poisson process, in which each point process is a univariate Poisson process. Another example is the multivariate Cox process, which consists of doubly stochastic Poisson processes (so the conditional intensity itself is a stochastic process). Perry and Wolfe [58] used a Cox process to model e-mail interactions (the edges) among a set of users (the nodes).

Instead of modeling edges as Cox processes, Fox et al. [23] used multivariate Hawkes processes to model people (i.e., nodes) communicating with each other via e-mail. Their conditional intensity function has an exponential kernel and a nonparametric background function $\mu_u(t) \geq 0$ for each person (i.e., process) u . It is written as

$$(2.2) \quad \lambda_u(t) = \mu_u(t) + \sum_{t_i < t} K_{u_i u} \omega e^{-\omega(t-t_i)},$$

where $K_{uv} = \mathbf{K}(u, v) \geq 0$ is the expected number of events of person v that are triggered by one event of person u . One can estimate the set Θ of parameters by minimizing the negative log-likelihood function

$$(2.3) \quad -\log(L(\Theta)) = -\sum_{k=1}^N \log(\lambda_{u_k}(t_k)) + \sum_{u=1}^U \int_0^T \lambda_u(t) dt,$$

where $\log(x)$ denotes the natural logarithm of x . Recall that u_k is the point process associated with event k .

There are several variants of a multivariate Hawkes process. One is to add regularization terms to (2.3) to improve the accuracy of parameter estimation. Lewis and Mohler [39] used maximum-penalized likelihood estimation, which enforces some regularity on the model parameters, to infer Hawkes processes. Zhou, Zha, and Song [77] extended this idea and promoted the low-rank and sparsity properties of $\mathbf{K}$ by adding nuclear and L_1 norms of $\mathbf{K}$ to (2.3) with the conditional intensity function $\lambda_u(t)$ from (2.2). Linderman and Adams [40] added random-graph priors on $\mathbf{K}$ and developed a fully Bayesian multivariate Hawkes model. See [42] for theoretical guarantees on inferring Hawkes processes with a regularizer. Another research direction is to speed up the parameter estimation of point-process models. For example, Hall and Willett [26] tried to learn the triggering matrix $\mathbf{K}$ via an online learning framework for streaming data. In a very recent paper, Achab et al. [2] developed a fast moment-matching method (instead of using a likelihood-based method) to estimate the matrix $\mathbf{K}$. A recent paper by Chen et al. [14] extended temporal multivariate Hawkes processes to study inhibitory events, in addition to the usual excitatory ones. (For example, the spiking of one neuron can decrease the firing rates of other neurons.) They developed a new theoretical approach that exploits a thinning-process representation [54], instead of the common cluster representation [27], of temporal Hawkes processes. Using their approach, they derived a

concentration inequality for the second-order statistics of temporal Hawkes processes with inhibition and proposed a nonparametric model for network reconstruction.

2.2. Spatiotemporal point processes. Many real-world data sets include not only time stamps but also accompanying spatial information, which can be particularly important for correctly inferring and understanding the dynamics that are associated with such data [8]. In earthquakes, for example, most aftershocks usually occur geographically near the main shock [56]. In online social media, if two individuals often check in at the same location at closely proximate times, there is more likely to be a connection between them than if such “joint check-ins” occur rarely [15]. These situations suggest that it is important to examine spatiotemporal point processes, rather than just temporal ones. Indeed, there are myriad applications of spatiotemporal Hawkes processes, including crime forecasting [46], detection of anomalous seismicity [56], and inference of Twitter topics [36].

We characterize a spatiotemporal point process with an event-counting measure S (which is defined on a spatiotemporal region) via its conditional intensity $\lambda(t, x, y)$, which is the expected rate of the accumulation of points around a particular spatiotemporal location. Given the history $\mathcal{H}_t$ of all points up to time t , we write

$$\lambda(t, x, y) = \lim_{\Delta t, \Delta x, \Delta y \downarrow 0} \left(\frac{\mathbb{E}[S\{[t, t + \Delta t) \times [x, x + \Delta x) \times [y, y + \Delta y)\} | \mathcal{H}_t]}{\Delta t \Delta x \Delta y} \right).$$

For the purpose of modeling earthquakes, Ogata [56] used a self-exciting point process with a conditional intensity of the form

$$\lambda(t, x, y) = \mu(x, y) + K \sum_{t > t_i} g(x - x_i, y - y_i, t - t_i).$$

In this setting, if an earthquake occurs, aftershocks are more likely to occur locally in time and space. The choice of the triggering kernel $g(t, x, y)$ is inspired by physical properties of earthquakes. For example, Ogata [56] used a modified Omori formula (a power law) [55] to describe the frequency of aftershocks per unit time. In sociological applications, there is no direct theory to indicate appropriate choices for the kernel function. Some researchers have chosen specific kernels (e.g., exponential kernels) that are easy to compute. For example, Tita et al. [72] used a spatiotemporal point process to infer missing information about event participants. They modeled interactions between event participants as a combination of a spatial Gaussian mixture model and a temporal Hawkes process with an exponential kernel. A key problem is how to justify kernel choices in specific applications.

3. Spatiotemporal models for network reconstruction. Many network-reconstruction methods, such as the ones in [14, 23, 40], have used self-exciting point processes to infer time-independent relationships (i.e., edges) between entities (i.e., nodes) with corresponding (exclusively) temporal point processes. Entity (i.e., process) u is adjacent to entity v if $\mathbf{K}(u, v) > 0$, where one estimates the triggering matrix $\mathbf{K}$ from data. Entity u is not adjacent to v if the former’s point process does not cause the latter’s point process in time (in the Granger sense [20]). For many problems, it is desirable (or even crucial) to incorporate spatial information [8, 17]. For example, spatial information is an important part of online

fingerprints in human activity, and it has a significant impact on most other social networks. In crime modeling, for instance, there is a “near repeat” phenomenon in crime locations, indicating the necessity of including spatial information. Specifically, the spatial neighborhood of an initial burglary has a higher risk of repeat victimization than more-distant locations [65]. In our work, we propose multivariate spatiotemporal Hawkes processes to infer relationships in networks and provide a novel approach for analyzing spatiotemporal dynamics.

It is also important to consider the assumptions on triggering kernels for Hawkes processes. In seismology, for example, researchers attempt to use an underlying physical model to help determine a good kernel. However, it is much more difficult to validate such models in social networks than for physical or even biological phenomena [60]. The content of social data is often unclear, and typically there is little understanding of the underlying mechanisms that produce them. With less direct knowledge of possible triggering kernels, it is helpful to employ a data-driven approach for kernel selection. Using a kernel with an inappropriate decay rate may lead to either underestimation or overestimation of the elements in the triggering matrix $\mathbf{K}$, which may also include false negatives or false positives in the inferred relationships between entities.

Therefore, we ultimately choose to use a nonparametric approach to learn triggering kernels in various applications to avoid a priori assumptions about a specific parametrization. Specifically, we use histogram estimators with EM-type algorithms to maximize the likelihood, as has been done in applications in seismology and crime modeling [39, 43, 75]. An alternative approach [14] is a penalized regression scheme. With such a scheme, one can approximate a kernel as a sum of basis functions and minimize the squared-error loss of the intensity function with a group-lasso penalty. Although some of the goals of previous Hawkes-process models and our paper are similar, there are many key differences. For example, the authors [14] focused on theoretical development of Hawkes processes with inhibition; they did not consider spatial information. By contrast, the purpose of our paper is to investigate spatiotemporal data sets from social media and crime with a self-exciting Hawkes process. The theoretical guarantees of our model arise from the consistency of the Hawkes process with the nonnegative triggering matrix and kernels (which is necessary for the cluster representation of such processes).

A multivariate spatiotemporal Hawkes process is a sequence $\{(t_i, x_i, y_i, u_i)\}_{i=1}^N$ with N events, where t_i and (x_i, y_i) are spatiotemporal stamps and u_i is the point-process index of event i . Each of the U nodes is a *marginal process*. The conditional intensity function for node u is

$$(3.1) \quad \lambda_u(t, x, y) = \mu_u(x, y) + \sum_{t > t_i} K_{u_i u} g(x - x_i, y - y_i, t - t_i).$$

The above Hawkes process assumes that each node u has a background Poisson process that is constant in time but inhomogeneous in space with conditional intensity $\mu_u(x, y)$. There is also self-excitation, as past events increase the likelihood of subsequent events. We quantify the impact that events associated with node u_i have on subsequent events of node u_j with spatiotemporal kernels and the element $\mathbf{K}(u_i, u_j) = K_{u_i u_j}$ of the triggering matrix.

3.1. A parametric model. We first propose a multivariate Hawkes process with a specific parametric form. We use this model to generate spatiotemporal events on synthetic networks and provide a form of “ground truth” that we can use later.

The background rate μ_u and the triggering kernel g for (3.1) are given by

$$g(x, y, t) = g_1(t) \times g_2(x, y) = \omega \exp(-\omega t) \times \frac{1}{2\pi\sigma^2} \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right),$$

$$\mu_u(x, y) = \sum_{i=1}^N \frac{\beta_{u_i u}}{2\pi\eta^2 T} \times \exp\left(-\frac{(x - x_i)^2 + (y - y_i)^2}{2\eta^2}\right).$$

For simplicity, we use exponential decay in time [55] and a Gaussian kernel in space [45]. We let T denote the time window of a data set; $K_{u_i u}$ denote the mean number of the events in process u that are triggered by each event in process u_i ; the quantity $\beta_{u_i u}$ denotes the extent to which events in process u_i contribute to the background rate for events in process u ; and σ and η , respectively, denote the standard deviations in the triggering kernel and background rate. The value of σ determines the spreading scale of the triggering in space.

3.2. A nonparametric model. With the conditional intensity in (3.1), we estimate the triggering kernel $g(x, y, t) = g_1(t) \times g_2(x, y)$ nonparametrically using histogram estimators [43]. We assume that g_2 is isotropic, which entails that $g_2(x, y) = g_2(r)$, where $r = \sqrt{x^2 + y^2}$. We let $h(r)$ be the spatial triggering kernel in the radial coordinate: $h(r) = 2\pi r g_2(r)$. We extend the background rate that was proposed in [22] to the multivariate case and write

$$(3.2) \quad \mu_u(x, y) = \gamma_u \tau(x, y) = \frac{\gamma_u}{T} \sum_{i=1}^N \frac{p_{ii}}{2\pi d_i^2} \exp\left(-\frac{(x - x_i)^2 + (y - y_i)^2}{2d_i^2}\right),$$

where γ_u is a scaling factor for the process u and p_{ii} is the probability that event i is a background event (i.e., it is not triggered by any event). We compute d_i by determining the radius of the smallest disk centered at (x_i, y_i) that includes at least n_p other events and is at least as large as some small value ϵ (which represents the error in location).

Once we fit the model to spatiotemporal data, the triggering matrix $\mathbf{K}$ gives our inferences for the underlying relationships between entities. For two entities u and v , the matrix element $\mathbf{K}(u, v)$ indicates a mixture of temporal causality and spatial dependence between them. In inferring latent relationships in a network, we assume that entity u is not related to v if $\mathbf{K}(u, v) = 0$. We threshold the matrix $\mathbf{K}$ at a certain level: we set elements that are smaller than the threshold value to 0; and we either maintain the values of larger or equal elements to obtain a weighted network, or we set them to 1 to produce an unweighted network. We use $\tilde{\mathbf{K}}$ to denote the thresholded matrix $\mathbf{K}$. We interpret that there is no relation between two nodes u and v if $\tilde{\mathbf{K}}(u, v) = \tilde{\mathbf{K}}(v, u) = 0$.

3.3. Model estimation. We use an EM-type algorithm [75] to estimate the parameters and kernel functions of our model. This EM-type algorithm gives an iterative method to find maximum-likelihood estimates of the parameters. We assume that the original model depends on unobservable latent variables. Suppose that we have data X and want to estimate parameters Θ . One can view the likelihood function $L(\Theta; X)$ as the marginal likelihood function of $L(\Theta; Y, X)$, where Y is a latent variable. We call $L(\Theta; Y, X)$ the “complete-data likelihood function” and $L(\Theta; X)$ the “incomplete-data likelihood function”. Because both

Y and $L(\Theta; Y, X)$ are random variables, we cannot estimate them directly. Therefore, we consider the following expectation function:

$$(3.3) \quad \begin{aligned} Q(\Theta, \Theta^{i-1}) &= \mathbb{E} [\log(L(\Theta; Y, X)) | X, \Theta^{i-1}] \\ &= \int \log(L(\Theta; Y, X)) f(Y | X, \Theta^{i-1}) dY, \end{aligned}$$

where $f(Y | X, \Theta^{i-1})$ is the probability density function of Y , given the data X and Θ^{i-1} . We update parameters by solving the following equation:

$$\hat{\Theta}^i = \arg \max_{\Theta} Q(\Theta, \Theta^{i-1}).$$

Algorithm 3.1. EM-type algorithm for the parametric spatiotemporal Hawkes model.

- 1: **Inputs:** point process: $\{(u_i, t_i, x_i, y_i)\}_{i=1}^N$; initial guesses for parameters: $\Theta^{(0)} = (\{K_{uv}^{(0)}\}_{u,v=1}^U, \{\beta_{uv}^{(0)}\}_{u,v=1}^U, \sigma^{(0)}, \omega^{(0)})$ and $\{p_{ij}^{(0)}\}_{i,j=1}^N, \{p_{ij}^{b,(0)}\}_{i,j=1}^N$; termination threshold: ϵ .
 - 2: **Outputs:** model parameters $\Theta = (\{K_{uv}\}_{u,v=1}^U, \{\beta_{uv}\}_{u,v=1}^U, \sigma, \omega)$.
 - 3: Initialize $\delta = 1$ and $k = 0$.
 - 4: **while** $\delta > \epsilon$ **do**
 - 5: Let $\eta^{2,(k)}$ and $\sigma^{2,(k)}$ be the value of η^2 and σ^2 at the k th iteration.
 - 6: **Expectation step:** for $i, j \in \{1, 2, \dots, N\}$,
 - 7: $p_{ij}^{(k)} = (K_{u_j u_i} g(t_i - t_j, x_i - x_j, y_i - y_j)) / \lambda(x_i, y_i, t_i)$.
 - 8: $p_{ij}^{b,(k)} = \beta_{u_j u_i}^{(k)} \exp(-\frac{(x_j - x_i)^2 + (y_j - y_i)^2}{2\eta^{2,(k)}}) / (2\pi\eta^{2,(k)} T \lambda(x_i, y_i, t_i))$.
 - 9: **Maximization step:** for $u, \hat{u} \in \{1, 2, \dots, U\}$,
 - 10: $\omega^{(k+1)} = \frac{\sum_{j < i} p_{ij}^{(k)}}{\sum_{j < i} p_{ij}^{(k)}(t_i - t_j) + \sum_{u=1}^U \sum_{i=1}^N K_{uu}(T - t_i) e^{-\omega(T-t_i)}}.$
 - 11: Let N_u denote the number of events in point process u ; and let i_l^u , with $l \in \{1, \dots, N_u\}$, index the events for process u . We update
 - 12: $K_{\hat{u}u}^{(k+1)} = \sum_{l=1}^{N_u} \sum_{t_{i_l^{\hat{u}}} < t_{i_l^u}} p_{i_l^{\hat{u}} i_l^u}^{(k)} / \sum_{l=1}^{N_{\hat{u}}} (1 - \exp[-w(T - t_{i_l^{\hat{u}}})])$,
 - 13: $\beta_{\hat{u}u}^{(k+1)} = \sum_{l=1}^{N_u} \sum_{i=1}^{N_{\hat{u}}} p_{i_l^u i_l^{\hat{u}}}^{b,(k)} / N_{\hat{u}}$.
 - 14: $\sigma^{2,(k+1)} = \sum_{i,j=1}^N (p_{ij}^{b,(k)} + p_{ij}^{(k)}) ((x_i - x_j)^2 + (y_i - y_j)^2) / \sum_{i,j=1}^N 2(p_{ij}^{b,(k)} + p_{ij}^{(k)})$.
 - 15: $\eta^{2,(k+1)} = \sigma^{2,(k+1)}$.
 - 16: $\delta = \|\Theta^{(k)} - \Theta^{(k+1)}\|$.
 - 17: $k = k + 1$.
 - 18: **end while**
-

3.3.1. Parametric model. The log-likelihood for the parametric model defined in (3.1) in a spatial region R and time window $[0, T]$ is

$$(3.4) \quad \log(L(\Theta; X)) = \sum_{k=1}^N \log(\lambda_{u_k}(t_k, x_k, y_k)) - \sum_{u=1}^U \iiint_R \int_0^T \lambda_u(t, x, y) dt dx dy.$$

Algorithm 3.2. EM-type algorithm for the nonparametric spatiotemporal Hawkes model.

- 1: **Inputs:** point process: $\{(u_i, t_i, x_i, y_i)\}_{i=1}^N$; initial guesses of parameters: $\{K_{uv}^{(0)}\}_{u,v=1}^U$ and $\{p_{ij}^{(0)}\}_{i,j=1}^N$; termination threshold: ϵ .
 - 2: **Outputs:** model parameters: $\{K_{uv}\}_{u,v=1}^U$; triggering probability between events: $\{p_{ij}\}_{i,j=1}^N$; temporal triggering kernel: g_1 ; spatial triggering kernel: g_2 .
 - 3: Initialize $\delta = 1$ and $\eta = 0$.
 - 4: **while** $\delta > \epsilon$ **do**
 - 5: Update $\tau^\eta(x, y)$ using (3.2).
 - 6: $\gamma_u^{(\eta)} = \sum_{u_i=u} p_{ii}^{(\eta)} / Z^{(\eta)}$, where $Z^{(\eta)}$ satisfies $\int_0^T \iint_R \tau^\eta(x, y) ds dt = Z^{(\eta)}$ on a bounded spatial domain R for $u \in \{1, \dots, U\}$.
 - 7: $K_{uv}^{(\eta)} = \sum_{u_i=u} \sum_{u_j=v} p_{ij}^{(\eta)} / N_u$ for $u, v \in \{1, \dots, U\}$.
 - 8: $g_1^{(\eta)}(t) = \sum_{i,j \in C_k} p_{ij}^{(\eta)} / [\Delta t_k \sum_{i < j} p_{ij}^{(\eta)}]$ for t in the k th temporal bin.
 - 9: $h^{(\eta)}(r) = \sum_{i,j \in D_k} p_{ij}^{(\eta)} / [\Delta r_k \sum_{i < j} p_{ij}^{(\eta)}]$ for r in the k th spatial bin. Set $g_2^{(\eta)}(r) = h^{(\eta)}(r) / (2\pi r)$.
 - 10: $p_{ij}^{(\eta+1)} = K_{u_i u_j}^{(\eta)} g_1^{(\eta)}(t_j - t_i) g_2^{(\eta)}(r_{ij})$ for $i < j$ and $p_{jj}^{(\eta+1)} = \mu_{u_j}^{(\eta)}(x_j, y_j)$.
 - 11: Normalize $p_{ij}^{(\eta+1)}$ so that $\sum_{i=1}^N p_{ij}^{(\eta+1)} = 1$ for any j .
 - 12: $\delta = \max_{i,j} \|p_{ij}^{(\eta+1)} - p_{ij}^{(\eta)}\|$ and $\eta = \eta + 1$.
 - 13: **end while**
-

We define random variables Y_{ij} and Y_{ij}^b using the approach from [45]. If event j triggers event i via the kernel g , then $Y_{ij} = 1$; otherwise, $Y_{ij} = 0$. The equality $Y_{ij}^b = 1$ indicates that event i is triggered by event j at a background rate of μ . We define two expectation matrices: $\mathbf{P}(i, j) = p_{ij} = \mathbb{E}[Y_{ij}]$ and $\mathbf{P}^b(i, j) = p_{ij}^b = \mathbb{E}[Y_{ij}^b]$. We convert the incomplete-data log-likelihood function in (3.4) into the following complete-data log-likelihood function:

$$\begin{aligned} \log(L(\Theta; X, Y)) &= \sum_{j < i} Y_{ij} \log(K_{u_j u_i} g(t_i - t_j, x_i - x_j, y_i - y_j)) - \sum_{u=1}^U \sum_{i=1}^N \beta_{u_i u} \\ &\quad - \sum_{u=1}^U \sum_{i=1}^N K_{u_i u} \left(1 - e^{-w(T-t_i)}\right) + \sum_{i=1}^N \sum_{j=1}^N Y_{ij}^b \log(\mu_{u_i}). \end{aligned}$$

We then calculate the expectation function using (3.3) to obtain

$$\begin{aligned} Q(\Theta) &= \sum_{i=1}^N \sum_{j=1}^N p_{ij}^b \log \left(\frac{\beta_{u_j u_i}}{2\pi\eta^2 T} \exp \left(-\frac{(x_i - x_j)^2 + (y_i - y_j)^2}{2\eta^2} \right) \right) - \sum_{u=1}^U \sum_{i=1}^N \beta_{u_i u} \\ &\quad + \sum_{j < i} p_{ij} \log \left(\omega K_{u_j u_i} e^{-\omega(t_i - t_j)} \frac{1}{2\pi\sigma^2} \exp \left(-\frac{(x_i - x_j)^2 + (y_i - y_j)^2}{2\sigma^2} \right) \right) \\ &\quad - \sum_{u=1}^U \sum_{i=1}^N K_{u_i u} \left(1 - e^{-w(T-t_i)}\right). \end{aligned}$$

We perform the maximization step of the EM-type algorithm (a projected gradient ascent) [39] directly by taking derivatives with respect to the parameters and setting them to 0. For the expectation step, we use the “optimal” parameter values from the prior maximization step to update the probabilities p_{ij} and p_{ij}^b . By (alternately) iterating the expectation and maximization steps, we obtain Algorithm 3.1 for the parametric model. For initialization, we sample Θ^0 , p_{ij} , and p_{ij}^b uniformly at random. Note, additionally, that $p_{ij} = 0$ for $i \leq j$.

Algorithm 3.3. Simulation of a multivariate spatiotemporal Hawkes process.

- 1: **Inputs:** time-window size: T ; spatial region: $R \subset \mathbb{R}^2$; background rate: $\{\gamma_u\}_{u=1}^U$; triggering matrix: $\{K_{uv}\}_{u,v=1}^U$; temporal and spatial triggering kernels: $g_1(t)$, $g_2(x, y)$.
 - 2: **Output:** point process: $\mathbf{C} = \{(u_i, t_i, x_i, y_i)\}_{i=1}^N$.
 - 3: Initialize an empty set $\mathbf{C}$ and an empty stack $\mathbf{Q}$.
 - 4: **Generate background events:**
 - 5: Draw N_u^b , the number of background events of type u , from a Poisson distribution with parameter $\lambda = \gamma_u T$ for each $u \leq U$.
 - 6: Add each background event $i \leq \sum_{u=1}^U N_u^b$, with its associated tuple (x_i, y_i, t_i, u_i) , to the set $\mathbf{C}$ and the stack $\mathbf{Q}$, where we draw (x_i, y_i, t_i) from the uniform spatiotemporal distribution over the time interval $[0, T]$ and a bounded spatial region R .
 - 7: **Generate triggered events:**
 - 8: **while** $\mathbf{Q}$ is not empty **do**
 - 9: Remove the most recently added element (x_i, y_i, t_i, u_i) from the stack $\mathbf{Q}$.
 - 10: Draw n_i , the number of events triggered by event i , from a Poisson distribution with parameter $\lambda_i = \sum_{u'=1}^U K_{u_i u'}$.
 - 11: Generate events (x_k, y_k, t_k, u_k) for each $k \leq n_i$ as follows:
 - 12: Sample t_k , (x_k, y_k) , and u_k according to $g_1(t - t_i)$, $g_2(x - x_i, y - y_i)$, and $P(u_k = \tilde{u}) = \frac{K_{u_i \tilde{u}}}{\sum_{v=1}^U K_{u_i v}}$, respectively.
 - 13: Add (x_k, y_k, t_k, u_k) to the set $\mathbf{C}$.
 - 14: **if** $t_k \leq T$ **then**
 - 15: Add the element (x_k, y_k, t_k, u_k) to the stack $\mathbf{Q}$.
 - 16: **end if**
 - 17: **end while**
-

3.3.2. Nonparametric model. The log-likelihood function of the nonparametric model is the same as that for the parametric model in (3.4). We use a similar approach as before to derive an EM-type algorithm for the nonparametric model. The main differences are that (1) only Y_{ij} are latent variables and $Y_{ii} = 1$ signifies that event i is a background event, whereas $Y_{ji} = 1$ signifies that event i is triggered by event j ; and (2) we assume that the triggering kernels $g_1(t)$ and $g_2(r)$ are piecewise-constant functions. Note that the elements of the expectation matrix are $\mathbb{E}[Y_{ij}] = \mathbf{P}(i, j) = p_{ij}$. In this case, note that $p_{ij} = 0$ for $i > j$. We discretize space and time into n_t^{bins} temporal bins and n_r^{bins} spatial bins, and the kernel takes a constant value in each spatiotemporal bin.

To formally present our EM-type algorithm (see Algorithm 3.2), we borrow notation from [22]. Let C_k denote the set of event pairs (i, j) for which $t_j - t_i$ belongs to the k th temporal

bin, D_k denote the set of event pairs (i, j) for which r_{ij} (the distance between nodes i and j) belongs to the k th spatial bin, N_u denote the number of events for node u , the parameter Δt_k denote the size of the k th temporal bin, and Δr_k denote the size of the k th spatial bin.

3.4. Simulations. To generate synthetic data for model comparisons, we need to simulate self-exciting point processes with the conditional intensity in (3.1) for each process u . We use the branching structures [79] of self-exciting point processes to develop Algorithm 3.3 for our simulations.

4. Numerical experiments and results. We apply our algorithms³ to both synthetic and real-world data sets to demonstrate the usefulness of (1) incorporating spatial information into Hawkes processes and (2) our nonparametric approach. We consider a synthetic data set in subsection 4.1, a social-media data set from Gowalla in subsection 4.2, a crime-topic network data set in subsection 4.3, and a violent gang-crime data set in subsection 4.4. Using the first three of these data sets, we compare our nonparametric spatiotemporal Hawkes model (“nonparametric Hawkes”) with the Bayesian Hawkes model⁴ from [40] (“Bayesian Hawkes”), the exclusively temporal Hawkes model with kernel $g(t) = \omega \exp(-\omega t)$ from [23] (“temporal Hawkes”), and the parametric spatiotemporal Hawkes model that we detailed in subsection 3.1 (“parametric Hawkes”). We make comparisons by examining how well the following properties are recovered when we infer a triggering matrix: (1) symmetry and reciprocity, (2) existence of edges, and (3) community structure. We also demonstrate the ability of both our parametric and nonparametric algorithms to infer the triggering kernel g . Using the fourth data set (see subsection 4.4), we study a network of crime events to illustrate our methodology with an additional application. We examine relations between crime events and repeated triggering patterns.

4.1. Synthetic data. We first generate synthetic triggering matrices $\mathbf{K}$ using a weighted stochastic block model (WSBM) [3, 57]. We assign a network’s nodes to four sets (called “communities”) and assign edges to adjacency-matrix blocks based on the set memberships of the nodes. Two of the communities consist of ten nodes each, and the other two communities consist of five nodes each. For each edge, we first draw a Bernoulli random variable to determine whether it exists, and we then draw an exponential random variable to determine the weight of the edge (if it exists). An edge between nodes from the same community exists with probability 0.68, and an edge between nodes from different communities exists with probability 0.2. The decay-rate parameter for the exponential random variable in these two situations is 0.1 and 0.01, respectively. By construction, our triggering matrices are symmetric.

The triggering matrices that we generate in this way are not guaranteed to satisfy the stability condition for Hawkes processes; this condition is that the largest-magnitude eigenvalue of $\mathbf{K}$ is smaller than 1 [18]. When this condition is satisfied, it is almost surely true that each event has finitely many subsequent events as “offspring”. In our work, we discard

³Our code is available at <https://github.com/ybcmath/MultiSTHP>.

⁴Linderman and Adams [40] used a sparse log-Gaussian Cox process to model the background rate, a logistic-normal density for the temporal kernel, an Aldous–Hoover graph prior for the existence of entries in $\mathbf{K}$, and gamma prior for the weights of those entries. In all of our experiments with their approach, we use the default hyperparameters and priors that come with the published code at <https://github.com/slinderman/pyhawkes>.

any simulated adjacency matrices that do not satisfy the stability condition, and we generate a new one to replace it. (With our choices of the parameters, we discard about 65% of the generated adjacency matrices.)

With each triggering matrix $\mathbf{K}$, we use Algorithm 3.3 to simulate a multivariate spatiotemporal Hawkes process with our parametric model in subsection 3.1 with $\omega = 0.6$, $\sigma^2 = 0.3$, $T = 250$, $R = [0, 1] \times [0, 1]$, and a homogeneous background rate $\mu_u = 0.2$ for all nodes u . We then reconstruct the underlying networks and the triggering kernels from the simulated data. As a result, our parametric spatiotemporal Hawkes model serves as a “ground-truth model,” and we expect it to have the best performance for synthetic data, given that we use the same model to produce the data.

4.1.1. Symmetry and reciprocity. As we noted in subsection 4.1, our simulated triggering matrices are symmetric, but our reconstructed adjacency matrices generally are not symmetric. Without prior information about symmetry, measuring deviation from symmetry gives one way to evaluate the performance of our inference methods. We use various reciprocity measures to quantify such deviation.

We conduct two sets of experiments. In the first set, we fix a single synthetic triggering matrix and simulate ten multivariate spatiotemporal Hawkes point processes. We then estimate the triggering matrix $\mathbf{K}$ from each point process using various methods, which we thereby compare with each other. In a second set of experiments, instead of fixing a single triggering matrix, we generate ten different triggering matrices using the same WSBM and the same parameters; and we simulate one point process for each triggering matrix.

There is no standard way of measuring reciprocity in a weighted network. In our calculations, we use diagnostics that were proposed in [68] and [4]. First, as in [68], we compute the reciprocated edge weight $K_{uv}^{\leftrightarrow} = \min\{K_{uv}, K_{vu}\}$, and we then calculate a network-level reciprocity score R_1 as the ratio between the total reciprocated weight $W^{\leftrightarrow} = \sum_{u \neq v} K_{uv}^{\leftrightarrow}$ and the total weight $W = \sum_{u \neq v} K_{uv}$. That is, the “reciprocity” is $R_1 := W^{\leftrightarrow}/W$. Second, Akoglu, de Melo, and Faloutsos [4] proposed three edge-level measures of reciprocity: (1) the “ratio” $R_{\text{ratio}} := \min\{K_{uv}, K_{vu}\}/\max\{K_{uv}, K_{vu}\}$; (2) “coherence” $R_{\text{coher}} = 2\sqrt{K_{uv}K_{vu}}/(K_{uv} + K_{vu})$; and (3) “entropy” $R_{\text{entropy}} := -r_{uv} \log_2(r_{uv}) - r_{vu} \log_2(r_{vu})$, where $r_{uv} = K_{uv}/(K_{uv} + K_{vu})$. These last three measures of reciprocity are measured at an edge level (as they are defined for a pair of nodes), whereas R_1 is a network-level measure. For each edge-level measure, we obtain a network-level measure by calculating the score for each pair of nodes in a network and then taking a mean over all pairs of nodes. Each of the above quantities gives a score between 0 and 1, where a larger value indicates a stronger tendency for the nodes in a network to reciprocate. In a perfectly symmetric and reciprocal network, each of the four methods gives a value of 1. As another measure of reciprocity, we also calculate the Pearson correlation between the triggering matrix and its transpose.

In Table 1, we report the mean reciprocity and the standard deviation over ten simulations with the same triggering matrix. In Table 2, we report the mean results from ten different triggering matrices. Both spatiotemporal models give higher scores than the exclusively temporal models, which is what we expected, as the temporal models discard spatial information. According to these measures of success, the nonparametric model has the best performance, even over the ground-truth (parametric) model that generated the data in the first place.

Table 1

Reciprocity of the triggering matrices that we infer using different methods: a nonparametric spatiotemporal Hawkes model, a temporal Hawkes model, a parametric spatiotemporal Hawkes model, and a fully Bayesian Hawkes model. We report the mean and standard deviation (in parentheses) over ten simulations that use the same (ground-truth) triggering matrix.

	Nonparametric	Temporal	Parametric	Bayesian
R_1	0.59 (0.05)	0.29 (0.06)	0.54 (0.03)	0.36 (0.03)
Correlation	0.84 (0.05)	0.36 (0.16)	0.79 (0.05)	0.30 (0.14)
Ratio	0.55 (0.02)	0.37 (0.11)	0.58 (0.02)	0.32 (0.02)
Coherence	0.75 (0.01)	0.63 (0.03)	0.71 (0.02)	0.68 (0.02)
Entropy	0.71 (0.01)	0.59 (0.03)	0.68 (0.02)	0.60 (0.02)

Table 2

Reciprocity of the triggering matrices that we infer using different methods: a nonparametric spatiotemporal Hawkes model, a temporal Hawkes model, a parametric spatiotemporal Hawkes model, and a fully Bayesian Hawkes model. We report the mean and standard deviation (in parentheses) over ten simulations, each with a different (ground-truth) triggering matrix.

	Nonparametric	Temporal	Parametric	Bayesian
R_1	0.61 (0.12)	0.36 (0.12)	0.55 (0.10)	0.40 (0.05)
Correlation	0.81 (0.16)	0.48 (0.27)	0.76 (0.15)	0.23 (0.14)
Ratio	0.63 (0.04)	0.43 (0.06)	0.62 (0.03)	0.33 (0.03)
Coherence	0.78 (0.04)	0.62 (0.03)	0.72 (0.03)	0.70 (0.03)
Entropy	0.75 (0.05)	0.58 (0.03)	0.69 (0.03)	0.62 (0.04)

4.1.2. Edge reconstruction. We also evaluate the reconstruction methods based on their ability to recover the existence of edges. This is particularly relevant if we want to know whether there is a connection between two entities. We will discuss this application in detail using the Gowalla data set (see subsection 4.2).

We consider an edge to exist if the corresponding weighted entry in an inferred triggering matrix exceeds a certain threshold. For different threshold levels, we compute the numbers of true positives (TP), false positives (FP), true negatives (TN), and false negatives (FN) for a given ground-truth triggering matrix. We summarize our results in a receiver operating characteristic (ROC) plot (see Figure 1), in which we plot the true-positive rate (TPR) (where $\text{TPR} = \text{TP}/(\text{TP} + \text{FN})$) versus the false-positive rate (FPR) (where $\text{FPR} = \text{FP}/(\text{FP} + \text{TN})$). A better inference of a triggering matrix gives a larger value of TPR for a fixed FPR.

Based on the ROC plot in Figure 1, we conclude that the spatiotemporal models—both the parametric and nonparametric spatiotemporal Hawkes models that we proposed in section 3—outperform the exclusively temporal ones. Therefore, incorporating spatial information improves the quality of our reconstructed unweighted (i.e., binary) networks, at least according to this measure of success. Unsurprisingly, the best results are from our parametric (ground-truth) model. The performance of our nonparametric model is very close to that of the parametric model, confirming its effectiveness at inferring the existence of edges.

4.1.3. Inferred kernels. We report inferred kernels for our synthetic networks of the nonparametric spatiotemporal Hawkes model, parametric spatiotemporal Hawkes model, and temporal Hawkes model in Figure 2. Recall that the ground-truth kernels that we use to sim-




Figure 1. Model comparison using synthetic networks. We show the mean ROC curves with error bars (averaged over ten simulations, each with a different triggering matrix) on edge reconstruction. The ROC curve of a better reconstruction should be closer to 1 for a larger range of horizontal-axis values, such that it has a larger area under the curve (AUC), which is equal to the probability that a uniformly-randomly chosen existing edge in a ground-truth network has a larger weight than a uniformly-randomly chosen missing edge in the inferred network.

Table 3

The L_1 errors of the inferred spatial and temporal kernels for a nonparametric spatiotemporal Hawkes model, a temporal Hawkes model, and a parametric spatiotemporal Hawkes model. We simulate ten point processes with the same triggering matrix and triggering kernel. We report the mean and standard deviation (in parentheses) of the L_1 errors averaged over ten simulations that use the same triggering kernel and matrix. Note that the exclusively temporal model does not estimate a spatial kernel.

	Nonparametric	Temporal	Parametric
Temporal kernel	0.07 (0.02)	0.20 (0.06)	0.02 (0.02)
Spatial kernel	0.06 (0.02)	-	0.12 (0.02)

ulate point processes are $g_1(t) = \omega \exp(-\omega t)$ and $h(r) = 2\pi r g_2(r) = \frac{r}{\sigma^2} \exp(-\frac{r^2}{2\sigma^2})$, where $r^2 = x^2 + y^2$, $\omega = 0.6$, and $\sigma^2 = 0.3$. Let $\hat{g}_1$ and $\hat{h}$ denote the inferred temporal and spatial kernels, respectively.

We calculate the L_1 errors $\int |g_1(t) - \hat{g}_1(t)| dt$ and $\int |h(r) - \hat{h}(r)| dr$. We report these errors in Table 3 and present visualizations of the inferred kernels in Figure 2. As expected, the two spatiotemporal Hawkes models give more accurate kernel inference than the exclusively temporal model. The nonparametric Hawkes model does not use any information about the ground-truth kernels. Surprisingly, it is more accurate (in terms of the L_1 error) at inferring the spatial trigger kernel than the parametric model, whose kernel has the same parametric form as the ground-truth kernel.

4.1.4. Community-structure recovery. We also evaluate the quality of the inferred networks based on their community structure, in which dense sets of nodes in a network are connected sparsely to other dense sets of nodes [21, 61]. Recall that we planted a four-community structure in the synthetic triggering matrices (see subsection 4.1). We apply the community-detection methods from [3] (an inference method for a WSBM), [35] (symmetric




Figure 2. Model comparison using synthetic networks: Inferred (left) temporal and (right) spatial kernels using three different methods: temporal Hawkes, parametric spatiotemporal Hawkes, and nonparametric spatiotemporal Hawkes. The dashed curves are (ground-truth) kernels that we used to generate the synthetic data.

nonnegative matrix factorization; NMF), and [29, 48, 50, 52] (modularity maximization).⁵ The WSBM that we infer for community detection is the same one that we used to construct the synthetic adjacency matrices. (See our earlier discussion in [subsection 4.1](#).) To evaluate our inferred community structure, we use the square-root variant of *normalized mutual information* (NMI) [70] between the inferred community assignment and “ground-truth” community labels. Specifically, let S_1 and S_2 be community assignments of the U nodes to C_1 and C_2 communities, respectively; and let $S_{\ell k}$, with $\ell \in \{1, 2\}$ and $k \in \{1, 2, \dots, C_\ell\}$, denote the set of nodes in the k th community in assignment S_ℓ . The NMI between S_1 and S_2 is

$$\text{NMI}(S_1, S_2) = \frac{I(S_1, S_2)}{\sqrt{H(S_1)H(S_2)}} \in [0, 1],$$

where $I(S_1, S_2) = \sum_{i=1}^{C_1} \sum_{j=1}^{C_2} \frac{|S_{1i} \cap S_{2j}|}{U} \log \left(\frac{|S_{1i} \cap S_{2j}|/U}{|S_{1i}| |S_{2j}|/U^2} \right)$ (where $|J|$ denotes the cardinality of the set J) and the entropy is $H(S_\ell) = - \sum_{i=1}^{C_\ell} \frac{|S_{\ell i}|}{N} \log \left(\frac{|S_{\ell i}|}{N} \right)$ (with $\ell \in \{1, 2\}$). Intuitively, NMI measures the amount of information that is shared by two community assignments. If they are the same after permuting community labels, the NMI is equal to 1. A larger NMI score implies that the inferred community assignment shares more information with the ground-truth labels. See [73] for a discussion of other approaches for comparing different community assignments in networks.

There are numerous approaches for detecting communities in networks [21, 57, 61], and we use methods with readily available code. As we show in [Table 4](#), all of these community-detection methods perform better when we infer triggering matrices using both spatial and temporal information than when we use exclusively temporal information. One can, of course, repeat our experiments using other methods.

⁵For modularity maximization, we use the implementation of a (locally greedy) Louvain-like [9] method (called GENLOUVAIN) from [29] with the default resolution-parameter value of 1 and the Newman–Girvan null model.

Table 4

Normalized mutual information between the outputs of different community-detection methods applied to the inferred networks (from four different types of Hawkes models) and the ground-truth community structure (averaged over ten simulations, each with a different triggering matrix).

	Nonparametric	Temporal	Parametric	Bayesian
Weighted SBM	0.80	0.38	0.83	0.36
Symmetric NMF	0.62	0.31	0.66	0.19
Modularity maximization	0.64	0.47	0.71	0.28

4.2. Gowalla friendship network. Gowalla is a location-based social-media platform in which users share their locations by checking in. We use a Gowalla data set—collected in [15] using Gowalla’s public application program interface (API)—of a “friendship” network with 196,591 users, 950,327 edges, and a total of 6,442,890 check-ins of these users between February 2009 and October 2010. The data set also includes the latitude and longitude coordinates and the time (with a precision of one second) of each check-in. Similar to a Facebook “friendship” network, the Gowalla friendship network is undirected. The mean number of friends per user is 9.7, the median is 3, and the maximum is 14,730. We study several subnetworks in the Gowalla data set; see Appendix A for details. We view the spatiotemporal check-ins of Gowalla users within each subnetwork as events in a multivariate point process and infer relationships between these users.

We compare our nonparametric Hawkes method with the Bayesian Hawkes method and the exclusively temporal Hawkes method with respect to how well our inferred edges match the Gowalla friendships. Because a Gowalla “friendship” relationship is undirected in nature, we first symmetrize the inferred triggering matrix (by calculating $\tilde{\mathbf{K}} = (\mathbf{K} + \mathbf{K}^T)/2$) to obtain an undirected network. We then calculate FPRs and TPRs in the same fashion as in subsection 4.1.2 using the matrices $\tilde{\mathbf{K}}$ from associated “ground-truth” friendship networks, and we generate the corresponding ROC curves. As we can see in the ROC curves of three different cities in Figure 3, we obtain the best results when using our nonparametric Hawkes model that incorporates spatial information. In the examined subnetworks, the mean AUCs are 0.4277 (with a standard deviation of 0.1042) for the temporal Hawkes method; 0.5301 (with a standard deviation of 0.0585) for the Bayesian Hawkes method; 0.5816 (with a standard deviation of 0.0525) for our parametric spatiotemporal Hawkes method; and 0.6692 (with a standard deviation of 0.0421) for our nonparametric spatiotemporal Hawkes method.

4.3. Crime-topic network. In a recent paper on crime classification, Kuang, Brantingham, and Bertozzi [33] performed topic modeling (see [34] for a review) on short narrative (i.e., text) descriptions of all crimes, with spatial coordinates and time stamps (with a precision of one minute), that were reported to and officially recorded by the Los Angeles Police Department (LAPD) between 1 January 2009 and 19 July 2014. The premise in their work was that crime topics, sets of words that co-occur frequently in the same crime narrative, better reflect the ecological circumstances of crime than official crime classifications that are based on legal codes. Targeting discovery of up to twenty topics, they found six topics related to violent crime, eight topics related to property crime, and six topics that seem to be related to deception-based crime. This classifies the twenty crime topics into three types.




Figure 3. ROC curves of four different Hawkes models for reconstructing three Gowalla friendship networks. (See [Appendix A](#) for details about the networks.) We show results for nonparametric spatiotemporal Hawkes (dashed blue curves), temporal Hawkes (dash-dotted yellow curves), Bayesian Hawkes (solid red curves), and parametric spatiotemporal Hawkes (dotted purple curves).

In the present case study, we extend this work by modeling the above data set as a crime-topic network. We associate each crime topic with a node, and we infer edges based on whether crime events of one topic trigger events of other topics. That is, we discover latent relationships between different crime topics based on associated crime events. Inspired by previous research on point-process models of crime events [46], we model crime events of different topics via a multivariate point process and infer connections between the crime topics using our nonparametric spatiotemporal Hawkes method. To evaluate our approach, we compare the communities that we detect in the reconstructed network with the three crime classes in [33].

4.3.1. Community detection. We infer crime-topic networks directly from crime events in individual Los Angeles neighborhoods⁶ using our nonparametric spatiotemporal Hawkes method, our parametric spatiotemporal Hawkes method, the Bayesian Hawkes method, and the exclusively temporal Hawkes method. We investigate the 100 neighborhoods with the most reported crime events among all 296 neighborhoods of Los Angeles (LA). On average, there are 4,140 crime events in the top-100 neighborhoods and 8,750 such events in the top-10 neighborhoods. We then apply the community-detection methods that we mentioned in [subsection 4.1.4](#) to the reconstructed networks; this assigns crime topics to communities. We quantify the difference between these community assignments and the crime-topic classifications from [33] by calculating NMI. We also visualize the crime-topic networks of the Westwood and Wingfoot neighborhoods in [Figure 4](#); they are located in West LA and South LA, respectively. From [Table 5](#), we see that using spatial information combined with a nonparametric

⁶We use the Zillow neighborhood boundaries from <https://www.zillow.com>.




Figure 4. Crime-topic networks generated by the nonparametric spatiotemporal Hawkes and temporal Hawkes methods colored by community assignments from modularity maximization: (a) nonparametric spatiotemporal Hawkes in Westwood, (b) temporal Hawkes in Westwood, (c) nonparametric spatiotemporal Hawkes in Wingfoot, and (d) temporal Hawkes in Wingfoot.

Table 5

Mean NMI (with one standard deviation reported in parentheses) between community assignments from several community-detection methods and the classifications from [33] in the 100 neighborhoods in Los Angeles with the most recorded crime events between 1 January 2009 and 19 July 2014.

	Nonparametric	Temporal	Parametric	Bayesian
Symmetric NMF	0.25 (0.11)	0.12 (0.084)	0.084 (0.12)	0.12 (0.080)
Weighted SBM	0.24 (0.12)	0.085 (0.086)	0.078 (0.079)	0.076 (0.075)

kernel leads to the best mean NMI score among the methods that we examine.

4.4. Network of crime events. In the previous examples, we studied relationships between entities based on spatiotemporal events that are associated with them. To examine connections between events, we now define an event network, which is both weighted and directed, in which each event is a node and $\mathbf{A}$ denotes the adjacency matrix of this network. That is, $\mathbf{A}(i, j)$ is the probability that event j is triggered by event i , and $\mathbf{A}(i, i)$ is the probability that event i is a background event. From this definition, we see that $\mathbf{A}$ is equal to the expectation matrix $\mathbf{P}$ from subsection 3.3.2. The weight of an edge reflects a triggering effect between two events, and the direction points from the earlier event to the later one. (Removing weights yields an unweighted directed acyclic graph.) For example, we can build a crime-event network in which each node is a crime incident (i.e., an event), and we infer edges between events using our nonparametric spatiotemporal Hawkes model.

4.4.1. Stochastic declustering. With an event network, a natural question is whether one can differentiate between “true” background events and triggered events. Such differentiation using the probability p_{ii} is called *stochastic declustering* [78]. To determine whether event i is a background event, we compare p_{ii} with a uniformly random sample from the interval $(0, 1)$. If p_{ii} is larger than the random number, we consider this event to be from the background; otherwise, we consider it to be triggered by other events.

We perform declustering experiments on synthetic data; we simulate ten synthetic point processes using a fixed triggering matrix that we generate from a WSBM. (See subsection 4.1 for details of the WSBM.) Recall from Algorithm 3.3 that we retain causality information (i.e., which events cause which others and which events are from the background) in our simulations, giving a notion of “ground truth” about the ancestors of each event. One way

Table 6

Comparison of our stochastic declustering results for the nonparametric spatiotemporal Hawkes, parametric spatiotemporal Hawkes, and temporal Hawkes methods using synthetic point-process data with networks from a WSBM (see subsection 4.1) and background labels from simulations using Algorithm 3.3. (We do not include results for Bayesian Hawkes, because it does not provide $\mathbf{P}$ directly.) We report the mean and the standard deviation (in parentheses) of the branching-ratio error, precision, and recall over ten simulations (which we do for ten point processes with the same triggering kernels and triggering matrix $\mathbf{K}$). For each simulation, each calculation is the mean over 20 runs of stochastic declustering.

	Nonparametric	Parametric	Temporal
Branching-ratio error	0.039 (0.0050)	0.01 (0.011)	0.022 (0.019)
Recall	0.75 (0.0098)	0.65 (0.027)	0.60 (0.035)
Precision	0.70 (0.0082)	0.64 (0.0093)	0.59 (0.0086)

to measure the quality of declustering is by comparing the inferred branching ratio [67] with the one from the ground-truth data. The branching ratio is defined as $1 - N_b/N$, where N_b is the number of background events and N is the total number of events. However, the difference in branching ratios itself typically does not completely reflect reconstruction errors. For example, in an extreme case, stochastic declustering can erroneously misclassify some number of background events as triggered ones and the same number of triggered events erroneously as background ones, although the branching ratio is the same as the true branching ratio in this scenario. To resolve this problem, we view declustering as a binary classification problem that assigns events to be either background or triggered events. We use measurements such as recall and precision to evaluate our declustering results. Recall that “recall” is the ratio between the number of background events that are correctly recovered by the declustering methods (i.e., the true positives) to the total number of background events; and “precision” is the ratio between the number of true positives to the number of events that are labeled as background events by stochastic declustering. From the results in Table 6, we see that the temporal Hawkes method performs worse than the nonparametric spatiotemporal Hawkes and parametric spatiotemporal Hawkes methods. Our nonparametric spatiotemporal Hawkes method has the best recall and precision (with the smallest variations as well), and our parametric spatiotemporal Hawkes method has the smallest branching-ratio error.

4.4.2. Motif analysis. Declustering methods can help differentiate between background and triggered events in an event network. To further examine spatiotemporal dynamics, we consider causality information between events. Similar to a relational-event model [13], one can obtain causality information from the matrix $\mathbf{P}$, because p_{ij} is the probability that event j is triggered by event i . We focus on repeated patterns to obtain information about local causal structure. Specifically, we examine network motifs [44], which are recurrent (and often statistically significant) patterns in a network.

We find that motif analysis is insightful for studying gang-crime event networks. Gang crimes are often characterized by retaliations (triggered crime events) between rival gangs; this can lead to a series of tit-for-tat reciprocal crimes. To find significant gang-retaliation patterns, we use a gang-crime data set (provided by the LAPD) from 2014–2015 with 4,158 events in Los Angeles. Using these data, we generate an event network with our nonparametric spatiotemporal Hawkes method. We then threshold the network, by keeping only edges whose




Figure 5. All possible three-node motifs for a network of events in the form of a directed acyclic graph (DAG). The DAG structure arises from the temporal information in the events. We highlight the nodes in the feedforward-loop motif (D) in red.

weight is at least 0.1 and then binarizing them, to yield an unweighted network. We use the motif-detection method and code [1] from [44], including their null model.⁷

We find, for thresholds ranging from 0.5 to 0.001, that a three-node feedforward-loop motif [41] occurs more significantly than by chance (with z-scores that are larger than 2) in both the city-wide data set and in the South LA⁸ subset (which consists of 1,912 events) of the data set. Davies and Marchione [19] found that the same three-node motif is significant in networks that they constructed (using different methods from ours both for network construction and for motif detection) using data sets from maritime piracy and residential burglaries.

We focus on the South LA area because it is the center of a gang intervention program [10]. Establishing which causal structures are statistically significant has important implications for countering gang violence, and a fast response to a gang crime may reduce the potential that it triggers a future retaliation. Knowing that feedforward-loop network motifs occur at rates that are larger than chance suggests that disrupting retaliation may require assessment of trade-offs in how to allocate intervention resources. For example, in a simple triggering chain (see Figure 5C), one can expect that intervention following an initial triggering event will have a direct effect on the second event and an indirect effect on the third event, although the effect on the third event may be attenuated by the intervening event. By contrast, we expect that intervention following the first event in the feedforward-loop motif (see Figure 5D) will have a direct effect on the second event and both a direct and indirect effect on the third event. Given the feedforward structure, it is possible that the third event is more likely to be disrupted from intervention following the first event than would be the case with only direct intervention following the second event.

5. Conclusions and discussion. In this paper, we studied the role of spatial information and nonparametric techniques in network reconstruction. We used point-process models to

⁷For each of our networks, we produce 100 “randomized” networks. To produce one such network, we use the default edge-swapping approach from [44]. This entails making a number of random swaps equal to about 100–200 times the number of edges. For each node in a network, we require that the randomized network preserves its numbers of in-edges, out-edges, and bidirectional edges.

⁸We use the term “South LA” to designate a specific area of Los Angeles that is defined in a recent Gang Reduction and Youth Development (GRYD) report [10].

infer latent networks from synthetic and real-world spatiotemporal data sets. We then applied tools from network analysis to examine the inferred networks.

As we have illustrated, it is very important to incorporate spatial information in network reconstruction. However, using such information effectively requires making a good choice of spatiotemporal triggering kernels. We achieved this using a nonparametric approach. Through experiments on synthetic data sets, we showed that a nonparametric spatiotemporal Hawkes method is capable of doing a good job of successfully recovering spatial and temporal triggering kernels. Moreover, our approach is able to infer a network structure that better recovers—compared to other network-reconstruction methods that we studied—symmetry and reciprocity, edge reconstruction, and community structure. Through experiments on real-world data sets, we illustrated that our approach yields meaningful inferred networks, in the sense that they have large positive correlations with some metadata.


Our approach is not limited to the Euclidean distance (for the spatial variables) that has been used commonly in seismology [56] and crime applications [46]. In other words, although the spatial triggering kernel $g_2(r)$ in this paper is a function of Euclidean distance, one can potentially use any notion of “distance” between two entities. For example, in a network, one can measure a distance between two entities based on the length of shortest paths between them. In a recent paper, Green, Horel, and Papachristos [25] proposed a social-contagion model in which they assumed, using a parametric form, that the strength of triggering in a Hawkes-process model depends on the shortest-path distance. With our approach, we can nonparametrically estimate such dependence. To give another example, consider a point process in which each event is associated with textual information. For instance, in a Twitter data set, one can consider each tweet (a time-stamped body of text) as an event in a point process. One can measure a distance between two tweets based on their text. It will be interesting to apply our nonparametric spatiotemporal point-process approach to these and other applications.

Naturally, our network-reconstruction method is not without limitations. It uses $O(U^2)$ parameters for U nodes. To avoid underfitting, it requires a large number of observed events. The computational complexity and memory requirement scale at least quadratically with the number of events, so the current EM-type algorithm is not ideal for analyzing large data sets. Therefore, it will be important to improve our inference method for network reconstruction, especially for large data sets.

Appendix A. Preprocessing of the Gowalla data. In this section, we detail how we preprocess the Gowalla data that were collected and studied in [15]. We examine data from three cities: New York City, Los Angeles, and San Francisco. In Figure 6, we visualize the networks that we use in this paper.

A.1. New York City (NYC). We study check-ins in New York City (NYC) during the period April–October 2010. We use a bounding box (with a north latitude of 40.92, a south latitude of 40.48, an east longitude of -73.70 , and a west longitude of -74.26)⁹ to locate check-ins in NYC. We consider “active” users, who have at least 100 check-ins during the

⁹We obtain latitude and longitude coordinates from the Google Maps geocoding service through the website <http://www.mapdevelopers.com>.



(a) A subnetwork of the largest connected component of the Gowalla network in NYC. (b) A subnetwork of the largest connected component of the Gowalla network in LA. (c) The 1-ego network of a user of Gowalla in SF.

Figure 6. Three different friendship networks (from New York City, Los Angeles, and San Francisco) in the Gowalla data set. We compare different network-reconstruction methods for these networks.

period. To alleviate our computational burden, we also only consider users who have at most 500 check-ins during the period to reduce the number of users and the total number of check-ins. Our inference process requires computing a triggering probability for each pair of events (i.e., check-ins); this results in a full upper-triangular matrix. The number of nonzero entries in this matrix scales with the square of the total number of events, so the memory requirement also scales quadratically with the number of events. We perform experiments only for cases in which the total number of events is at most 10,000 to be able to store triggering probabilities for all pairs of events in 4 GB of memory. There are 5,801 unique users with at least one check-in in NYC during the period, and there are 101,329 check-ins in total. After removing “inactive” users (i.e., those with strictly fewer than 100 check-ins) and overly active users (i.e., those with strictly more than 500 check-ins), we are left with 160 users and a total of 29,118 check-ins. We also restrict ourselves to users in the largest connected component (LCC) of the network. This yields 46 users and 8,495 check-ins, on which we apply our inference methodology.

A.2. Los Angeles (LA). We apply the same procedure as in [Appendix A.1](#) on the check-in data for Los Angeles (LA). The bounding box that we use for LA has a north latitude of 34.34, a south latitude of 33.70, an east longitude of -118.16 , and a west longitude of -118.67 . We restrict the area of LA to be the same as that of NYC, although LA’s geographic area is much larger than that of NYC. After selecting only users in the LCC of the Gowalla network among users who are active (with at least 150 check-ins) but not overly active (with at most 1,000 check-ins) users, we are left with 23 users and 6,203 check-ins.

A.3. San Francisco (SF). To look at a different type of example, we also study the 1-ego network of the most popular user (who has 14 friends) in San Francisco (SF). (A 1-ego network [74] of a node is an induced subgraph that includes a focal node—the ego—and its neighbors, which are its adjacent nodes.) The bounding box that we use for SF has a north latitude of 37.93, a south latitude of 37.64, an east longitude of -122.28 , and a west longitude of -123.17 . In this 1-ego network, there are 9,887 check-ins.

Acknowledgments. We thank Chandan Dhal and Jialin Liu for helpful discussions and preliminary work on community detection. We also thank Jialin Liu for making the plots in Figure 4.

REFERENCES

- [1] MFINDER 1.21, <http://www.weizmann.ac.il/mcb/UriAlon/download/network-motif-software>.
- [2] M. ACHAB, E. BACRY, S. GAÏFFAS, I. MASTROMATTEO, AND J.-F. MUZY, *Uncovering causality from multivariate Hawkes integrated cumulants*, J. Mach. Learn. Res., 18 (2017), pp. 6998–7025.
- [3] C. AICHER, A. Z. JACOBS, AND A. CLAUSET, *Learning latent block structure in weighted networks*, J. Complex Netw., 3 (2014), pp. 221–248.
- [4] L. AKOGLU, P. O. V. DE MELO, AND C. FALOUTSOS, *Quantifying reciprocity in large weighted communication networks*, in Pacific-Asia Conference on Knowledge Discovery and Data Mining, Springer-Verlag, 2012, pp. 85–96.
- [5] E. BACRY, I. MASTROMATTEO, AND J.-F. MUZY, *Hawkes processes in finance*, Mark. Microstructure Liq., 1 (2015), 1550005.
- [6] D. BALCAN, V. COLIZZA, B. GONÇALVES, H. HU, J. J. RAMASCO, AND A. VESPIGNANI, *Multiscale mobility networks and the spatial spreading of infectious diseases*, Proc. Natl. Acad. Sci. USA, 106 (2009), pp. 21484–21489.
- [7] H. BARBOSA, M. BARTHELEMY, G. GHOSHAL, C. R. JAMES, M. LENORMAND, T. LOUAIL, R. MENEZES, J. J. RAMASCO, F. SIMINI, AND M. TOMASINI, *Human mobility: Models and applications*, Phys. Rep., 734 (2018), pp. 1–74.
- [8] M. BARTHELEMY, *Morphogenesis of Spatial Networks*, Springer International Publishing, Cham, 2018.
- [9] V. D. BLONDEL, J.-L. GUILLAUME, R. LAMBIOTTE, AND E. LEFEBVRE, *Fast unfolding of communities in large networks*, J. Stat. Mech. Theory Exp., 2008 (2008), P10008.
- [10] P. J. BRANTINGHAM, N. SUNDBACK, B. YUAN, AND K. CHAN, *GRYD Intervention Incident Response & Gang Crime 2017 Evaluation Report*, 2017; available online at paleo.sscnet.ucla.edu/GRYD_IR_and_Gang_Crime_Report_FINALv2.pdf.
- [11] D. BROCKMANN, L. HUFNAGEL, AND T. GEISEL, *The scaling laws of human travel*, Nature, 439 (2006), pp. 462–465.
- [12] E. N. BROWN, R. E. KASS, AND P. P. MITRA, *Multiple neural spike train data analysis: State-of-the-art and future challenges*, Nature Neuroscience, 7 (2004), pp. 456–461.
- [13] C. T. BUTTS, *A relational event framework for social action*, Sociol. Methodol., 38 (2008), pp. 155–200.
- [14] S. CHEN, A. SHOJAIE, E. SHEA-BROWN, AND D. WITTEN, *The Multivariate Hawkes Process in High Dimensions: Beyond Mutual Excitation*, preprint, <https://arxiv.org/abs/1707.04928>, 2017.
- [15] E. CHO, S. A. MYERS, AND J. LESKOVEC, *Friendship and mobility: User movement in location-based social networks*, in Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, New York, 2011, pp. 1082–1090.
- [16] M. D. CONOVER, C. DAVIS, E. FERRARA, K. MCKELVEY, F. MENCZER, AND A. FLAMMINI, *The geospatial characteristics of a social movement communication network*, PloS ONE, 8 (2013), e55957.
- [17] N. CRESSIE, *Statistics for Spatial Data*, John Wiley & Sons, Hoboken, NJ, 2015.
- [18] D. J. DALEY AND D. VERE-JONES, *An Introduction to the Theory of Point Processes: Volume II: General Theory and Structure*, Springer-Verlag, Berlin, 2007.
- [19] T. DAVIES AND E. MARCHIONE, *Event networks and the identification of crime pattern motifs*, PloS ONE, 10 (2015), e0143638.
- [20] M. EICHLER, R. DAHLHAUS, AND J. DUECK, *Graphical modeling for multivariate Hawkes processes with nonparametric link functions*, J. Time Series Anal., 38 (2017), pp. 225–242.
- [21] S. FORTUNATO AND D. HRIC, *Community detection in networks: A user guide*, Phys. Rep., 659 (2016), pp. 1–44.
- [22] E. W. FOX, F. P. SCHOENBERG, AND J. S. GORDON, *Spatially inhomogeneous background rate estimators and uncertainty quantification for nonparametric Hawkes point process models of earthquake occurrences*, Ann. Appl. Stat., 10 (2016), pp. 1725–1756.

- [23] E. W. FOX, M. B. SHORT, F. P. SCHOENBERG, K. D. CORONGES, AND A. L. BERTOZZI, *Modeling e-mail networks and inferring leadership using self-exciting point processes*, J. Amer. Statist. Assoc., 111 (2016), pp. 564–584.
- [24] C. W. J. GRANGER, *Investigating causal relations by econometric models and cross-spectral methods*, Econometrica, 37 (1969), pp. 424–438.
- [25] B. GREEN, T. HOREL, AND A. V. PAPACHRISTOS, *Modeling contagion through social networks to explain and predict gunshot violence in Chicago, 2006 to 2014*, JAMA Intern. Med., 177 (2017), pp. 326–333.
- [26] E. C. HALL AND R. M. WILLETT, *Tracking dynamic point processes on networks*, IEEE Trans. Inform. Theory, 62 (2016), pp. 4327–4346.
- [27] A. G. HAWKES AND D. OAKES, *A cluster process representation of a self-exciting process*, J. Appl. Probability, 11 (1974), pp. 498–503.
- [28] P. HOLME, *Modern temporal network theory: A colloquium*, Eur. Phys. J. B, 88 (2015), 234.
- [29] L. G. JEUB, M. BAZZI, I. S. JUTLA, AND P. J. MUCHA, *A Generalized Louvain Method for Community Detection Implemented in MATLAB* (2011–2017; Version 2.1.1), <http://netwiki.amath.unc.edu/GenLouvain>.
- [30] M. KARSAI, H.-H. JO, AND K. KASKI, *Bursty Human Dynamics*, SpringerBriefs in Complexity, Springer International Publishing, Cham, 2018.
- [31] M. KIVELÄ, A. ARENAS, M. BARTHELEMY, J. P. GLEESON, Y. MORENO, AND M. A. PORTER, *Multilayer networks*, J. Complex Netw., 2 (2014), pp. 203–271.
- [32] M. KIVELÄ AND M. A. PORTER, *Estimating interevent time distributions from finite observation periods in communication networks*, Phys. Rev. E, 92 (2015), 052813.
- [33] D. KUANG, P. J. BRANTINGHAM, AND A. L. BERTOZZI, *Crime topic modeling*, Crime Science, 6 (2017), 12.
- [34] D. KUANG, J. CHOO, AND H. PARK, *Nonnegative matrix factorization for interactive topic modeling and document clustering*, in Partitional Clustering Algorithms, Springer International Publishing, Cham, 2015, pp. 215–243.
- [35] D. KUANG, C. DING, AND H. PARK, *Symmetric nonnegative matrix factorization for graph clustering*, in Proceedings of the 2012 SIAM International Conference on Data Mining, SIAM, Philadelphia, 2012, pp. 106–117, <https://doi.org/10.1137/1.9781611972825.10>.
- [36] E. L. LAI, D. MOYER, B. YUAN, E. FOX, B. HUNTER, A. L. BERTOZZI, AND P. J. BRANTINGHAM, *Topic time series analysis of microblogs*, IMA J. Appl. Math., 81 (2016), pp. 409–431.
- [37] S. L. LAURITZEN, *Graphical Models*, Clarendon Press, Oxford, 1996.
- [38] D. LAZER, A. PENTLAND, L. ADAMIC, S. ARAL, A.-L. BARABÁSI, D. BREWER, N. CHRISTAKIS, N. CONTRACTOR, J. FOWLER, M. GUTMANN, T. JEBARA, G. KING, M. MACY, D. ROY, AND M. VAN ALSTYNE, *Computational social science*, Science, 323 (2009), pp. 721–723.
- [39] E. LEWIS AND G. O. MOHLER, *A Nonparametric EM Algorithm for Multiscale Hawkes Processes*, preprint, 2011, available online at http://paleo.sscnet.ucla.edu/Lewis-Molher-EM_Preprint.pdf.
- [40] S. LINDERMAN AND R. ADAMS, *Discovering latent network structure in point process data*, in International Conference on Machine Learning, 2014, pp. 1413–1421.
- [41] S. MANGAN AND U. ALON, *Structure and function of the feed-forward loop network motif*, Proc. Natl. Acad. Sci. USA, 100 (2003), pp. 11980–11985.
- [42] B. MARK, G. RASKUTTI, AND R. WILLETT, *Network Estimation from Point Process Data*, preprint, <https://arxiv.org/abs/1802.04838>, 2018.
- [43] D. MARSAN AND O. LENGLINE, *Extending earthquakes' reach through cascading*, Science, 319 (2008), pp. 1076–1079.
- [44] R. MILO, S. SHEN-ORR, S. ITZKOVITZ, N. KASHTAN, D. CHKLOVSKII, AND U. ALON, *Network motifs: Simple building blocks of complex networks*, Science, 298 (2002), pp. 824–827.
- [45] G. O. MOHLER, *Marked point process hotspot maps for homicide and gun crime prediction in Chicago*, Int. J. Forecast, 30 (2014), pp. 491–497.
- [46] G. O. MOHLER, M. B. SHORT, P. J. BRANTINGHAM, F. P. SCHOENBERG, AND G. E. TITA, *Self-exciting point process modeling of crime*, J. Amer. Statist. Assoc., 106 (2011), pp. 100–108.
- [47] G. O. MOHLER, M. B. SHORT, S. MALINOWSKI, M. JOHNSON, G. E. TITA, A. L. BERTOZZI, AND P. J. BRANTINGHAM, *Randomized controlled field trials of predictive policing*, J. Amer. Statist. Assoc., 110 (2015), pp. 1399–1411.

- [48] P. J. MUCHA, T. RICHARDSON, K. MACON, M. A. PORTER, AND J.-P. ONNELA, *Community structure in time-dependent, multiscale, and multiplex networks*, Science, 328 (2010), pp. 876–878.
- [49] J. A. NELDER AND R. J. BAKER, *Generalized Linear Models*, John Wiley & Sons, Hoboken, NJ, 1972.
- [50] M. E. J. NEWMAN, *Finding community structure in networks using the eigenvectors of matrices*, Phys. Rev. E, 74 (2006), 036104.
- [51] M. E. J. NEWMAN, *Networks*, 2nd ed., Oxford University Press, Oxford, UK, 2018.
- [52] M. E. J. NEWMAN AND M. GIRVAN, *Finding and evaluating community structure in networks*, Phys. Rev. E, 69 (2004), 026113.
- [53] A. NOULAS, S. SCCELLATO, R. LAMBIOTTE, M. PONTIL, AND C. MASCOLO, *A tale of many cities: Universal patterns in human urban mobility*, PLoS ONE, 7 (2012), e37027.
- [54] Y. OGATA, *On Lewis' simulation method for point processes*, IEEE Trans. Inform. Theory, 27 (1981), pp. 23–31.
- [55] Y. OGATA, *Statistical models for earthquake occurrences and residual analysis for point processes*, J. Amer. Statist. Assoc., 83 (1988), pp. 9–27.
- [56] Y. OGATA, *Space-time point-process models for earthquake occurrences*, Ann. Inst. Statist. Math., 50 (1998), pp. 379–402.
- [57] T. P. PEIXOTO, *Bayesian Stochastic Blockmodeling*, preprint, <https://arxiv.org/abs/1705.10225>, 2018; in Advances in Network Clustering and Blockmodeling, P. Doreian, V. Batagelj, A. Ferligoj, eds., John Wiley & Sons, New York, in press.
- [58] P. O. PERRY AND P. J. WOLFE, *Point process modelling for directed interaction networks*, J. R. Stat. Soc. Ser. B. Stat. Methodol., 75 (2013), pp. 821–849.
- [59] M. A. PORTER AND J. P. GLEESON, *Dynamical Systems on Networks: A Tutorial*, Front. Appl. Dyn. Syst. Rev. Tutor. 4, Springer International Publishing, Cham, 2016.
- [60] M. A. PORTER AND S. D. HOWISON, *The Role of Network Analysis in Industrial and Applied Mathematics*, preprint, <https://arxiv.org/abs/1703.06843>, 2017.
- [61] M. A. PORTER, J.-P. ONNELA, AND P. J. MUCHA, *Communities in networks*, Notices Amer. Math. Soc., 56 (2009), pp. 1082–1097, 1164–1166.
- [62] A. REINHART, *Rejoinder: A review of self-exciting spatio-temporal point processes and their applications*, Statist. Sci., 33 (2018), pp. 330–333.
- [63] S. SCCELLATO, A. NOULAS, AND C. MASCOLO, *Exploiting place features in link prediction on location-based social networks*, in Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, New York, 2011, pp. 1046–1054.
- [64] F. P. SCHOENBERG, D. R. BRILLINGER, AND P. GUTTORP, *Point processes, spatial-temporal*, in Encyclopedia of Environmetrics 3, John Wiley & Sons, Chichester, UK, 2002, pp. 1573–1577.
- [65] M. B. SHORT, P. J. BRANTINGHAM, A. L. BERTOZZI, AND G. E. TITA, *Dissipation and displacement of hotspots in reaction-diffusion models of crime*, Proc. Natl. Acad. Sci. USA, 107 (2010), pp. 3961–3965.
- [66] A. SIMMA AND M. I. JORDAN, *Modeling events with cascades of Poisson processes*, in Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence, AUAI Press, Arlington, VA, 2010, pp. 546–555.
- [67] D. SORNETTE AND S. UTKIN, *Limits of declustering methods for disentangling exogenous from endogenous events in time series with foreshocks, main shocks, and aftershocks*, Phys. Rev. E, 79 (2009), 061110.
- [68] T. SQUARTINI, F. PICCIOLO, F. RUZZENENTI, AND D. GARLASCHELLI, *Reciprocity of weighted networks*, Sci. Rep., 3 (2013), 02729.
- [69] A. STOMAKHIN, M. B. SHORT, AND A. L. BERTOZZI, *Reconstruction of missing data in social networks based on temporal patterns of interactions*, Inverse Problems, 27 (2011), 115013.
- [70] A. STREHL AND J. GHOSH, *Cluster ensembles—A knowledge reuse framework for combining multiple partitions*, J. Mach. Learn. Res., 3 (2002), pp. 583–617.
- [71] P. SUNY, J. LI, Y. MAO, R. ZHANG, AND L. WANG, *Inferring Multiplex Diffusion Network via Multivariate Marked Hawkes Process*, preprint, <https://arxiv.org/abs/1809.07688>, 2018.
- [72] G. TITA, P. J. BRANTINGHAM, A. GALSTYAN, AND Y.-S. CHO, *Latent self-exciting point process model for spatial-temporal networks*, Discrete Contin. Dyn. Syst. Ser. B, 19 (2014), pp. 1335–1354.

- [73] A. L. TRAUD, E. D. KELSIC, P. J. MUCHA, AND M. A. PORTER, *Comparing community structure to characteristics in online collegiate social networks*, SIAM Rev., 53 (2011), pp. 526–543, <https://doi.org/10.1137/080734315>.
- [74] J. UGANDER, B. KARRER, L. BACKSTROM, AND C. MARLOW, *The Anatomy of the Facebook Social Graph*, preprint, <https://arxiv.org/abs/1111.4503>, 2011.
- [75] A. VEEN AND F. P. SCHOENBERG, *Estimation of space–time branching process models in seismology using an EM-type algorithm*, J. Amer. Statist. Assoc., 103 (2008), pp. 614–624.
- [76] B. WANG, X. LUO, F. ZHANG, B. YUAN, A. L. BERTOZZI, AND P. J. BRANTINGHAM, *Graph-Based Deep Modeling and Real Time Forecasting of Sparse Spatio-Temporal Data*, preprint, <https://arxiv.org/abs/1804.00684>, 2018.
- [77] K. ZHOU, H. ZHA, AND L. SONG, *Learning social infectivity in sparse low-rank networks using multi-dimensional Hawkes processes*, in Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics, 2013, pp. 641–649.
- [78] J. ZHUANG, Y. OGATA, AND D. VERE-JONES, *Stochastic declustering of space-time earthquake occurrences*, J. Amer. Statist. Assoc., 97 (2002), pp. 369–380.
- [79] J. ZHUANG, Y. OGATA, AND D. VERE-JONES, *Analyzing earthquake clustering features by using stochastic reconstruction*, J. Geophys. Res., 109 (2004), B05301.