

Convex Parameter Recovery for Interacting Marked Processes

Anatoli Juditsky, Arkadi Nemirovski, Liyan Xie, *Student Member, IEEE*, and Yao Xie[✉], *Member, IEEE*

Abstract—We introduce a new general modeling approach for multivariate discrete event data with categorical interacting marks, which we refer to as marked Bernoulli processes. In the proposed model, the probability of an event of a specific category to occur in a location may be influenced by past events at this and other locations. We do not restrict interactions to be positive or decaying over time as it is commonly adopted, allowing us to capture an arbitrary shape of influence from historical events, locations, and events of different categories. In our modeling, prior knowledge is incorporated by allowing general convex constraints on model parameters. We develop two parameter estimation procedures utilizing the constrained Least Squares (LS) and Maximum Likelihood (ML) estimation, which are solved using variational inequalities with monotone operators. We discuss different applications of our approach and illustrate the performance of proposed recovery routines on synthetic examples and a real-world police dataset.

Index Terms—Optimization methods, estimation theory, stochastic processes.

I. INTRODUCTION

DISCRETE events are a type of sequential data, where each data point is a tuple consisting of event time, location, and possibly category. Such event data is ubiquitous in modern applications, such as police data [1], electronic health records [2], and social network data [3], [4]. In modeling discrete events, we are particularly interested in estimating the interactions of events, such as triggering or inhibiting effects of past events on future events. For example, in crime event modeling, the triggering effect has been empirically verified; when a crime event happens, it makes future events more likely to happen in the neighborhood. Similar empirical observations have been made for other applications such as in biological neural networks, social networks [5], [6], financial networks [7], and spatio-temporal epidemiological processes [8].

A popular model for capturing *interactions* between discrete events is the so-called Hawkes processes [9]–[12].

Manuscript received May 11, 2020; revised November 12, 2020; accepted November 20, 2020. Date of publication November 30, 2020; date of current version January 7, 2021. The work of Anatoli Juditsky and Arkadi Nemirovski was supported in part by the MIAI Grenoble Alpes under Grant ANR-19-P3IA-0003; in part by the CNPq under Grant 401371/2014-0; and in part by the NSF under Grant CCF-1523768. The work of Liyan Xie and Yao Xie was supported by NSF CAREER under Grant CCF-1650913, Grant DMS-1938106, and Grant DMS-1830210. (*Corresponding author: Yao Xie*)

Anatoli Juditsky is with the LJK, Université Grenoble Alpes, Campus de Saint-Martin-d'Hères, 38401 Grenoble, France (e-mail: anatoli.juditsky@univ-grenoble-alpes.fr).

Arkadi Nemirovski, Liyan Xie, and Yao Xie are with the School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: arkadi.nemirovski@isye.gatech.edu; lxie49@gatech.edu; yao.xie@isye.gatech.edu).

This article has supplementary downloadable material available at <https://doi.org/10.1109/JSAIT.2020.3040999>, provided by the authors.

Digital Object Identifier 10.1109/JSAIT.2020.3040999

The Hawkes process is a type of mutually-exciting non-homogeneous point process with intensity function consisting of a deterministic part and a stochastic part depending on the past event. The stochastic part of the intensity function can capture the interactions of past events and the current event, and it may be parameterized in different ways. In a certain sense, Hawkes processes may be viewed as a point process analog to classical autoregression in time series analysis. Hawkes process has received much attention since it is quite general and can conveniently model interactions. For instance, in a network Hawkes process,¹ interactions between nodes are modeled using a directed weighted graph in which direction and magnitude of edges indicate direction and strength of influence of one node on another. Along this line, there are various generalizations that allow for other types of point process modeling, where different “link” functions are considered, such as self-correcting process, reactive process, and specialized process (see [12] for an overview).

Estimating the *interactions* of the past events and the current event is a fundamental problem for Bernoulli processes since it reveals the underlying temporal and spatial structures and predicts future events. There has been much prior work in estimating model parameters, assuming that interactions are shift-invariant and captured through *kernel functions*. Furthermore, various simplifying assumptions are typically made for the kernel functions, e.g., that the spatio-temporal interactions are decoupled (e.g., [5]), implying that the interaction kernel function is a product of the interaction over time and interaction over locations and can be estimated separately. It is often assumed that the temporal kernel function decays exponentially over time with an unknown decay rate [6], or it is completely specified [13]; thus, the problem focus is on estimating spatial interaction between locations. It is also commonly assumed that the interactions are positive, i.e., the interaction triggers rather than inhibit future events [14]. Such simplification, however, may impede capturing complex interaction effects between events. For instance, negative interaction or inhibition is well known to play a major role in neuronal connectivity [15]. The study of more complex modeling of spatial aspects, especially jointly with discrete marks, is still in infancy.

In this article, we present a general computational framework for estimating marked spatio-temporal processes with categorical marks. Motivated by Hawkes processes, we consider a model of a discrete-time process on a finite spatio-temporal grid, which we refer to as Bernoulli processes. A brief description of the proposed modeling is as follows. At each time t at a site k of the grid of the *M-state Bernoulli process* can be

¹When space is discretized, the spatio-temporal point process of a grid can be modeled as a network point process.

in one of $M + 1$ states – a ground state, in which “nothing happens,” or an event state if an event of one of M given types at every (discrete) time instant t takes place at the site. We assume that the probability distribution of the events at each location at time t is a (linear or nonlinear) function on the process history – past events at different sites at times from $t - d$ to $t - 1$, d being the memory parameter of the process. For instance, each site of a 1-state linear (vanilla) Bernoulli process can be in one of two states – 0 (no event) or 1 (event takes place). From the point of view of time series, this process can be seen as a vector autoregressive process. The observations at sites of the grid at time t are Bernoulli random variables with the conditional expectation (what is the same as the conditional probability of an event to take place) being a linear combination of historical states of the process at times $t - d$ to $t - 1$. The linear combination coefficients are unknown process parameters. This model can be seen as a natural simplification of the continuous-time Hawkes process, where spatio-temporal cells are so small that one can ignore the chances for two or more events occurring in a cell. A notable feature of our model is that prior information on the structure of interactions is represented by general convex constraints on the parameters,² allowing for very general types of structures of interactions. For instance, we can relax the nonnegativity restrictions on interaction parameters and/or avoid assumptions of monotone or exponential time decay of interactions commonly used in the literature. When the situation has a “network component” allowing to assume that interacting sites are pairs of neighboring nodes in a known graph, we can incorporate this information, for instance, by restricting the interaction coefficients for non-neighboring pairs of sites to be zero.

The considered model is related to information diffusion processes over continuous time, for example, non-linear Hawkes model [16], self-exciting processes over networks (see [12] for an overview), information diffusion networks [17], and multivariate stationary Hawkes processes [15]. Compared to these well-known models, time and space discretization leading to the spatio-temporal Bernoulli process is a considerable simplification that, nonetheless, leads to practical estimation routines that can be used in “real world” scenarios.

Various approaches to parametric and nonparametric estimation of spatio-temporal processes have been proposed in the literature. A line of work [14], [18], [19] consider non-parametric Hawkes process estimation based on the Expectation-Maximization (EM) algorithms and the Kernel method. Least-square estimates for link functions of continuous-time multivariate stationary Hawkes process are studied in [15]. There is also much work [20]–[22] considering the estimation in the Bayesian framework. In particular, [23] considers estimation in a Bernoulli model similar to the one we promote in this article using the Bayesian approach and impose prior distributions on parameters. Several authors consider the problem of sparse model estimation for point processes, see, e.g., [24], etc.

An important feature of the proposed models is that they allow for simple “computation-friendly” statistical inferences. Our approach to processing the resulting estimation problems is based on convex optimization, which leads to computationally efficient procedures. Our primary tools here are variational

inequalities (VI) with monotone operators.³ Specifically, we show that the parameters of spatio-temporal models can be recovered in a computationally efficient fashion by solving inequalities of this type, both in the cases of linear models (Sections II-A–II-E) and of nonlinear models satisfying certain monotonicity restrictions (Section II-F). In the linear case, our approach results in the usual Least Squares estimate⁴ which, of course, could be motivated without any references to VI’s. However, these references explain how to act in the nonlinear cases, where Least Squares, if applied, “as is” typically lose computational tractability. Aside from the VI-based approach, we consider the standard Maximum Likelihood estimation (Section III). In the linear case, computing the Maximum Likelihood estimate amounts to solving a convex optimization problem and thus is computationally efficient. On the other hand, in the nonlinear case, Maximum Likelihood estimation typically becomes problematic computationally, including the cases where the VI-based approach remains computation-friendly. (A notable exception is the spatio-temporal logistic model in which the Maximum Likelihood estimation reduces to solving the convex problem.) It should be added that finite-sample theoretical results on the statistical performance of the estimates we develop do not favor Maximum Likelihood as compared with the VI-based estimation.

Finally, we also demonstrate the good performance of our method on synthetic and real data. In particular, we study a real crime dataset in Atlanta, USA, to demonstrate the promise of our methods to recover interesting structures from real-data and predict the probability of crime incidents.

This article is organized as follows. We discuss the LS estimation of the network Bernoulli process in Section II, by introducing the model of the simplest Bernoulli process with $\{0, 1\}$ -valued mark in Section II-A. We then derive the Least Squares estimate (which under the circumstances is what VI-estimate boils down to) for this model in Section II-C and building data-driven confidence sets for the estimated parameters in Section II-D. We describe the general model of the M -state Bernoulli process and discuss the Least Squares estimation in Section II-E. The nonlinear modeling of the process and the corresponding VI estimate are presented in Section II-F. In Section III, we discuss the Maximum Likelihood estimate of parameters of the general Bernoulli process. The application of the proposed approach is illustrated by various simulation examples in Section IV-A. Finally, Section IV-B shows an application of our modeling to “real-world” data analysis of crime events in Atlanta.

II. ESTIMATING PARAMETERS OF SPATIO-TEMPORAL BERNOUlli PROCESS

Here we consider spatio-temporal Bernoulli process with discrete-time over discrete locations. Specifically, we assume that the discrete-time and location grid we deal with is fine enough so that we can neglect the possibility for more than one event to occur in a cell of the grid. We will model the interactions of these events in the grid. In Sections II-C–II-E we develop and process *linear* models; nonlinear extensions are considered in Section II-F.

³Utilizing VI’s with monotone operators for statistical estimation is the main novelty in this article; to the best of our knowledge in statistics, this approach was used only once (see paper [25] on Generalized Linear Models).

⁴This, in hindsight, is resembling but not identical to what is done in [15].

²Convexity is assumed for the sake of computational tractability.

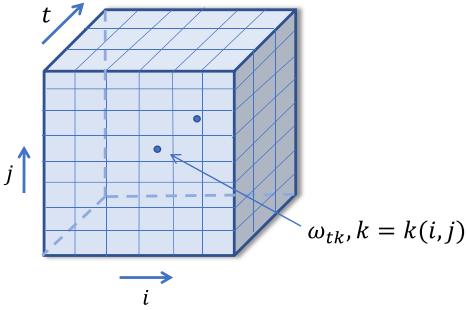


Fig. 1. Illustration of the discretized process. Observation ω_{tk} , at the location of a three-dimensional spatio-temporal grid.

A. Single-State Model

Define a *spatio-temporal Bernoulli process with memory depth d* as follows. We assume the memory depth is a pre-specified hyper-parameter (e.g., it can be estimated using cross-validation as explained in Section IV-B when we study real data). We observe on discrete time horizon $\{t : -d+1 \leq t \leq N\}$ random process as follows. At time t we observe Boolean vector $\omega_t \in \mathbf{R}^K$ with entries $\omega_{tk} \in \{0, 1\}$, $1 \leq k \leq K$. Here $\omega_{tk} = 1$ and $\omega_{tk} = 0$ mean, respectively, that at time t in location k an event took/did not take place. We set

$$\begin{aligned}\omega^t &= \{\omega_{sk}, -d+1 \leq s \leq t, 1 \leq k \leq K\} \in \mathbf{R}^{(t+d) \times K}, \\ \omega_\tau^t &= \{\omega_{sk}, \tau \leq s \leq t, 1 \leq k \leq K\} \in \mathbf{R}^{(t-\tau+1) \times K}.\end{aligned}$$

In other words, ω^t denotes all observations (at all locations) until current time t , and ω_τ^t contains observations on time horizon from τ to t .

We assume that for $t \geq 1$ the conditional probability of the event $\omega_{tk} = 1$, given the history ω^{t-1} , is specified as

$$\beta_k + \sum_{s=1}^d \sum_{\ell=1}^K \beta_{k\ell}^s \omega_{(t-s)\ell}, \quad 1 \leq k \leq K, \quad (1)$$

where $\beta = \{\beta_k, \beta_{k\ell}^s : 1 \leq s \leq d, 1 \leq k, \ell \leq K\}$ is a collection of coefficients. Here

- β_k corresponds to the *baseline intensity* at the k -th location (i.e., the intrinsic probability for an event to happen at a location without the exogenous influence, also called the birthrate);
- $\beta_{k\ell}^s$ captures the magnitude of the *influence* of an event that occurs at time $t-s$ at the ℓ -th location on chances for an event to happen at time t in the k -th location; so the sum in (1) represents the cumulative influence of past events at the k -th location.

Since the probability of occurrence is between 0 and 1, we require the coefficients to satisfy

$$\begin{aligned}0 &\leq \beta_k + \sum_{s=1}^d \sum_{\ell=1}^K \min[\beta_{k\ell}^s, 0], \quad \forall k \leq K, \\ 1 &\geq \beta_k + \sum_{s=1}^d \sum_{\ell=1}^K \max[\beta_{k\ell}^s, 0], \quad \forall k \leq K.\end{aligned} \quad (2)$$

Note that constraints in (2) allow some of the coefficients $\beta_{k\ell}^s$ to be negative, permitting the corresponding model to capture the inhibitive effect of past events. Fig. 2 illustrates a realization of the sample path of a simple Bernoulli process in the considered setting with different memory depths (5 for

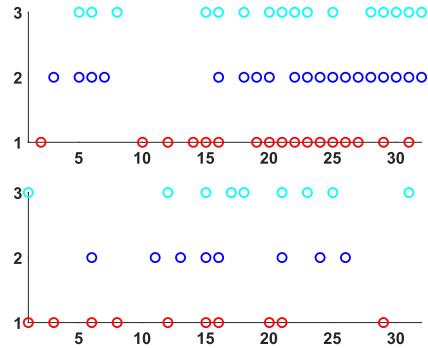


Fig. 2. Realizations of spatio-temporal Bernoulli processes with memory depths 5 (top) and 0 (bottom) on time horizon $N = 32$ with three locations represented with y-axis 1, 2, and 3. “1” events in different locations are marked by different colors.

the top figure and 0 for the bottom). Note that in the bottom plot, the events are more spread out due to the memoryless nature of the process.

Our goal is to recover the collection of parameters β using a set of observations ω^N over a time horizon N .

B. Preliminaries on Variational Inequalities With Monotone Operators

Variational inequalities (VI's) with monotone operators is the principal computational tool of the approach we are about to describe. We start with the related preliminaries. A vector field $F : \mathcal{X} \rightarrow \mathbf{R}^N$ defined on a nonempty convex subset \mathcal{X} of \mathbf{R}^N is called *monotone*, if $\langle F(x) - F(y), x - y \rangle \geq 0$ whenever $x, y \in \mathcal{X}$. When $N = 1$, monotonicity means that the scalar function F is nondecreasing on \mathcal{X} ; a basic example (by far not the only useful one) of a multivariate monotone vector field is the gradient field of a differentiable convex function $f : \mathcal{X} \rightarrow \mathbf{R}$. We say that $\alpha \geq 0$ is a *modulus of strong monotonicity* of vector field F , when

$$\langle F(x) - F(y), x - y \rangle \geq \alpha \|x - y\|_2^2 \quad \forall x, y \in \mathcal{X};$$

when $\alpha > 0$, F is called *strongly monotone*. A pair (\mathcal{X}, F) comprised of nonempty convex domain \mathcal{X} and monotone vector field F on this domain gives rise to *variational inequality* $VI(F, \mathcal{X})$. A *weak solution* to this VI is any point $\bar{x} \in \mathcal{X}$ such that

$$\langle F(x), x - \bar{x} \rangle \geq 0 \quad \forall x \in \mathcal{X}.$$

Whenever F is strongly monotone, weak solution, if exists, is unique.

A *strong solution* is a point $\bar{x} \in \mathcal{X}$ such that

$$\langle F(\bar{x}), x - \bar{x} \rangle \geq 0 \quad \forall x \in \mathcal{X}.$$

Every strong solution is a weak one; when F is continuous on \mathcal{X} , the inverse also is true. When \mathcal{X} is a convex *compact* set, $VI(F, \mathcal{X})$ always has weak solutions. When F is the gradient field of a continuously differentiable convex function f on \mathcal{X} , the weak and the strong solutions to $VI(F, \mathcal{X})$ are exactly the minimizers of f on \mathcal{X} . Finally, we should stress that variational inequalities with monotone operators are the most general “problems with convex structure;” under mild computability assumptions, that can be efficiently solved to a high accuracy.

C. Least Squares (LS) Estimation

As applied to the simple spatio-temporal model described in Section II-A the VI-based approach we are developing boils down to the Least Squares (LS) estimation. Let $\kappa = K + dK^2$; we arrange all reals from the collection β in (1) into a column vector (still denoted β):

$$\beta = \left[\beta_1, \dots, \beta_K, \beta_{11}^1, \dots, \beta_{11}^d, \beta_{1K}^1, \dots, \beta_{1K}^d, \dots, \beta_{KK}^1, \dots, \beta_{KK}^d \right]^T \in \mathbb{R}^\kappa.$$

Note that constraints (2) above state that β must reside in the polyhedral set \mathcal{B} given by explicit polyhedral representation.⁵ Assume that we are given a convex compact set $\mathcal{X} \subset \mathcal{B}$ such that $\beta \in \mathcal{X}$; we introduce this set to account for additional to the obvious inclusion $\beta \in \mathcal{B}$ a priori information, if any, on the vector of model's parameters. Our model says that for $t \geq 1$, the conditional expectation of ω_t given ω^{t-1} is $\eta^T(\omega_{t-d}^{t-1})\beta$,

$$\text{Prob}_{\omega^{t-1}}\{\omega_t = 1\} = \eta^T(\omega_{t-d}^{t-1})\beta,$$

with a known to us function $\eta(\cdot)$ which is defined on the set of all zero-one arrays $\omega_{t-d}^{t-1} \in \{0, 1\}^{d \times K}$ and takes values in the matrix space $\mathbb{R}^{\kappa \times K}$:

$$\eta^T(\omega_{t-d}^{t-1}) = \left[I_K, I_K \otimes \text{vec}(\omega_{t-d}^{t-1})^T \right] \in \mathbb{R}^{K \times \kappa}, \quad (3)$$

where I_K is a $K \times K$ identity matrix, \otimes denotes the standard Kronecker product, and $\text{vec}(\cdot)$ vectorizes a matrix by stacking all columns. Note that the matrix $\eta(\omega_{t-d}^{t-1})$ is Boolean and has at most one nonzero entry in every row.⁶

Consider a vector field $F : \mathcal{X} \rightarrow \mathbb{R}^\kappa$, defined as

$$F(x) = \frac{1}{N} \mathbf{E}_{\omega^N} \left\{ \sum_{t=1}^N \left[\eta(\omega_{t-d}^{t-1}) \eta^T(\omega_{t-d}^{t-1}) x - \eta(\omega_{t-d}^{t-1}) \omega_t \right] \right\} : \mathcal{X} \rightarrow \mathbb{R}^\kappa,$$

where \mathbf{E}_{ω^N} denotes expectation taken with respect to the distribution of ω^N (notation \mathbf{E}_{ω^t} is similarly defined). Below, all expectations and probabilities are conditional given a specific realization of the initial fragment ω_{-d+1}^0 of observations.

Observe that we have

$$\begin{aligned} \langle F(x) - F(y), x - y \rangle &= \frac{1}{N} \sum_{t=1}^N \mathbf{E}_{\omega^N} \\ &\times \left\{ (x-y)^T \eta(\omega_{t-d}^{t-1}) \eta^T(\omega_{t-d}^{t-1}) (x-y) \right\} \geq 0 \quad \forall x, y \in \mathcal{X}. \end{aligned}$$

⁵Polyhedral representation of a set $X \subset \mathbb{R}^n$ is a representation of the form

$$X = \{x \in \mathbb{R}^n : \exists w \in \mathbb{R}^m : Px + Qw \leq r\},$$

that is, representation of X as a projection of the solution set of a system of linear inequalities in the space of (x, w) -variables on the plane of x -variables. When X is polyhedrally representable, it automatically is polyhedral — can be represented by a finite system of linear inequalities in x -variables only. This system, however, can be much larger than the one in the polyhedral representation in question, making explicit polyhedral representations the standard descriptions of polyhedral sets in optimization.

⁶Indeed, (1) says that a particular entry in β , β_k or $\beta_{k\ell}^s$, affects at most one entry in $\eta^T(\omega_{t-d}^{t-1})\beta$, namely, the k -th entry, implying that each column of $\eta^T(\cdot)$ has at most one nonzero entry.

Thus, the vector field F is *monotone*. Moreover, we have $F(\beta) = 0$, since

$$\begin{aligned} F(\beta) &= \frac{1}{N} \mathbf{E}_{\omega^N} \left\{ \sum_{t=1}^N \eta(\omega_{t-d}^{t-1}) \left[\eta^T(\omega_{t-d}^{t-1})\beta - \omega_t \right] \right\} \\ &= \frac{1}{N} \sum_{t=1}^N \mathbf{E}_{\omega^t} \left\{ \eta(\omega_{t-d}^{t-1}) \left[\eta^T(\omega_{t-d}^{t-1})\beta - \mathbf{E}_{|\omega^{t-1}}\{\omega_t\} \right] \right\} \\ &= \frac{1}{N} \sum_{t=1}^N \mathbf{E}_{\omega^{t-1}} \left\{ \eta(\omega_{t-d}^{t-1}) \left[\eta^T(\omega_{t-d}^{t-1})\beta - \eta^T(\omega_{t-d}^{t-1})\beta \right] \right\} \\ &= 0, \end{aligned}$$

where $\mathbf{E}_{|\omega^{t-1}}$ denotes the conditional expectation given ω^{t-1} . Therefore, $\beta \in \mathcal{X}$ is a zero of the monotone operator F and therefore it is a solution to the variational inequality $\text{VI}[F, \mathcal{X}]$.

Now consider the empirical version

$$\begin{aligned} F_{\omega^N}(x) &= \underbrace{\left[\frac{1}{N} \sum_{t=1}^N \eta(\omega_{t-d}^{t-1}) \eta^T(\omega_{t-d}^{t-1}) \right] x}_{A[\omega^N]} \\ &\quad - \underbrace{\frac{1}{N} \sum_{t=1}^N \eta(\omega_{t-d}^{t-1}) \omega_t}_{a[\omega^N]} \end{aligned} \quad (4)$$

of vector field F . Note that $F_{\omega^N}(x)$ monotone and affine, and its expected value is $F(x)$ at every point x .

We propose to use, as an estimate of β , a weak solution to the Sample Average Approximation of $\text{VI}[F, \mathcal{X}]$, i.e., the variational inequality

$$\text{find } z \in \mathcal{X} : \langle F_{\omega^N}(w), w - z \rangle \geq 0, \quad \forall w \in \mathcal{X}. \text{ VI}[F_{\omega^N}, \mathcal{X}]$$

The monotone vector field $F_{\omega^N}(\cdot)$ is continuous (even affine), so that weak solutions to $\text{VI}[F_{\omega^N}, \mathcal{X}]$ are exactly the same as strong solutions, i.e., points $\bar{x} \in \mathcal{X}$ such that $\langle F_{\omega^N}(\bar{x}), x - \bar{x} \rangle \geq 0$ for all $x \in \mathcal{X}$. Moreover, the empirical vector field $F_{\omega^N}(x)$ is just the gradient field of the convex quadratic function

$$\Psi_{\omega^N}(x) = \frac{1}{2N} \sum_{t=1}^N \left\| \eta^T(\omega_{t-d}^{t-1})x - \omega_t \right\|_2^2, \quad (5)$$

so that weak (same as strong) solutions to $\text{VI}[F_{\omega^N}, \mathcal{X}]$ are just minimizers of this function on \mathcal{X} . In other words, our estimate based on solving variational inequality is an optimal solution to the Least Squares (LS) formulation: the constrained optimization problem

$$\min_{x \in \mathcal{X}} \Psi_{\omega^N}(x) \quad (6)$$

with a convex quadratic objective. Problem (6), the same as a general variational inequality with a monotone operator, can be routinely and efficiently solved by convex optimization algorithms.

D. Toward Performance Guarantees

Our objective in this section is to construct non-asymptotic confidence sets for parameter estimates built in the previous section. Utilizing concentration inequalities for martingales,

we can express these sets in terms of the process observations in the spirit of results of [24], [26], [27].

Observe that the vector of true parameters β underlying our observations not only solves variational inequality $\text{VI}[F, \mathcal{X}]$, but also solves the variational inequality $\text{VI}[\bar{F}_{\omega^N}, \mathcal{X}]$, where

$$\bar{F}_{\omega^N}(x) = A[\omega^N]x - \underbrace{\frac{1}{N} \sum_{t=1}^N \eta(\omega_{t-d}^{t-1}) \eta^T(\omega_{t-d}^{t-1}) \beta}_{\bar{a}[\omega^N]}$$

with $A[\omega^N]$ defined in (4).

In fact, β is just a root of $\bar{F}_{\omega^N}(x) : \bar{F}_{\omega^N}(\beta) = 0$. Moreover, the monotone affine operators $F_{\omega^N}(x)$ and $\bar{F}_{\omega^N}(x)$ differ only in the value of constant term: in $F_{\omega^N}(x)$ this term is $a[\omega^N]$, and in $\bar{F}_{\omega^N}(x)$ this term is $\bar{a}[\omega^N]$. Thus, equivalently, β is the minimizer on \mathcal{X} of the quadratic form

$$\Psi_{\omega^N}(x) := \frac{1}{2N} \sum_{t=1}^N \left\| \eta^T(\omega_{t-d}^{t-1})x - \eta^T(\omega_{t-d}^{t-1})\beta \right\|_2^2,$$

and the functions Ψ in (5) and $\bar{\Psi}$ above differ only in the constant terms (which do not affect the results of minimization) and in the linear terms. Moreover, the difference of the vectors of coefficients of linear terms is given by (due to $\bar{F}_{\omega^N}(\beta) = 0$):

$$\begin{aligned} \Delta_F &:= F_{\omega^N}(\beta) - \bar{F}_{\omega^N}(\beta) = F_{\omega^N}(\beta) = \bar{a}[\omega^N] - a[\omega^N] \\ &= \frac{1}{N} \sum_{t=1}^N \underbrace{\eta(\omega_{t-d}^{t-1})}_{\xi_t} \left[\eta^T(\omega_{t-d}^{t-1})\beta - \omega_t \right]. \end{aligned} \quad (7)$$

Note that this is the same as the difference of constant terms in $F_{\omega^N}(\cdot)$ and $\bar{F}_{\omega^N}(\cdot)$.

Concentration bounds for $F_{\omega^N}(\beta)$ can be obtained by applying general Bernstein-type inequalities for martingales.

Lemma 1: For all $\epsilon \in (0, 1)$ vector $F_{\omega^N}(\beta) = \Delta_F$ in (7) satisfies

$$\text{Prob}_{\omega^N} \left\{ \|F_{\omega^N}(\beta)\|_{\infty} \geq \sqrt{\frac{\ln(2\kappa/\epsilon)}{2N}} + \frac{\ln(2\kappa/\epsilon)}{3N} \right\} \leq \epsilon. \quad (8)$$

Proof: Since the conditional expectation of ω_t given ω^{t-1} is $\eta^T(\omega_{t-d}^{t-1})\beta$, we have $\mathbf{E}_{|\omega^{t-1}}[\xi_t] = 0$. Thus, ξ_t is a *martingale-difference*. Also, because both ω_t and $\eta^T(\omega_{t-d}^{t-1})\beta$ are vectors with nonnegative entries not exceeding 1, we have $\|\eta^T(\omega_{t-d}^{t-1})\beta - \omega_t\|_{\infty} \leq 1$. Besides this, $\eta(\omega_{t-d}^{t-1})$ is a Boolean matrix with at most one nonzero in every row, whence $\|\eta^T(\omega_{t-d}^{t-1})z\|_{\infty} \leq \|z\|_{\infty}$ for all z . The bottom line is that $\|\xi_t\|_{\infty} \leq 1$. Furthermore, the conditional variance of components of ω_t is bounded by 1/4, so, applying the Azuma-Hoeffding inequality [28] to components $(F_{\omega^N}(\beta))_k$, $k = 1, \dots, \kappa$, of $F_{\omega^N}(\beta)$ we conclude that

$$\begin{aligned} \text{Prob}_{\omega^N} \left\{ |(F_{\omega^N}(\beta))_k| \geq \sqrt{\frac{x}{2N}} + \frac{x}{3N} \right\} &\leq 2 \exp\{-x\} \\ \forall 1 \leq k \leq \kappa, x \geq 0. \end{aligned}$$

The latter bound results in (8) by application of the total probability formula. \blacksquare

A somewhat finer analysis allows to establish more precise data-driven deviation bounds for components of $F_{\omega^N}(\beta)$.

Lemma 2: For all $y > 1$ entries $F_{\omega^N}(\beta)_k$, $k = 1, \dots, \kappa$, of $F_{\omega^N}(\beta)$ satisfy, with probability at least $1 - 2e(y[\ln((y-1)N) + 2] + 2) e^{-y}$,

$$\begin{aligned} a[\omega^N]_k - \bar{\psi}(a[\omega^N]_k, N; y) \\ \leq F_{\omega^N}(\beta)_k \leq a[\omega^N]_k - \underline{\psi}(a[\omega^N]_k, N; y) \end{aligned} \quad (9)$$

where $a[\omega^N]_k$ is the k -th component of $a[\omega^N]$ as in (4) and lower and upper functions $\underline{\psi}(\cdot)$, $\bar{\psi}(\cdot)$ are defined in relation (??), see the Appendix.

Proof of Lemma 2 is postponed till the Appendix. We are about to extract from this lemma upper bounds on the accuracy of recovered coefficients.

1) *Upper-Bounding Risk of Recovery:* Recall that our estimate $\hat{\beta} := \hat{\beta}(\omega^N)$ solves the variational inequality $\text{VI}[F_{\omega^N}, \mathcal{X}]$ with $F_{\omega^N}(x) = A[\omega^N]x - a[\omega^N]$, see (4). Note that $A[\omega^N]$ is positive semidefinite (we write $A \succeq 0$, and we write $A > 0$ for positive definite A). Given $A \in \mathbf{R}^{\kappa \times \kappa}$, $A \succeq 0$, and $p \in [1, \infty]$, define the “condition number”

$$\theta_p[A] := \max \left\{ \theta \geq 0 : g^T A g \geq \theta \|g\|_p^2 \quad \forall g \in \mathbf{R}^{\kappa} \right\}. \quad (10)$$

Observe that $\theta_p[A] > 0$ whenever $A > 0$, and that for $p, p' \in [1, \infty]$ one has

$$\begin{aligned} g^T A g &\geq \frac{1}{2} \left\{ \theta_p[A] \|g\|_p^2 + \theta_{p'}[A] \|g\|_{p'}^2 \right\} \\ &\geq \sqrt{\theta_p[A] \theta_{p'}[A]} \|g\|_p \|g\|_{p'}. \end{aligned} \quad (11)$$

The following result is immediate:

Theorem 1 (Bounding ℓ_p estimation error): For every $p \in [1, \infty]$ and every ω^N one has

$$\|\hat{\beta}(\omega^N) - \beta\|_p \leq \|F_{\omega^N}(\beta)\|_{\infty} / \sqrt{\theta_p[A[\omega^N]] \theta_1[A[\omega^N]]}. \quad (12)$$

As a result, for every $\epsilon \in (0, 1)$, the probability of the event

$$\|\hat{\beta}(\omega^N) - \beta\|_p \leq (\theta_p[A[\omega^N]] \theta_1[A[\omega^N]])^{-1} \times \left(\sqrt{\frac{\ln(2\kappa/\epsilon)}{2N}} + \frac{\ln(2\kappa/\epsilon)}{3N} \right) \quad \forall p \in [1, \infty] \quad (13)$$

is at least $1 - \epsilon$.

Proof: Let us fix ω^N and set $\hat{\beta} = \hat{\beta}(\omega^N)$, $A = A[\omega^N]$. Since $F_{\omega^N}(\cdot)$ is continuous and $\hat{\beta}$ is a weak solution to $\text{VI}[F_{\omega^N}, \mathcal{X}]$, $\hat{\beta}$ is also a strong solution: $\langle F_{\omega^N}(\hat{\beta}), z - \hat{\beta} \rangle \geq 0$ for all $z \in \mathcal{X}$; in particular, $\langle F_{\omega^N}(\hat{\beta}), \beta - \hat{\beta} \rangle \geq 0$. On the other hand, $F_{\omega^N}(\hat{\beta}) = F(\beta) - A(\beta - \hat{\beta})$. As a result, $0 \leq \langle F_{\omega^N}(\hat{\beta}), \beta - \hat{\beta} \rangle = \langle F_{\omega^N}(\beta) - A(\beta - \hat{\beta}), \beta - \hat{\beta} \rangle$, whence

$$\begin{aligned} (\beta - \hat{\beta})^T A(\beta - \hat{\beta}) &\leq \langle F_{\omega^N}(\beta), \beta - \hat{\beta} \rangle \\ &\leq \|F_{\omega^N}(\beta)\|_{\infty} \|\beta - \hat{\beta}\|_1. \end{aligned} \quad (14)$$

Setting $p' = 1$ in (11), we obtain

$$(\beta - \hat{\beta})^T A(\beta - \hat{\beta}) \geq \sqrt{\theta_1[A] \theta_p[A]} \|\beta - \hat{\beta}\|_1 \|\beta - \hat{\beta}\|_p.$$

This combines with (14) to imply (12); then (12) together with (8) imply (13). \blacksquare

Remark 1: (Evaluating the condition number): To assess the upper bound (13) one needs to compute “condition numbers” $\theta_p[A]$ of a positive definite matrix A . The computation

is easy when $p = 2$, in which case $\theta_2[A]$ is the minimal eigenvalue of A , and when $p = \infty$:

$$\theta_\infty[A] = \min_{1 \leq i \leq \kappa} \{x^T A x : \|x\|_\infty \leq 1, x_i = 1\}$$

is the minimum of κ efficiently computable quantities. In general, $\theta_1[A]$ is difficult to compute, but this quantity admits an efficiently computable tight within the factor $\pi/2$ lower bound. Specifically, for a symmetric positive definite A , $\min_{z \in \mathbb{R}^N} z^T A z : \|z\|_1 = 1\}$ is the largest $r > 0$ such that the ellipsoid $\{z : z^T A z \leq r\}$ is contained in the unit $\|\cdot\|_1$ -ball, or, passing to polars, the largest r such that the ellipsoid $y^T A^{-1} y \leq r^{-1}$ contains the unit $\|\cdot\|_\infty$ -ball. Because of this, the definition of $\theta_1[A]$ in (10) is equivalent to $\theta_1[A] = [\max_{\|x\|_\infty \leq 1} x^T A^{-1} x]^{-1}$. It remains to note that when Q is a symmetric positive semidefinite $\kappa \times \kappa$ matrix, the efficiently computable by semidefinite relaxation upper bound on $\max_{\|x\|_\infty \leq 1} x^T Q x$, given by

$$\min_{\lambda} \left\{ \sum_i \lambda_i : \lambda_i \geq 0 \quad \forall i; \text{Diag}\{\lambda_1, \dots, \lambda_\kappa\} \succeq Q \right\},$$

is tight within the factor $\pi/2$, see [29].

Under favorable circumstances, we can expect that for large N the minimal eigenvalue of $A[\omega^N]$ will be of the order of one with overwhelming probability implying that the lengths of the confidence intervals (16) go to 0 as $N \rightarrow \infty$ at the rate $O(1/\sqrt{N})$. Note, however, that inter-dependence of the “regressors” $\eta(\omega_{t-d}^{t-1})$ across t makes it difficult to prove something along these lines.

2) *Estimating Linear Forms of β :* We can use concentration bounds of Lemmas 1 and 2 to build confidence intervals for linear functionals of β . For instance, inequality (9) of Lemma 2 leads to the following estimation procedure of the linear form $e(\beta) = e^T \beta$, $e \in \mathbb{R}^\kappa$. Given $y > 1$, consider the pair of optimization problems

$$\begin{aligned} \underline{e}[\omega^N, y] &= \min_x \left\{ e^T x : \begin{array}{l} x \in \mathcal{X}, \\ \underline{\psi}(a[\omega^N]_k, N; y) \leq (A[\omega^N]x)_k \\ \leq \bar{\psi}(a[\omega^N]_k, N; y), \quad k = 1, \dots, \kappa, \end{array} \right\} \\ \bar{e}[\omega^N, y] &= \max_x \left\{ e^T x : \begin{array}{l} x \in \mathcal{X}, \\ \underline{\psi}(a[\omega^N]_k, N; y) \leq (A[\omega^N]x)_k \\ \leq \bar{\psi}(a[\omega^N]_k, N; y), \quad k = 1, \dots, \kappa \end{array} \right\} \end{aligned} \quad (15)$$

where $\underline{\psi}(\cdot)$ and $\bar{\psi}(\cdot)$ are defined as in (??) of the Appendix. These problems clearly are convex, so $\underline{e}[\omega^N, y]$ and $\bar{e}[\omega^N, y]$ are efficiently computable. Immediately, we have the following.

Lemma 3: Given $y > 1$, the probability of the event

$$\underline{e}[\omega^N, y] \leq e^T \beta \leq \bar{e}[\omega^N, y] \quad \forall e, \quad (16)$$

is at least $1 - 2\kappa e(y[\ln((y-1)N) + 2] + 2)e^{-y}$.

Indeed, when events

$$\begin{aligned} a[\omega^N]_k - \bar{\psi}(a[\omega^N]_k, N; y) &\leq F_{\omega^N}(\beta)_k \\ &\leq a[\omega^N]_k - \underline{\psi}(a[\omega^N]_k, N; y), \quad k = 1, \dots, \kappa, \end{aligned}$$

take place, β is a feasible solution to optimization problems in (15). Due to Lemma 2, this implies that (16) takes place with probability at least $1 - 2\kappa e(y[\ln((y-1)N) + 2] + 2)e^{-y}$.

E. Estimating Parameters of Multi-State Spatio-Temporal Processes

In this section, we consider the multi-state spatio-temporal process in which an event outcome contains additional information about its category [19]. So far, we considered the case where at every time instant t every location k maybe be either in the state $\omega_{tk} = 0$ (“no event”), or $\omega_{tk} = 1$ (“event”). We are now extending the model by allowing the state of a location at a given time instant to take $M \geq 2$ “nontrivial” values on the top of the zero value “no event.” In other words, observation of the multi-state Bernoulli process is categorical — we can either observe no event or observe one of M possible event outcomes.

We define *M-state spatio-temporal process with memory depth d* as follows:

- We observe a random process on time horizon $\{t : -d + 1 \leq t \leq N\}$, observation at time t being

$$\omega_t = \{\omega_{tk} \in \{0, 1, \dots, M\}, 1 \leq k \leq K\}.$$

- For every $t \geq 1$, the conditional, $\omega^{t-1} = (\omega_{-d+1}, \omega_{-d+2}, \dots, \omega_{t-1})$ given, distribution of ω_{tk} is defined as follows. With every location k , we associate an array of (baseline) parameters $\beta_k = \{\beta_k(p), 1 \leq p \leq M\}$, and with every pair of locations k, ℓ and every $s \in \{1, \dots, d\}$ — an array of (interaction) parameters $\beta_{k\ell}^s = \{\beta_{k\ell}^s(p, q), 1 \leq p \leq M, 0 \leq q \leq M\}$. Then induced by ω^{t-1} probability of ω_{tk} to be of category p , $1 \leq p \leq M$, is given by

$$\text{Prob}_{\omega^{t-1}}\{\omega_{tk} = p\} = \beta_k(p) + \sum_{s=1}^d \sum_{\ell=1}^K \beta_{k\ell}^s(p, \omega_{(t-s)\ell}), \quad (17)$$

and the probability for ω_{tk} to take value 0 (no event or “ground event”) is the complementary probability

$$\begin{aligned} \text{Prob}_{\omega^{t-1}}\{\omega_{tk} = 0\} &= 1 - \sum_{p=1}^M \left[\beta_k(p) + \sum_{s=1}^d \sum_{\ell=1}^K \beta_{k\ell}^s(p, \omega_{(t-s)\ell}) \right]. \end{aligned}$$

In other words, $\beta_{k\ell}^s(p, q)$ is the contribution of the location ℓ in state $q \in \{0, 1, \dots, M\}$ at time $t-s$ to the probability for the location k to be in state $p \in \{1, \dots, M\}$ at time t , and $\beta_k(p)$, $p \in \{1, \dots, M\}$ is the “endogenous” component of the probability of the latter event. Of course, for this description to make sense, the β -parameters should guarantee that for every ω^{t-1} , that is, for every collection $\{\omega_{\tau\ell} \in \{0, 1, \dots, M\} : \tau < t, 1 \leq \ell \leq K\}$, the prescribed by (17) probabilities are nonnegative and their sum over $p = 1, \dots, M$ is ≤ 1 . Thus, the β -parameters should satisfy the system of constraints

$$\begin{aligned} 0 &\leq \beta_k(p) + \sum_{s=1}^d \sum_{\ell=1}^K \min_{0 \leq q \leq M} \beta_{k\ell}^s(p, q), \\ 1 &\leq p \leq M, \quad 1 \leq k \leq K, \\ 1 &\geq \sum_{p=1}^M \beta_k(p) + \sum_{s=1}^d \sum_{\ell=1}^K \max_{0 \leq q \leq M} \sum_{p=1}^M \beta_{k\ell}^s(p, q), \\ 1 &\leq k \leq K. \end{aligned} \quad (18)$$

The solution set \mathcal{B} of this system is a polyhedral set given by explicit polyhedral representations.

- We are given convex compact set \mathcal{X} in the space of parameters $\beta = \{\beta_k, \beta_{k\ell}^s(p, q), 1 \leq s \leq d, 1 \leq k, \ell \leq K, 1 \leq p \leq M, 0 \leq q \leq M\}$ such that \mathcal{X} contains the true parameter β of the process we are observing, and \mathcal{X} is contained in the polytope \mathcal{B} given by constraints (18).

We arrange the collection of β -parameters associated with a M -state spatio-temporal process with memory depth d into a column vector (still denoted β) and denote by κ the dimension of β .⁷ Note that (17) says that the M -dimensional vector of conditional probabilities for ω_{tk} to take values $p \in \{1, \dots, M\}$ given ω^{t-1} is

$$\left[\eta_k^T(\omega_{t-d}^{t-1}) \beta \right]_p$$

with known to us function $\eta_k(\cdot)$ defined on the set of arrays $\omega_{t-d}^{t-1} \in \{0, 1, \dots, M\}^{d \times K}$ and taking values in the space of $\kappa \times M$ matrices. Note that the value of ω_{tk} is the index of the category, and does not mean magnitude. Same as above, $\eta_k(\omega_{d-1}^{t-1})$ is a Boolean matrix.

To proceed, for $0 \leq q \leq M$, let $\chi_q \in \mathbf{R}^M$ be defined as follows: $\chi_0 = 0 \in \mathbf{R}^M$, and $\chi_q, 1 \leq q \leq M$, is the q -th vector of the standard basis in \mathbf{R}^M . In particular, the state ω_{tk} can be encoded by vector $\bar{\omega}_{tk} = \chi_{\omega_{tk}}$, and the state of our process at time t — by the block vector $\bar{\omega}_t \in \mathbf{R}^{MK}$ with blocks $\bar{\omega}_{tk} \in \mathbf{R}^M, k = 1, \dots, K$. In other words: the k -th block in $\bar{\omega}_t$ is an M -dimensional vector which is the p -th basic orth of \mathbf{R}^M when $\omega_{tk} = p \geq 1$, and is the zero vector when $\omega_{tk} = 0$. Arranging $\kappa \times M$ matrices $\eta_k(\cdot)$ into a matrix

$$\eta(\cdot) = [\eta_1(\cdot), \dots, \eta_K(\cdot)] \in \{0, 1\}^{\kappa \times MK},$$

we obtain

$$\mathbf{E}_{|\omega^{t-1}}\{\bar{\omega}_t\} = \eta^T(\omega_{t-d}^{t-1})\beta \in \mathbf{R}^{MK},$$

where $\mathbf{E}_{|\omega^{t-1}}$ is the conditional expectation given ω^{t-1} . Note that similarly to Section II-A, (17) says that every particular entry in β , $\beta_k(p)$ or $\beta_{k\ell}^s(p, q)$, affects at most one of the entries in the block vector $[\eta_1^T(\omega_{t-d}^{t-1})\beta; \dots; \eta_K^T(\omega_{t-d}^{t-1})\beta]$ specifically, the p -th entry of the k -th block, so that the Boolean matrix $\eta(\omega_{t-d}^{t-1})$ has at most one nonzero entry in every row.

Note that the spatio-temporal Bernoulli process with memory depth d , as defined in Section II-A, is a special case of M -state ($M = 1$) spatio-temporal process with memory depth d , the case where state 0 at a location contributes nothing to probability of state 1 in another location at a later time, that is, $\beta_{k\ell}^s(1, 0) = 0$ for all s, k, ℓ .

Motivating example: Different types of crime events: As an illustration, consider a spatio-temporal model of crime events of different types, e.g., burglary and robbery, in a geographic area of interest. We split the area into K non-overlapping cells, which will be our locations. Selecting the time step in such a way that we can ignore the chances for two or more crime events to occur in the same spatio-temporal cell, we can model the history of crime events in the area as a $M = 2$ -state spatio-temporal process, with additional to (18) convex restrictions on

⁷In general, $\kappa = KM + dK^2M^2$. However, depending on application, it could make sense to postulate that some of the components of β are zeros, thus reducing the actual dimension of β ; for example, we could assume that $\beta_{k\ell}(\cdot, \cdot) = 0$ for some “definitely non-interacting” pairs k, ℓ of locations.

the vector of parameters β expressing our *a priori* information on the probability $\beta_k(p)$ of a “newborn” crime event of category p to occur at time instant t at location k and on the contribution $\beta_{k\ell}^s(p, q)$ of a crime event of category q in spatio-temporal cell $\{t - s, \ell\}$ to the probability of crime event of category p , $p \geq 1$, to happen in the spatio-temporal cell $\{t, k\}$.

The problem of estimating parameters β of the M -state spatio-temporal process from observations of this process can be processed exactly as in the case of the single state spatio-temporal Bernoulli process. Specifically, observations ω^N give rise to two monotone and affine vector fields on \mathcal{X} , the first observable and the second unobservable:

$$\begin{aligned} F_{\omega^N}(x) &= \underbrace{\left[\frac{1}{N} \sum_{t=1}^N \eta(\omega_{t-d}^{t-1}) \eta^T(\omega_{t-d}^{t-1}) \right] x - \frac{1}{N} \sum_{t=1}^N \eta(\omega_{t-d}^{t-1}) \bar{\omega}_t}_{A[\omega^N]} \\ \bar{F}_{\omega^N}(x) &= A[\omega^N]x - A[\omega^N]\beta. \end{aligned} \quad (19)$$

The two fields differ only in constant term, β is a root of the second field, and the difference of constant terms, same as the vector $F_{\omega^N}(\beta)$ due to $\bar{F}_{\omega^N}(\beta) = 0$, are zero-mean satisfying, for exactly the same reasons as in Section II-D, concentration bounds (8) and (9) of Lemmas 1 and 2. To recover β from observations, we may use the Least Squares (LS) estimate obtained by solving variational inequality $\text{VI}[F_{\omega^N}, \mathcal{X}]$ with the just defined F_{ω^N} , or, which is the same, by solving

$$\min_{x \in \mathcal{X}} \left\{ \Psi_{\omega^N}(x) := \frac{1}{2N} \sum_{t=1}^N \left\| \eta^T(\omega_{t-d}^{t-1})x - \bar{\omega}_t \right\|_2^2 \right\}. \quad (20)$$

Note that (8) and (9), by the same argument as in Section II-D, imply the validity in our present situation of Theorem 1 and Lemma 3.

F. Nonlinear Link Function

So far, our discussion focused on “linear” link functions, where past events contribute additively to the probability of a specific event in a given spatio-temporal cell. We now consider the case of non-linear link functions. This generalizes our model to allow more complex spatio-temporal interactions.

1) *Single-State Process:* Let $\phi(\cdot) : D \rightarrow \mathbf{R}^K$ be a continuous monotone vector field defined on a closed convex domain $D \subset \mathbf{R}^K$ such that

$$y \in D \Rightarrow 0 \leq \phi(y) \leq [1; \dots; 1].$$

For example, we may consider “sigmoid field” $\phi(u) = [\phi_1(u); \dots; \phi_K(u)]$ with

$$[\phi(u)]_k = \frac{\exp\{u_k\}}{1 + \exp\{u_k\}}, \quad k \leq K, \quad D = \mathbf{R}^K.$$

Given positive integer N , we define a *spatio-temporal Bernoulli process with memory depth d and link function ϕ* as a random process with realizations $\{\omega_{tk} \in \{0, 1\}, k \leq K, -d + 1 \leq t \leq N\}$ in the same way it was done in Section II-A with assumptions of Section II-A replaced with the following:

- we are given a convex compact set $\mathcal{X} \subset \mathbf{R}^K$ such that the vector of parameters β underlying the observed process belongs to \mathcal{X} and every $\beta \in \mathcal{X}$ satisfies

$$\eta^T(\omega_{t-d}^{t-1})\beta \in D \quad \forall 1 \leq t \leq N \quad (21)$$

with given functions $\eta(\omega_{t-d}^{t-1})$ taking values in the space of $\kappa \times K$ matrices;

- the conditional expectation of $\omega_t \in \{0, 1\}^K$ given ω^{t-1} is $\phi(\eta^T(\omega_{t-d}^{t-1})\beta)$.

Let us set

$$\begin{aligned} F(x) &= \frac{1}{N} \mathbf{E}_{\omega^N} \left\{ \sum_{t=1}^N \left[\eta(\omega_{t-d}^{t-1}) \phi(\eta^T(\omega_{t-d}^{t-1})x) \right. \right. \\ &\quad \left. \left. - \eta(\omega_{t-d}^{t-1}) \omega_t \right] \right\} : \mathcal{X} \rightarrow \mathbf{R}^\kappa, \\ F_{\omega^N}(x) &= \underbrace{\frac{1}{N} \sum_{t=1}^N \eta(\omega_{t-d}^{t-1}) \phi(\eta^T(\omega_{t-d}^{t-1})x)}_{A_{\omega^N}(x)} \\ &\quad - \underbrace{\frac{1}{N} \sum_{t=1}^N \eta(\omega_{t-d}^{t-1}) \omega_t}_{a[\omega^N]} : \mathcal{X} \rightarrow \mathbf{R}^\kappa, \\ \bar{F}_{\omega^N}(x) &= A_{\omega^N}(x) \\ &\quad - \underbrace{\frac{1}{N} \sum_{t=1}^N \eta(\omega_{t-d}^{t-1}) \phi(\eta^T(\omega_{t-d}^{t-1})\beta)}_{\bar{a}[\omega^N]} : \mathcal{X} \rightarrow \mathbf{R}^\kappa. \end{aligned} \quad (22)$$

We are now essentially in the situation of Section II-C (where we considered the special case $\phi(z) \equiv z$ of our present situation). Specifically, $F(\cdot)$ is a monotone (albeit not affine) vector field on \mathcal{X} , $F(\beta) = 0$. The empirical version $F_{\omega^N}(x)$, for every $x \in \mathcal{X}$, is a monotone on \mathcal{X} vector field which is an unbiased estimate of $F(x)$. Besides this, $\bar{F}_{\omega^N}(x)$ is a monotone on \mathcal{X} vector field, and the true vector of parameters β underlying our observations solves the variational inequality $\text{VI}[\bar{F}_{\omega^N}, \mathcal{X}]$ (is a root of \bar{F}_{ω^N}). These observations suggest estimating β by weak solution to the variational inequality $\text{VI}[F_{\omega^N}, \mathcal{X}]$.

Note that, same as above, vector fields F_{ω^N} and \bar{F}_{ω^N} differ only in the constant terms, and this difference is nothing but $F_{\omega^N}(\beta)$ due to $\bar{F}_{\omega^N}(\beta) = 0$; moreover $\xi_t = \eta(\omega_{t-d}^{t-1})\omega_t - \eta^T(\omega_{t-d}^{t-1})\beta$ is a martingale difference. Though deviation probabilities for $F_{\omega^N}(\beta)$ do not obey the same bound as in the case of $\phi(z) \equiv z$ (since the matrices $\eta(\omega_{t-d}^{t-1})$ now not necessarily are Boolean with at most one nonzero in a row), the reasoning which led us to (8) demonstrates that the vector $F_{\omega^N}(\beta)$ in our present situation does obey the bound

$$\text{Prob}_{\omega^N} \left\{ \|F_{\omega^N}(\beta)\|_\infty \geq \Theta \left[\sqrt{\frac{\ln(2\kappa/\epsilon)}{2N}} + \frac{\ln(2\kappa/\epsilon)}{3N} \right] \right\} \leq \epsilon \quad \forall \epsilon \in (0, 1), \quad (23)$$

where Θ is the maximum, over all possible ω_{d-1}^{t-1} , of the $\|\cdot\|_1$ -norm of rows of $\eta(\omega_{t-d}^{t-1})$. Note that in the situation of this section, our $O(1/\sqrt{N})$ exponential bounds on large deviations of $F_{\omega^N}(\beta)$ from zero, while being good news, do *not* result in easy-to-compute on-line upper-risk bounds and confidence intervals for linear functions of β . Indeed, in order to adjust to our present situation Theorem 1, we need to replace the condition numbers $\theta_p[\cdot]$ with constants of strong monotonicity of the vector field $F_{\omega^N}(\cdot)$ on \mathcal{X} . On the other hand, to adopt the result of Lemma 3 in the present setting, we need

to replace the quantities \bar{e} and \underline{e} , see (15), with the maximum (resp., minimum) of the linear form $e^T x$ over the set $\{x \in \mathcal{X} : \|F_{\omega^N}(x)\|_\infty \leq \delta\}$. Both these tasks for a *nonlinear* operator $F_{\omega^N}(\cdot)$ seem to be problematic.

2) *Multi-State Processes*: The construction in the previous paragraph can be extended to M -state processes. Below, with a slight abuse of notation, we redefine notation for the multi-state processes.

Let us identify two-dimensional $K \times M$ array $\{a_{k\ell} : 1 \leq k \leq K, 1 \leq \ell \leq M\}$ with KM -dimensional block vector with K blocks $[a_{k1}; a_{k2}; \dots; a_{kM}]$, $1 \leq k \leq K$, of dimension M each. With this convention, a parametric $K \times M$ array $\psi(z) = \{\psi_{kp}(z) \in \mathbf{R} : k \leq K, 1 \leq p \leq M\}$ depending on KM -dimensional vector z of parameters becomes a vector field on \mathbf{R}^{KM} . Assume that we are given an array $\phi(\cdot) = \{\phi_{kp}(\cdot) \in \mathbf{R} : k \leq K, 1 \leq p \leq M\}$ of the outlined structure such that vector field $\phi(\cdot)$ is continuous and monotone on a closed convex domain $D \subset \mathbf{R}^{KM}$, and for all $y \in D$

$$0 \leq \phi_{kp}(y) \leq 1, \quad 1 \leq p \leq M, \quad 1 \leq k \leq K \quad \& \quad \sum_{p=1}^M \phi_{kp}(y) \leq 1, \\ 1 \leq k \leq K. \quad (24)$$

We assume that the conditional probability for location k at time t to be in state $p \in \{1, \dots, M\}$ (i.e., to have $\omega_{tk} = p$) given ω^{t-1} is

$$\phi_{kp}(\eta^T(\omega_{t-d}^{t-1})\beta)$$

for some vector of parameters $\beta \in \mathbf{R}^\kappa$ and known to us function $\eta(\cdot)$ taking values in the space of $\kappa \times KM$ matrices and such that $\eta^T(\omega_{d-1}^{t-1})\beta \in D$ whenever $\omega_{\tau k} \in \{0, 1, \dots, M\}$ for all τ and k . As a result, the conditional probability to have $\omega_{tk} = 0$ is

$$1 - \sum_{p=1}^M \phi_{kp}(\eta^T(\omega_{t-d}^{t-1})\beta).$$

In addition, we assume that we are given a convex compact set $\mathcal{X} \subset \mathbf{R}^\kappa$ such that $\beta \in \mathcal{X}$ and for all such β

$$\eta^T(\omega_{t-d}^{t-1})\beta \in D \quad \forall \{\omega_{\tau k} \in \{0, 1, \dots, M\} \quad \forall \tau, k\}.$$

Same as in Section II-E, we encode the collection $\{\omega_{tk} : 1 \leq k \leq K\}$ of locations' states at time t by block vector $\bar{\omega}_t$ with K blocks of dimension M each, with the k -th block equal to the ω_{tk} -th vector of the standard basis in \mathbf{R}^M when $\omega_{tk} > 0$ and equal to 0 when $\omega_{tk} = 0$. We clearly have

$$\mathbf{E}_{|\omega^{t-1}}\{\bar{\omega}_t\} = \phi(\eta^T(\omega_{t-d}^{t-1})\beta).$$

Setting

$$\begin{aligned} F(x) &= \frac{1}{N} \mathbf{E}_{\omega^N} \left\{ \sum_{t=1}^N \left[\eta(\omega_{t-d}^{t-1}) \phi(\eta^T(\omega_{t-d}^{t-1})x) \right. \right. \\ &\quad \left. \left. - \eta(\omega_{t-d}^{t-1}) \bar{\omega}_t \right] \right\} : \mathcal{X} \rightarrow \mathbf{R}^\kappa, \end{aligned}$$

$$F_{\omega^N}(x) = \underbrace{\frac{1}{N} \sum_{t=1}^N \eta(\omega_{t-d}^{t-1}) \phi(\eta^T(\omega_{t-d}^{t-1})x)}_{A_{\omega^N}(x)} : \mathcal{X} \rightarrow \mathbf{R}^\kappa$$

$$\begin{aligned}
& - \underbrace{\frac{1}{N} \sum_{t=1}^N \eta(\omega_{t-d}^{t-1}) \bar{\omega}_t}_{a[\omega^N]} : \mathcal{X} \rightarrow \mathbf{R}^\kappa \\
\bar{F}_{\omega^N}(x) &= A_{\omega^N}(x) \\
& - \underbrace{\frac{1}{N} \sum_{t=1}^N \eta(\omega_{t-d}^{t-1}) \phi(\eta^T(\omega_{t-d}^{t-1}) \beta)}_{a[\omega^N]} : \mathcal{X} \rightarrow \mathbf{R}^\kappa,
\end{aligned} \tag{25}$$

(see equation (22)), we can repeat word by word the comment at the end of Section II-F1.

III. MAXIMUM LIKELIHOOD ESTIMATE

In the previous sections, we have discussed the Least Squares estimate of the parameter vector β . Now, we consider commonly used in statistics alternative approach based on the Maximum Likelihood (ML) estimation. ML estimate is obtained by maximizing over $\beta \in \mathcal{X}$ the conditional likelihood of what we have observed, the condition being the actually observed values of ω_{tk} for $-d+1 \leq t \leq 0$ and $1 \leq k \leq K$. In this section, we study the properties of the ML estimate and show that its calculation reduces to a convex optimization problem.

A. ML Estimation: Case of Linear Link Function

a) *Single state model:* Assume, in addition to what has been already assumed, that for every t random variables ω_{tk} are conditionally independent across k given ω^{t-1} . Then the negative log-likelihood, conditioned by the value of ω^0 , is given by

$$\begin{aligned}
L(\beta) = & \frac{1}{N} \sum_{t=1}^N \sum_{k=1}^K \left[-\omega_{tk} \ln \left(\beta_k + \sum_{s=1}^d \sum_{\ell=1}^K \beta_{k\ell}^s \omega_{(t-s)\ell} \right) \right. \\
& \left. - (1 - \omega_{tk}) \ln \left(1 - \beta_k - \sum_{s=1}^d \sum_{\ell=1}^K \beta_{k\ell}^s \omega_{(t-s)\ell} \right) \right].
\end{aligned}$$

Note that $L(\cdot)$ is a convex function, so the ML estimate in our model reduces to the convex program

$$\min_{x \in \mathcal{X}} L(x). \tag{26}$$

b) *Multi-state model:* Assume that states ω_{tk} at locations k at time t are conditionally independent across $k \leq K$ given ω^{t-1} . Then the ML estimate is given by minimizing, over $\beta \in \mathcal{X}$, the conditional negative log-likelihood of collection ω^N of observations (the condition being the initial segment ω^0 of the observation). The objective in this minimization problem is the convex function

$$L_{\omega^N}(\beta) = -\frac{1}{N} \sum_{t=1}^N \sum_{k=1}^K \psi_{tk}(\beta, \omega^N),$$

where

$$\psi_{tk}(\beta, \omega^N) = \begin{cases} \ln \left(\left[\eta_k^T(\omega_{t-d}^{t-1}) \beta \right]_{\omega_{tk}} \right), & \omega_{tk} \in \{1, \dots, M\}, \\ \ln \left(1 - \sum_{j=1}^M \left[\eta_k^T(\omega_{t-d}^{t-1}) \beta \right]_j \right), & \omega_{tk} = 0. \end{cases} \tag{27}$$

c) *Toward performance guarantees:* We are about to show that the ML estimate has a structure similar to the LS estimator that we have dealt within Section II, and obeys bounds similar to (23). Given a small positive tolerance ϱ , consider M -state spatio-temporal process with K locations and vector of parameters $\beta \in \mathbf{R}^\kappa$, as defined in Section II-E, restricted to reside in the polyhedral set B_ϱ cut off \mathbf{R}^κ by “ ϱ -strengthened” version of constraints (18), specifically, the constraints

$$\begin{aligned}
\varrho &\leq \beta_k(p) + \sum_{s=1}^d \sum_{\ell=1}^K \min_{0 \leq q \leq M} \beta_{k\ell}^s(p, q), \\
1 \leq p &\leq M, 1 \leq k \leq K, \\
1 - \varrho &\geq \sum_{p=1}^{M-1} \beta_k(p) + \sum_{s=1}^d \sum_{\ell=1}^K \max_{0 \leq q \leq M} \sum_{p=1}^M \beta_{k\ell}^s(p, q), \\
1 \leq k &\leq K.
\end{aligned} \tag{28}$$

The purpose of strengthening the constraints on β is to make the maximum likelihood, to be defined below, continuously differentiable on the given parameter domain.

In what follows, we treat vectors from \mathbf{R}^{KM} as block vectors with K blocks of dimension M each. For such a vector z , $[z]_{kp}$ stands for the p -th entry in the k -th block of z . Let

$$Z_0 = \left\{ \omega \in \mathbf{R}^{MK} : \omega \geq 0, \sum_{p=1}^M [\omega]_{kp} \leq 1 \quad \forall k \leq K \right\}.$$

Similarly, for a small positive tolerance ϱ , define

$$Z_\varrho = \left\{ z \in \mathbf{R}^{MK} : [z]_{kp} \geq \varrho \quad \forall k, p, \sum_{p=1}^M [z]_{kp} \leq 1 - \varrho \quad \forall k \right\} \subset Z_0.$$

We associate with a vector $w \in Z_0$ the convex function $\mathcal{L}_w : Z_\varrho \rightarrow \mathbf{R}$,

$$\begin{aligned}
\mathcal{L}_w(z) := & - \sum_{k=1}^K \left[\sum_{p=1}^M [w]_{kp} \ln([z]_{kp}) \right. \\
& \left. + \left[1 - \sum_{p=1}^M [w]_{kp} \right] \ln \left(1 - \sum_{p=1}^M [z]_{kp} \right) \right]. \tag{29}
\end{aligned}$$

From now on, assume that we are given a convex compact set $\mathcal{X} \subset B_\varrho$ known to contain the true vector β of parameters. Then the problem of minimizing the negative log-likelihood becomes

$$\min_{x \in \mathcal{X}} \left\{ L_{\omega^N}(x) = \frac{1}{N} \sum_{t=1}^N \mathcal{L}_{\bar{\omega}_t} \left(\eta^T(\omega_{t-d}^{t-1}) x \right) \right\}, \tag{30}$$

where $\bar{\omega}_t = \bar{\omega}_t(\omega^t)$ encodes, as explained in Section II-E, the observations at time t , and $\eta(\omega_{t-d}^{t-1})$ are as defined in Section II-E.

Note that by construction, $\bar{\omega}_t$ belongs to Z_0 . Moreover, by construction, we have $\eta^T(\omega_{t-d}^{t-1}) x \in Z_\varrho$ whenever $x \in B_\varrho$ and $\omega_{tk} \in \{0, 1, \dots, M\}$ for all t and k . Now, minimizers of $L_{\omega^N}(x)$ over $x \in \mathcal{X}$ are exactly the solutions of the variational inequality stemming from \mathcal{X} and the monotone and smooth vector

field (the smoothness property is due to $L_{\omega^N}(x)$ being convex and smooth on \mathcal{X}):

$$F_{\omega^N}(x) = \nabla_x L_{\omega^N}(x) = \frac{1}{N} \sum_{t=1}^N \eta\left(\omega_{t-d}^{t-1}\right) \theta\left(\eta^T\left(\omega_{t-d}^{t-1}\right) x, \bar{\omega}_t(\omega^t)\right)$$

with

$$\begin{aligned} \theta(z, \omega) &= \nabla_z \mathcal{L}_w(z) \\ &= - \sum_{k=1}^K \left[\sum_{p=1}^M \frac{[w]_{kp}}{[z]_{kp}} e^{kp} - \frac{1 - \sum_{p=1}^M [w]_{kp}}{1 - \sum_{p=1}^M [z]_{kp}} \sum_{p=1}^M e^{kp} \right], \\ &\quad [w \in Z_0] \end{aligned}$$

where $e^{kp} \in \mathbf{R}^{KM}$ is the block-vector with the p -th vector of the standard basis in \mathbf{R}^M as the k -th block and all other blocks equal to 0.

Note that we clearly have

$$w \in Z_\varrho \Rightarrow \phi_w(w) = 0. \quad (31)$$

Let us show that $F_{\omega^N}(\beta)$ is “typically small”: its magnitude obeys the large deviation bounds similar to (8) and (23). Indeed, let us set $\bar{z}_t(\omega^{t-1}) = \eta^T\left(\omega_{t-d}^{t-1}\right) \beta$, so that $\bar{z}_t \in Z_0$ due to $\beta \in B_\varrho$. Invoking (31) with $w = \bar{z}_t(\omega^{t-1})$, we have

$$F_{\omega^N}(\beta) = \frac{1}{N} \sum_{t=1}^N \underbrace{\eta\left(\omega_{t-d}^{t-1}\right) \vartheta_t[\omega^t]}_{\xi_t},$$

where

$$\begin{aligned} \vartheta_t[\omega^t] &= - \sum_{k=1}^K \left[\sum_{p=1}^M \frac{[\bar{\omega}_t(\omega^t)]_{kp} - [\bar{z}_t(\omega^{t-1})]_{kp}}{[\bar{z}_t(\omega^{t-1})]_{kp}} e^{kp} \right. \\ &\quad \left. + \frac{\sum_{p=1}^M [\bar{z}_t]_{kp} - [\bar{\omega}_t(\omega^t)]_{kp}}{1 - \sum_{p=1}^M [\bar{z}_t(\omega^{t-1})]_{kp}} \sum_{p=1}^M e^{kp} \right]. \end{aligned}$$

Since the conditional expectation of $[\bar{\omega}_t(\omega^t)]_{kp}$ given ω^{t-1} equals $[\bar{z}_t(\omega^{t-1})]_{kp}$ the conditional expectation of ξ_t given ω^{t-1} is zero. Besides this, random vectors ξ_t take their values in a bounded set (of size depending on ϱ). As a result, $\|F_{\omega^N}(\beta)\|_\infty$ admits bound on probabilities of large deviations of the form (23), with properly selected (and depending on ϱ) factor Θ . However, for the reasons presented in Section II-F, extracting from this bound meaningful conclusions on the accuracy of the ML estimate is a difficult task, and it remains an open problem.

Remark 2 (Decomposition of LS and ML estimation): In the models we have considered, the optimization problems (6), (20), (26), and (30), we aim to solve when building the LS and the ML estimates under mild assumptions are decomposable (in spite of the fact that the observations are dependent). Indeed, vector

$$\begin{aligned} \beta &= \{\beta_{kp}, \beta_{k\ell}^s(p, q), 1 \leq k, \ell \leq K, \\ &\quad 1 \leq p \leq M, 0 \leq q \leq M, 1 \leq s \leq d\} \end{aligned}$$

of the model parameters can be split into K subvectors

$$\begin{aligned} \beta^k &= \{\beta_{kp}, \beta_{k\ell}^s(p, q), 1 \leq \ell \leq K, 1 \leq p \leq M, \\ &\quad 0 \leq q \leq M, 1 \leq s \leq d\}, \quad k = 1, \dots, K. \end{aligned}$$

It is immediately seen that the objectives to be minimized in the problems in question are sums of K terms, with the k -th term depending only on x^k . As a result, if the domain \mathcal{X} summarizing our a priori information on β is decomposable: $\mathcal{X} = \{x : x^k \in \mathcal{X}_k, 1 \leq k \leq K\}$, the optimization problems yielding the LS and the ML estimates are collections of K uncoupled convex optimization problems in variables x^k . Moreover, under favorable circumstances optimization problem (20) admits even finer decomposition. Namely, splitting β^k into subvectors

$$\beta^{kp} = \{\beta_{kp}, \beta_{k\ell}^s(p, q), 1 \leq \ell \leq K, 1 \leq s \leq d, 0 \leq q \leq M\},$$

it is easily seen that the objective in (20) is the sum, over $k \leq K$ and $p \leq M$, of functions $\Psi_{\omega^N}^{kp}(x^{kp})$. As a result, when $\mathcal{X} = \{x : x^{kp} \in \mathcal{X}_{kp}, 1 \leq k \leq K, 1 \leq p \leq M\}$, (20) is a collection of KM uncoupled convex problems $\min_{x^{kp} \in \mathcal{X}_{kp}} \Psi_{\omega^N}^{kp}(x^{kp})$.

The outlined decompositions may be used to accelerate the solution process.

B. ML Estimate: General Link Functions

Let us now derive ML estimate for the case of nonlinear link function considered in Section II-F2. In this situation, we strengthen constraints (24) on D to

$$\begin{aligned} y \in D \Rightarrow \varrho &\leq \phi_{kp}(y), \quad \sum_{p=1}^M \phi_{kp}(y) \leq 1 - \varrho, \\ &1 \leq k \leq K, 1 \leq p \leq M, \end{aligned}$$

with some $\varrho > 0$. Assuming that ω_{ik} 's are conditionally independent across k given ω^{t-1} , computing ML estimate for the general link-function reduces to solving problem (30) with $\mathcal{L}_w(z) : D \rightarrow \mathbf{R}$, $w \in Z_0$, given by

$$\begin{aligned} \mathcal{L}_w(z) &= - \sum_{k=1}^K \left[\sum_{p=1}^M [w]_{kp} \ln(\phi_{kp}(z)) \right. \\ &\quad \left. + \left[1 - \sum_{p=1}^M [w]_{kp} \right] \ln\left(1 - \sum_{p=1}^M \phi_{kp}(z)\right) \right]. \end{aligned}$$

Assuming ϕ continuously differentiable on D and $\mathcal{L}_w(\cdot)$ convex on D , we can repeat, with straightforward modifications, everything that was said above (that is, in the special case of $\phi(z) \equiv z$), including exponential bounds on probabilities of large deviations of $F_{\omega^N}(\beta)$. However, in general, beyond the case of affine $\phi_{kp}(\cdot)$, function $\mathcal{L}_w(\cdot)$ becomes nonconvex. This is due to the fact that convexity on D of functions

$$-\ln(\phi_{kp}(\cdot)), \quad -\ln\left(1 - \sum_p \phi_{kp}(\cdot)\right)$$

is a rare commodity. Nevertheless, convexity of these functions does take place in the case logistic link function

$$\phi_{kp}(z) = \frac{\exp\{a_{kp}(z)\}}{\sum_{q=0}^M \exp\{a_{kq}(z)\}}$$

with functions $a_{kq}(z)$, $0 \leq q \leq M$ that are affine in z .

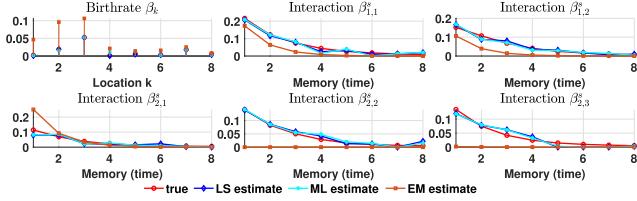


Fig. 3. Single-state process: estimates for baseline intensity β_k and interactions parameters $\beta_{k\ell}^s$ for one random instance.

TABLE I

SINGLE-STATE PROCESS: ERROR OF ML, LS, AND EM ESTIMATION FOR THE ONE INSTANCE SHOWN IN FIG. 3

Estimate	ℓ_1 error	ℓ_2 error	ℓ_∞ error
ML	1.7150 (22.57%)	0.1534 (17.67%)	0.0342 (13.64%)
LS	1.8849 (24.80%)	0.1714 (19.73%)	0.0372 (14.84%)
EM (exponential kernel)	6.3127 (83.06%)	0.6413 (73.83%)	0.2105 (83.97%)

IV. NUMERICAL EXPERIMENTS

A. Experiments With Simulated Data

This section presents the results of several simulation experiments illustrating applications of the proposed Bernoulli process models. We compare performances of Least Squares (LS) and Maximum Likelihood (ML) estimates in terms of ℓ_1 , ℓ_2 , and ℓ_∞ norms of the error of parameter vector recovery. We assume that d (or a reasonable upper bound on it) is known in our simulation examples. The bracket percentage inside the table below shows the norm of the error relative to the norm of the corresponding true parameter vector.

1) *Single State Spatio-Temporal Processes*: First, consider a single state setting with the memory depth $d = 8$ and the number of locations $K = 8$. The true parameter values are selected randomly from the set \mathcal{X}_0 as follows:

- $\beta_k \geq 0$, $\beta_{k\ell}^s \geq 0$; and $\beta_k + \sum_{s=1}^d \sum_{\ell=1}^K \beta_{k\ell}^s \leq 1 \quad \forall k$;
- $\beta_{k\ell}^s = 0$ when $|k - \ell| > 1$ (interactions are local);
- For every $1 \leq k, \ell \leq K$, $\beta_{k\ell}^s$ is a non-increasing convex function of s .⁸

Note that we have imposed additional to (2) constraints on β .

We report the performance of the LS estimate (obtained by solving $\text{VII}[F_{\omega^N}, \mathcal{X}]$) and the ML estimate (obtained by solving (26)). To ensure a fair comparison, we do not introduce any additional constraints on the interaction coefficients in our estimation procedure (meaning that the LS and ML estimates do not have any prior knowledge about \mathcal{X}_0 and their assumed admissible set \mathcal{X} is much larger than \mathcal{X}_0). Utilizing the MATLAB implementation [30] of the EM algorithm, we also compute estimations of parameters of the commonly used model of Hawkes process with exponential temporal kernel (see, e.g., [5]). The latter is equivalent to assuming that $\beta_{k\ell}^s = a_{k\ell} s^{-\tau_s}$, $s = 1, 2, \dots$, where $\tau > 0$ is the decay rate parameter and $a_{k\ell} > 0$ represents the interactions between two locations.

Fig. 3 shows the recovered interaction coefficients using various methods with $N = 10,000$ observations, for a single (randomly generated) instance. The associated error metrics are presented in Table I. The confidence intervals in Fig. 4 are

⁸Here, the convexity of a function $f(s)$ in $s \in G = \{1, \dots, d\}$ means that the function is the restriction of a convex function on the segment $[1, d]$ onto the grid G or, which is the same, that $f(s-1) - 2f(s) + f(s+1) \geq 0$, $s = 2, 3, \dots, d-1$. This translates into the constraint $\beta_{k,\ell}^{s-1} - 2\beta_{k,\ell}^s + \beta_{k,\ell}^{s+1} \geq 0$, $s = 2, 3, \dots, d-1 \quad \forall k, \ell$.

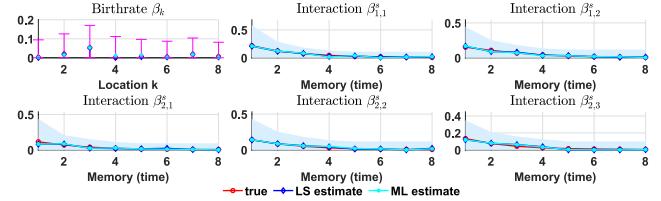


Fig. 4. Computed 90% confidence intervals corresponding to Fig. 3.

TABLE II
SINGLE-STATE PROCESS: ERROR OF ML, LS, AND EM ESTIMATION
AVERAGED OVER 100 TRIALS

Estimate	ℓ_1 error	ℓ_2 error	ℓ_∞ error
ML	1.1482 (15.11%)	0.1112 (12.60%)	0.0336 (11.87%)
LS	1.9776 (26.02%)	0.1831 (20.72%)	0.0472 (16.62%)
EM (exponential kernel)	6.4725 (85.16%)	0.6695 (75.72%)	0.2209 (75.17%)

computed according to (15) by letting e be standard basis vectors in \mathbf{R}^K and restricting the parameter space to \mathcal{X} . We also repeat the experiment 100 times (each time, generate new true parameters), and the average errors are reported in Table II. The experiments show that ML and LS estimates exhibit similar performance (ML outperforming slightly the LS estimates). Both of them outperform the recovery by EM algorithm based on the exponential kernel, which may be due to a more flexible parameterization of our model.

2) *Multi-State Spatio-Temporal Processes*: Now consider a multi-state spatio-temporal Bernoulli process with the number of states $M = 2$. Here the possible states $p = 0$ represents no event, $p = 1, 2$ represent the event of category 1 and 2, respectively. We assume memory depth $d = 8$ and the number of locations $K = 10$. The true parameters are randomly generated from the set \mathcal{X}_0 specified by (again, we impose additional constraints as in Section IV-A1):

- $\beta_k(p) \geq 0$, $\beta_{kl}^s(p, q) \geq 0$; $\sum_{p=1}^M \beta_k(p) + \sum_{s=1}^d \sum_{\ell=1}^K \max_{0 \leq q \leq M} \sum_{p=1}^M \beta_{kl}^s(p, q) \leq 1 \quad \forall k \leq K$;
- $\beta_{kl}^s(p, q) = 0$ when $|k - l| > 1 \quad \forall p, q$ (interactions are local);
- For every $1 \leq k, l \leq K$ and $1 \leq p \leq M$, $0 \leq q \leq M$, $\beta_{kl}^s(p, q)$ is a non-increasing convex function of s .

Furthermore, we consider two scenarios, with additional constraints on the parameters

- Scenario 1: events can only trigger future events of the same category, i.e., $\beta_{kl}^s(p, q) \equiv 0$, $q \neq p$;
- Scenario 2: events of category $q = 0, \dots, M$, only trigger events with category $p \leq q$. This can happen, for example, when modeling earthquakes aftershocks: events are marked using M categories according to their magnitudes: $u_1 < \dots < u_M$. Set $u_0 = 0$ and treat the event “no earthquake” as “earthquake of magnitude 0.” Then each earthquake can trigger “aftershocks” with the same or smaller magnitudes.

We generate a synthetic data sequence of length $N = 20,000$. For a single (randomly generated) instance, recovery of baseline and interaction parameters are presented in Fig. 5. The associated recovery errors of the LS estimate (solution to (20)) and the ML estimate (solution to (30)) are reported in Table III. In addition, we also report the recovery errors separately for (i) the baseline intensity vector (referred to as “birthrates”) $\beta_{\text{birth}} = \{\beta_k(p), k \leq K, 1 \leq p \leq M\} \in \mathbb{R}^{KM \times 1}$; and (ii) the vector of interactions between different locations

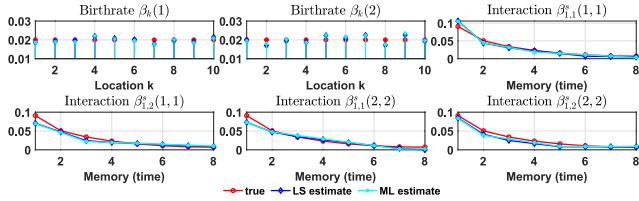


Fig. 5. Multi-state process: examples of LS and ML estimates for baseline intensity $\beta_k(p)$ and interactions parameters $\beta_{k\ell}^s(p, q)$.

TABLE III
MULTI-STATE PROCESS RECOVERY: NORMS OF RECOVERY ERROR FOR
LS ESTIMATE $\hat{\beta}_{LS}$ AND ML ESTIMATE $\hat{\beta}_{ML}$

Estimate	Scenario 1		Scenario 2	
	ℓ_1 error	ℓ_2 error	ℓ_1 error	ℓ_2 error
$\hat{\beta}_{ML}$	0.3524 (4.7%)	0.0532 (2.5%)	1.0179 (13.6%)	0.1146 (5.9%)
$\hat{\beta}_{LS}$	0.4947 (6.6%)	0.0744 (3.4%)	1.0854 (14.5%)	0.1230 (6.3%)
$\hat{\beta}_{ML, \text{birth}}$	0.0106 (2.7%)	0.0028 (3.1%)	0.0226 (5.7%)	0.0060 (6.7%)
$\hat{\beta}_{LS, \text{birth}}$	0.0160 (4.0%)	0.0044 (5.0%)	0.0237 (5.9%)	0.0066 (7.4%)
$\hat{\beta}_{ML, \text{inter}}$	0.3419 (4.8%)	0.0531 (2.5%)	0.9952 (14.0%)	0.1144 (5.9%)
$\hat{\beta}_{LS, \text{inter}}$	0.4786 (6.7%)	0.0743 (3.4%)	1.0617 (15.0%)	0.1228 (6.3%)

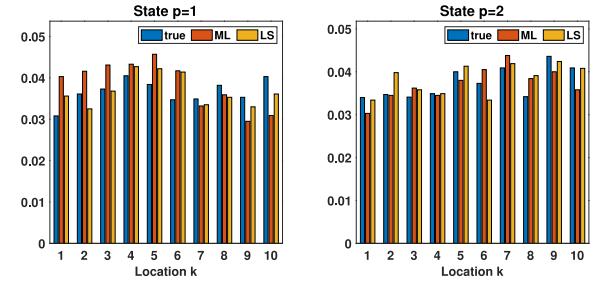
$\beta_{\text{inter}} = \{\beta_{k\ell}^s(p, q)\} \in \mathbb{R}^{dk^2M(M+1) \times 1}$. As shown in Table III, the ℓ_1 recovery error for estimating birthrate is smaller than that for the interaction parameters. Thus, the recovery error for β is dominated by the error for interaction parameters. This could be explained because the magnitude of the baseline intensity is higher than the influence parameters (which is usually needed to have stationary processes).

Finally, to assess the predictive capability of our model, we did the following experiment. Generate one sequence of discrete events, with length $N = 20,000$, using randomly selected parameters. We divide the sequence in half: use half for “training” and the other half for “testing”. In particular, we (1) use the first half of the sequence for estimating the Bernoulli process model parameter, (2) use the “trained” model to generate a new “synthetic” sequence of length $N/2$, and (3) compare the “synthetic” sequence with the “test” sequence, in terms of the frequency of events, for each category, at each location. The results in Fig. 6 show that the synthetic sequence has a reasonably good match with the testing sequence, based on the LS and the ML estimates.

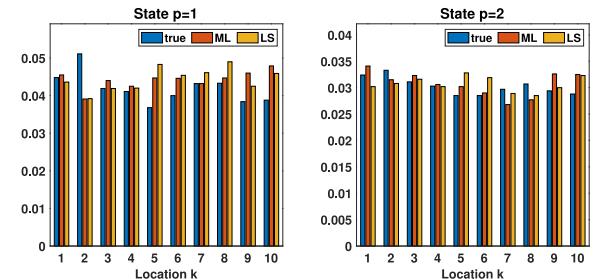
3) *Sparse Network Recovery With Negative and Non-Monotone Interactions*: In the last synthetic example, we consider an example to recover a network with “non-conventional” interactions: non-monotonic temporal interactions and negative interactions. Consider a sparse, directed, and non-planar graph (meaning that this cannot be embedded on a two-dimensional Euclidean space and, thus, this does not correspond to discretized space) with $K = 8$ nodes. The interaction functions are illustrated in Fig. 7.

The baseline intensities are all positive at all 8 nodes. The directed edge (arrows) means there is a one-directional “influence” from one node to its neighbor, e.g., $1 \rightarrow 5$. The self-edges, e.g., $2 \rightarrow 2$ and $5 \rightarrow 5$, denote that these nodes have a self-exciting effect: events happen at the node will trigger future events at itself. The true parameters of the model are generated as follows.

- Baseline parameters values at all locations are drawn independently from a uniform distribution on $[0, 0.2]$;



Scenario 1



Scenario 2

Fig. 6. Multi-state process: experiment to compare the frequency of events from a synthetic sequence (generated using models estimated from training sequence using LS and ML estimates) with that from the testing sequence.

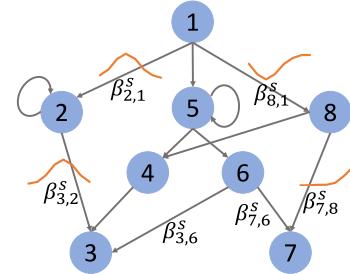


Fig. 7. Sparse non-planar graph with non-monotonic and negative interaction. Note that the interaction $1 \rightarrow 8$ is negative.

- For each directed edge $\ell \rightarrow k$, the interaction $\beta_{k,\ell}^s$ is given by $\beta_{k,\ell}^s = 0.05e^{-0.25(s-\tau_{k,\ell})^2}$, $s \geq 0$, and the peak $\tau_{k,\ell}$ is randomly chosen from $\{1, \dots, d\}$, except for one edge $1 \rightarrow 8$, whose interaction function is set to be negative: $\beta_{8,1}^s = -0.05e^{-0.25(s-\tau_{8,1})^2}$.

In our implementation, we consider two scenarios: (1) the graph structure is *unknown*: we do *not* impose sparsity constraints while obtaining the LS and ML estimates; (2) the graph structure is *known*, and then we impose the sparsity constraints by setting the interactions to be 0 when there is no edge; this illustrate the scenario when we have some prior information about the network structure. We report recovery errors for the two scenarios in Table IV and compare the recovery of interaction parameters under scenario (1) with the true values in Fig. 8. From the experiment results, we observe that both the LS and ML estimates match closely with the true parameters, even when the underlying graph structure is unknown. The comparison in Table IV shows a significant improvement in the estimation error when the graph structure is known *a priori*. This is consistent with our previous remark that knowing the network structure allows for a better choice of the feasible region resulting in reduced estimation error. Moreover, by

TABLE IV
SPARSE NETWORK RECOVERY WITH NON-CONVENTIONAL
INTERACTIONS: ERRORS OF LS AND ML ESTIMATES $\hat{\beta}_{\text{LS}}$, $\hat{\beta}_{\text{ML}}$

Estimate	Unknown Graph			Known Graph		
	ℓ_1 error	ℓ_2 error	ℓ_∞ error	ℓ_1 error	ℓ_2 error	ℓ_∞ error
$\hat{\beta}_{\text{ML}}$	1.7694 (58.71%)	0.1128 (24.65%)	0.0224 (13.79%)	0.4715 (15.64%)	0.0593 (12.95%)	0.0173 (10.68%)
$\hat{\beta}_{\text{LS}}$	1.8757 (62.23%)	0.1166 (25.48%)	0.0211 (13.01%)	0.4773 (15.84%)	0.0606 (13.23%)	0.0204 (12.58%)
$\hat{\beta}_{\text{ML, birth}}$	0.0367 (3.84%)	0.0162 (4.42%)	0.0111 (6.84%)	0.0126 (1.32%)	0.0068 (1.85%)	0.0061 (3.75%)
$\hat{\beta}_{\text{LS, birth}}$	0.0378 (3.95%)	0.0172 (4.69%)	0.0129 (7.94%)	0.0126 (1.32%)	0.0069 (1.89%)	0.0061 (3.75%)
$\hat{\beta}_{\text{ML, inter}}$	1.7327 (84.20%)	0.1117 (40.69%)	0.0224 (44.73%)	0.4589 (22.30%)	0.0589 (21.46%)	0.0173 (34.65%)
$\hat{\beta}_{\text{LS, inter}}$	1.8379 (89.31%)	0.1153 (42.02%)	0.0211 (42.19%)	0.4648 (22.58%)	0.0602 (21.92%)	0.0204 (40.81%)

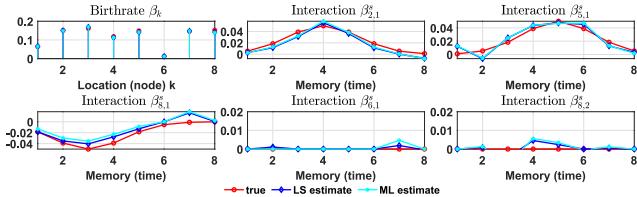


Fig. 8. Sparse network identification when graph is unknown: examples of LS and ML estimates of baseline intensity and vectors of interaction parameters; interactions $\beta_{6,1}$ and $\beta_{8,2}$ correspond to edges $1 \rightarrow 6$ and $2 \rightarrow 8$ which do not exist in the graph in Fig. 7.

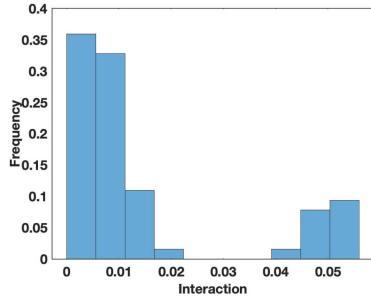


Fig. 9. Sparse network support recovery: histogram of the recovered interaction parameters $\{\max_{s=1}^d |\beta_{k,\ell}^s|, 1 \leq k, \ell \leq K\}$. Edges with non-zero interactions can be perfectly separated from those with zero interactions.

examining the histogram of the maximum interaction between each pair, i.e., $\{\max_{s=1}^d |\beta_{k,\ell}^s|, 1 \leq k, \ell \leq K\}$ as shown in Fig. 9, we observe that we can indeed accurately recover the support of the graph: the estimates of the edges with non-zero interactions, are completely separable from the estimates of the edges with zero interactions. This indicates that we can apply an appropriate threshold (in this case, e.g., 0.03) to recover precisely the unknown graph structure completely. This example also shows that even when prior information about the sparse structure of the underlying network is not available, LS and ML estimates can recover the underlying network reasonably well, which opens possibilities of applying the proposed approach to perform causal inference [31] using discrete-event data.

B. Real Data Studies: Crime in Atlanta

Finally, we study a real crime dataset in Atlanta, USA, to demonstrate the promise of our methods to recover interesting structures from real-data. We consider two categories of crime incidents, “burglary” and “robbery”. These incidents were reported to the Atlanta Police Department from January 1, 2015, to September 19, 2017. The dataset contains 47,245 “burglary” and 3,739 “robbery” incidents. As mentioned in the introduction, it is believed that crime incidents are related and have “self-exciting” patterns: once crime incidence happens, it triggers similar crimes more likely to happen in the

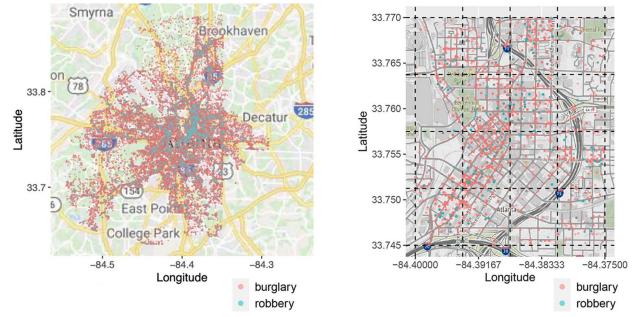


Fig. 10. Raw data map: burglary and robbery incidents in Atlanta. Left: the full map; Right: zoom-in around downtown Atlanta.

TABLE V
CRIME EVENT MODEL RECOVERY: FREQUENCY OF BURGLARY AND ROBBERY EVENTS AT EACH LOCATION

Locations	Burglary			Robbery		
	True	With constr	Without constr	True	With constr	Without constr
1	0.1499	0.1707	0.1766	0.0102	0.0195	0.0186
2	0.0284	0.0373	0.0445	0.0017	0.0203	0.0212
3	0.0483	0.0580	0.0606	0.0021	0.0254	0.0195
4	0.0407	0.0364	0.0356	0.0017	0.0178	0.0224
5	0.0508	0.0529	0.0648	0.0042	0.0220	0.0165
6	0.1957	0.2088	0.1834	0.0131	0.0208	0.0144
7	0.0970	0.1368	0.1224	0.0068	0.0229	0.0191
8	0.0419	0.0580	0.0563	0.0021	0.0127	0.0182
9	0.0148	0.0161	0.0220	0.0013	0.0165	0.0212
10	0.0584	0.0729	0.0805	0.0055	0.0258	0.0178
11	0.1266	0.1525	0.1529	0.0106	0.0195	0.0169
12	0.1364	0.1266	0.1186	0.0102	0.0191	0.0169
13	0.0322	0.0521	0.0445	0.0021	0.0229	0.0224
14	0.0627	0.0868	0.0834	0.0055	0.0212	0.0195
15	0.0208	0.0224	0.0280	0.0008	0.0241	0.0216
16	0.0144	0.0203	0.0178	0.0013	0.0203	0.0212

neighborhood in the near future [32]. Here, we model the data using a multi-state Bernoulli process with two states: no event ($p = 0$), burglary ($p = 1$), and robbery ($p = 2$).

We extract crime events around the Atlanta downtown area, as shown in Fig. 10, which contains 6031 “burglary” events and 454 “robbery” events. The whole time horizon (788 days) is split into discrete time intervals of four hours. The memory depth d is set to 6 in this example. This value was obtained using a simple “cross-validation-like” procedure utilizing predictions of frequencies of the burglary and robbery incidents in various spatial cells. The downtown region is divided uniformly into 16 sub-regions.

We compute the LS estimates of the parameters $\{\beta_k(p), \beta_{k,l}^s(p, q)\}$, in two different ways to set up the constraints: in the first setup, we do not impose additional constraints on the parameters apart from “basic” constraints (18); in the second setup, we impose constraints to only consider temporal interaction function, $\beta_{k,l}^s$, with monotonic and convex “shapes”.⁹ The estimated parameters are shown in Fig. 11. In the figure, the size of the red dot in each region is proportional to the magnitude of the estimated birthrate $\beta_k(p)$, $k = 1, \dots, K$, for Burglary/Robbery, respectively; the width of the arrow is proportional to the magnitude of the interaction $\beta_{k,l}^s(p, q)$ between

⁹Such constraints are routinely imposed when estimating parameters of Hawkes model, see, e.g., [12].

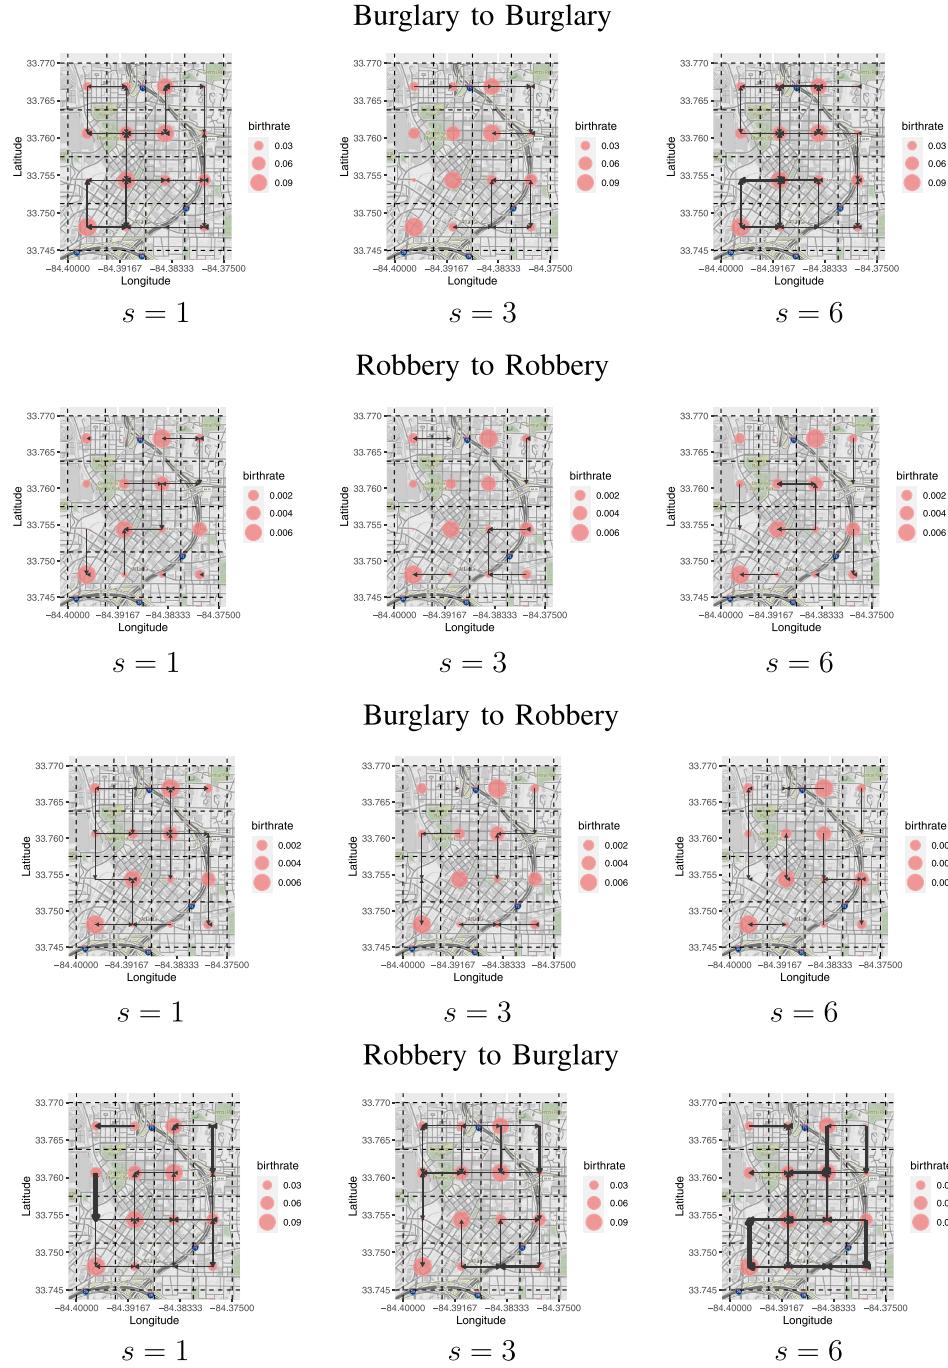


Fig. 11. Robbery and burglary in downtown Atlanta: recovered spatio-temporal interactions, using LS estimates without additional constraint on the shapes of the interaction functions.

locations. It is interesting to notice that our model recovers the dynamic of the interactions and how they change over time. There also seem to be strong interactions between burglary and robbery at different locations.

To validate the model, we experiment similar to we did for the simulated data in Section IV-A2. We take the two-year duration of data, divide the sequence in half, use the first half of the sequence to estimate a multi-state Bernoulli process model, generate a synthetic sequence, and compare with the second half of the sequence reserved for testing. We compare the frequency of Burglary and Robbery events across all locations, for the synthetic and testing sequence. The results are shown in Table V. The results look to be a reasonably good

match, considering that the crime events are relatively rare and with highly complex (and unknown) dynamics: predicting their frequency in the first place is a highly challenging task and an essential research task of criminology.

We also note that the prediction for burglary seems to be better since the frequencies from the synthetic sequence are very close, and the relative error is smaller. This is expected since the number of burglary cases is much larger than the number of robbery cases in our dataset, and the frequency of robbery cases is very small (typically below 0.01, as shown in Table V). The experiment serves as a sanity check and shows that for challenging and noisy real-world datasets, there could be a certain truth to the proposed methods.

REFERENCES

[1] 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. Stat. Assoc.*, vol. 106, no. 493, pp. 100–108, 2011.

[2] A. E. Johnson *et al.*, "MIMIC-III, a freely accessible critical care database," *Sci. Data*, vol. 3, no. 1, pp. 1–9, 2016.

[3] A. Stomakhin, M. B. Short, and A. L. Bertozzi, "Reconstruction of missing data in social networks based on temporal patterns of interactions," *Inverse Probl.*, vol. 27, no. 11, 2011, Art. no. 115013.

[4] E. L. Lai *et al.*, "Topic time series analysis of microblogs," *IMA J. Appl. Math.*, vol. 81, no. 3, pp. 409–431, 2016.

[5] K. Zhou, H. Zha, and L. Song, "Learning social infectivity in sparse low-rank networks using multi-dimensional Hawkes processes," in *Proc. Artif. Intell. Stat.*, 2013, pp. 641–649.

[6] S. Li, Y. Xie, M. Farajtabar, A. Verma, and L. Song, "Detecting changes in dynamic events over networks," *IEEE Trans. Signal Inf. Process. Netw.*, vol. 3, no. 2, pp. 346–359, Jun. 2017.

[7] P. Embrechts, T. Liniger, and L. Lin, "Multivariate Hawkes processes: An application to financial data," *J. Appl. Probab.*, vol. 48, pp. 367–378, Jul. 2016.

[8] M. Kuperman and G. Abramson, "Small world effect in an epidemiological model," *Phys. Rev. Lett.*, vol. 86, no. 13, p. 2909, 2001.

[9] A. G. Hawkes, "Spectra of some self-exciting and mutually exciting point processes," *Biometrika*, vol. 58, no. 1, pp. 83–90, 1971.

[10] A. G. Hawkes, "Point spectra of some mutually exciting point processes," *J. Roy. Stat. Soc. B (Methodol.)*, vol. 33, no. 3, pp. 438–443, 1971.

[11] A. G. Hawkes and D. Oakes, "A cluster process representation of a self-exciting process," *J. Appl. Probab.*, vol. 11, no. 3, pp. 493–503, 1974.

[12] A. Reinhart, "A review of self-exciting spatio-temporal point processes and their applications," 2017, *arXiv:1708.02647*.

[13] E. C. Hall and R. M. Willett, "Tracking dynamic point processes on networks," *IEEE Trans. Inf. Theory*, vol. 62, no. 7, pp. 4327–4346, Jul. 2016.

[14] B. Yuan, H. Li, A. L. Bertozzi, P. J. Brantingham, and M. A. Porter, "Multivariate spatiotemporal Hawkes processes and network reconstruction," *SIAM J. Math. Data Sci.*, vol. 1, no. 2, pp. 356–382, 2019.

[15] M. Eichler, R. Dahlhaus, and J. Dueck, "Graphical modeling for multivariate Hawkes processes with nonparametric link functions," *J. Time Series Anal.*, vol. 38, no. 2, pp. 225–242, 2017.

[16] S. Chen, A. Shojaie, E. Shea-Brown, and D. Witten, "The multivariate Hawkes process in high dimensions: Beyond mutual excitation," 2017, *arXiv:1707.04928*.

[17] M. Gomez-Rodriguez, J. Leskovec, and B. Schölkopf, "Modeling information propagation with survival theory," in *Proc. Int. Conf. Mach. Learn.*, 2013, pp. 666–674.

[18] 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.*, vol. 10, no. 3, pp. 1725–1756, 2016.

[19] J. Moller and R. P. Waagepetersen, *Statistical Inference and Simulation for Spatial Point Processes*. New York, NY, USA: CRC Press, 2003.

[20] G. Mohler *et al.*, "Modeling and estimation of multi-source clustering in crime and security data," *Ann. Appl. Stat.*, vol. 7, no. 3, pp. 1525–1539, 2013.

[21] S. Ertekin *et al.*, "Reactive point processes: A new approach to predicting power failures in underground electrical systems," *Ann. Appl. Stat.*, vol. 9, no. 1, pp. 122–144, 2015.

[22] J. Pitkin, I. Manolopoulou, and G. Ross, "Bayesian hierarchical modelling of sparse count processes in retail analytics," 2018, *arXiv:1805.05657*.

[23] A. Python, J. Illian, C. Jones-Todd, and M. Blangiardo, "A Bayesian approach to modelling fine-scale spatial dynamics of non-state terrorism: World study, 2002–2013," 2016, *arXiv:1610.01215*.

[24] N. R. Hansen, P. Reynaud-Bouret, and V. Rivoirard, "Lasso and probabilistic inequalities for multivariate point processes," *Bernoulli*, vol. 21, no. 1, pp. 83–143, 2015.

[25] A. B. Juditsky and A. Nemirovski, "Signal recovery by stochastic optimization," *Autom. Remote Control*, vol. 80, no. 10, pp. 1878–1893, 2019.

[26] V. Konev and S. Pergamenshchikov, "On asymptotic minimaxity of fixed accuracy estimators for autoregression parameters I. Stable process," *Math. Methods Stat.*, vol. 5, no. 2, pp. 125–153, 1996.

[27] R. Liptser and V. Spokoiny, "Deviation probability bound for martingales with applications to statistical estimation," *Stat. Probab. Lett.*, vol. 46, no. 4, pp. 347–357, 2000.

[28] K. Azuma, "Weighted sums of certain dependent random variables," *Tohoku Math. J.*, vol. 19, no. 3, pp. 357–367, 1967.

[29] Y. Nesterov, "Semidefinite relaxation and nonconvex quadratic optimization," *Optim. Methods Softw.*, vol. 9, nos. 1–3, pp. 141–160, 1998.

[30] H. Xu and H. Zha, "THAP: A MATLAB toolkit for learning with Hawkes processes," 2017, *arXiv:1708.09252*.

[31] A. Tank, E. B. Fox, and A. Shojaie, "Granger causality networks for categorical time series," 2017, *arXiv:1706.02781*.

[32] M. B. Short *et al.*, "A statistical model of criminal behavior," *Math. Models Methods Appl. Sci.*, vol. 18, no. s1, pp. 1249–1267, 2008.

[33] D. A. Freedman, "On tail probabilities for martingales," *Ann. Probab.*, vol. 3, no. 1, pp. 100–118, 1975.

[34] X. Fan, I. Grama, and Q. Liu, "Hoeffding's inequality for supermartingales," *Stochastic Process. Appl.*, vol. 122, no. 10, pp. 3545–3559, 2012.

[35] B. Bercu, B. Delyon, and E. Rio, *Concentration Inequalities for Sums and Martingales*. New York, NY, USA: Springer, 2015.



Anatoli Juditsky received the M.S.E.E. degree from the Moscow Institute of Physics and Technology in 1985, and the Ph.D. degree in electrical engineering from the Institute of Control Science, Moscow, USSR, in 1989. Since 1999, he has been a Professor with the Department of Mathematics and Informatics, University Joseph Fourier, Grenoble, France. He held a Research Position with INRIA-IRISA, Rennes, from 1991 to 1996, and then with INRIA, Grenoble, from 1996 to 1999. His current research interests include stochastic and large-scale convex optimization, statistical theory and their applications.

Arkadi Nemirovski, photograph and biography not available at the time of publication.



Liyan Xie (Student Member, IEEE) received the B.Sc. degree in statistics from the University of Science and Technology, China, in 2016. She is currently pursuing the Ph.D. degree with the H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology. She is also a Recipient of the 2020 ARC Student Fellowship and the IDEAS-TRIAD Research Scholarships from the Transdisciplinary Research Institute for Advancing Data Science (TRIAD) at Georgia Tech. Her research interests are the intersection between statistics and optimization, with a primary focus on sequential change detection and their applications in sensor networks and healthcare. She was awarded the finalist of the INFORMS QSR Student Paper Competition and Runner-up for the INFORMS Computing Society Student Paper Prize in 2019.



Yao Xie (Member, IEEE) received the B.Sc. degree in electrical engineering and computer science from the University of Science and Technology of China, the M.Sc. degree in electrical and computer engineering from the University of Florida, and the Ph.D. degree in electrical engineering (minor in Mathematics) from Stanford University. She is an Associate Professor and a Harold R. and Mary Anne Nash Early Career Professor with the Georgia Institute of Technology, H. Milton Stewart School of Industrial and Systems Engineering, and an Associate Director of the Machine Learning Center. She was a Research Scientist with Duke University. She has worked on such problems in sensor networks, social networks, power systems, crime data analysis, and wireless communications. Her research areas are statistics (in particular sequential analysis and sequential change-point detection), machine learning, and signal processing, providing the theoretical foundation and developing computationally efficient and statistically powerful algorithms. She received the National Science Foundation CAREER Award in 2017. She is currently an Associate Editor for *IEEE TRANSACTIONS ON SIGNAL PROCESSING*, *Sequential Analysis: Design Methods and Applications*, and *INFORMS Journal on Data Science*, and serves on the Editorial Board of *Journal of Machine Learning Research*.