
Learning Hawkes Processes from a Handful of Events

Farnood Salehi*
EPFL
farnood.salehi@epfl.ch

William Trouleau*
EPFL
william.trouleau@epfl.ch

Matthias Grossglauser
EPFL
matthias.grossglauser@epfl.ch

Patrick Thiran
EPFL
patrick.thiran@epfl.ch

Abstract

Learning the causal-interaction network of multivariate Hawkes processes is a useful task in many applications. Maximum-likelihood estimation is the most common approach to solve the problem in the presence of long observation sequences. However, when only short sequences are available, the lack of data amplifies the risk of overfitting and regularization becomes critical. Due to the challenges of hyper-parameter tuning, state-of-the-art methods only parameterize regularizers by a single shared hyper-parameter, hence limiting the power of representation of the model. To solve both issues, we develop in this work an efficient algorithm based on variational expectation-maximization. Our approach is able to optimize over an extended set of hyper-parameters. It is also able to take into account the uncertainty in the model parameters by learning a posterior distribution over them. Experimental results on both synthetic and real datasets show that our approach significantly outperforms state-of-the-art methods under short observation sequences.

1 Introduction

In many real-world applications, including finance, computational biology, social-network studies, criminology, and epidemiology, it is important to gain insight from the interactions of multivariate time series of discrete events. For example, in finance, changes in the price of a stock might affect the market [4]; and in epidemiology, individuals infected by an infectious disease might spread the disease to their neighbors [15]. Such networks of time series often exhibit mutually exciting patterns of diffusion. Hence, a recurring issue is to learn in an unsupervised way the causal structure of interacting networks. This task is usually tackled by defining a so-called causal graph of entities where an edge from a node i to a node j means that events in node j depend on the history of node i . Such causal interactions are typically learned with either directed information [24, 23], transfer entropy [26], or Granger causality [2, 11].

A widely used model for capturing mutually exciting patterns in a multi-dimensional time series is the Multivariate Hawkes process (MHP), a particular type of temporal point process where an event in one dimension can affect future arrivals in other dimensions. It has been shown that learning the excitation matrix of an MHP encodes the causal structure between the processes, both in terms of Granger causality [12] and directed information [13]. Most studies focus on the scalability of MHPs to large datasets. However, in many applications, data can be very expensive to collect, or simply not available. For example, in economic and public health studies, collecting survey data is usually an expensive process. Similarly, in the case of epidemic modeling, it is critical to learn as fast as possible the patterns of diffusion of a spreading disease. As a result, the amount of data

*The first two authors contributed equally to this work.

available is intrinsically limited. MHPs are known to be sensitive to the amount of data used for training, and the excitation patterns learned by MHPs from short sequences can be unreliable [29]. In such settings, the likelihood becomes noisy and regularization becomes critical. Nevertheless, as most hyper-parameter tuning algorithms such as grid search, random search, and even Bayesian optimization become challenging when the number of hyper-parameters is large, state-of-the-art methods only parameterize regularizers by a single shared hyper-parameter, hence limiting the power of representation of the model.

In this work, we address both the small data and hyper-parameter tuning issues by considering the parameters of the model as latent variables and by developing an efficient algorithm based on variational expectation-maximization. By estimating the evidence – rather than the likelihood – the proposed approach is able to optimize, with minimal computational complexity, over an extended set of hyper-parameters. Our approach is also able to take into account the uncertainty in the model parameters by fitting a posterior distribution over them. Therefore, rather than just providing a point estimate, this approach can provide an estimation of uncertainty on the learned causal graph. Experimental results on synthetic and real datasets show that, as a result, the proposed approach significantly outperforms state-of-the-art methods under short observation sequences, and maintains the same performance in the large-data regime.

2 Related Works

The most common approaches to learning the excitation matrix of MHPs are based on variants of regularized maximum-likelihood estimation (MLE). Zhou et al. [31] propose regularizers that enforce sparse and low-rank structures, along with an efficient algorithm based on the alternating-direction method of multipliers. To mitigate the parametric assumption, Xu et al. [28] represent the excitation functions as a series of basis functions, and to achieve sparsity under this representation they propose a sparse group-lasso regularizer. Such estimation methods are often referred to as non-parametric as they enable more flexibility on the shape of the excitation functions [19, 16]. To estimate the excitation matrix without any parametric modeling, fully non-parametric approaches were developed [32, 1]. However, these methods focus on scalability and target settings where large-scale datasets are available.

Bayesian methods go beyond the classic approach of MLE by enabling a probabilistic interpretation of the model parameters. A few studies tackled the problem of learning MHPs from a Bayesian perspective. Linderman and Adams [20] use a Gibbs sampling-based approach, but the convergence of the proposed algorithm is slow. To tackle this problem, Linderman and Adams [21] discretize the time, which introduces noise in the model. In a different setting where some of the events or dimensions are hidden, Linderman et al. [22] use an expectation maximization algorithm to marginalize over the unseen part of the network.

Bayesian probabilistic models are usually intractable and require approximate inference. To address the issue, variational inference (VI) approximates the high-dimensional posterior of the probabilistic model. It recently gained interest in many applications. VI is used, to name a few, for word embedding [5, 7], paragraph embedding [17], and knowledge-graph embedding [6]. For more details on this topic, we refer the reader to Zhang et al. [30] and Blei et al. [9]. Variational inference has also proven to be a successful approach to learning hyper-parameters [8, 6]. Building on recent advances in variational inference, we develop in this work a variational expectation-maximization algorithm by interpreting the parameters of an MHP as latent variables of a probabilistic model.

3 Preliminary Definitions

3.1 Multivariate Hawkes Processes

Formally, a D -dimensional MHP is a collection of D univariate counting processes $N_i(t)$, $i = 1, \dots, D$, whose realization over an observation period $[0, T)$ consists of a sequence of discrete events $\mathcal{S} = \{(t_n, i_n)\}$, where $t_n \in [0, T)$ is the timestamp of the n -th event and $i_n \in \{1, \dots, D\}$ is

its dimension. Each process has the particular form of conditional intensity function

$$\lambda_i(t) = \mu_i + \sum_{j=1}^D \int_0^t \phi_{ij}(t-\tau) dN_j(\tau), \quad (1)$$

where $\mu_i > 0$ is the constant exogenous part of the intensity of process i , and the excitation function $\phi_{ij} : \mathbb{R}_+ \mapsto \mathbb{R}_+$ encodes the effect of past events from dimension j on future events from dimension i . The larger the values of $\phi_{ij}(t)$, the more likely events in dimension j will trigger events in dimension i . It has been shown that the excitation matrix $[\phi_{ij}(t)]$ encodes the causal structure of the MHP in terms of Granger causality, *i.e.*, $\phi_{ij}(t) = 0$ if and only if the process j does not Granger-cause process i [13, 12].

Most of the literature uses a parametric form for the excitation functions. The most popular form is the exponential excitation function

$$\phi_{ij}(t) = w_{ij} \zeta e^{-\zeta t}. \quad (2)$$

However, in most applications the excitation patterns are unknown and this form might be too restrictive. Hence, to alleviate the assumption of a particular form for the excitation function, other approaches [19, 16, 28] over-parameterize the space and encode the excitation functions as a linear combination of M basis functions $\kappa_1(t), \kappa_2(t), \dots, \kappa_M(t)$ as

$$\phi_{ij}(t) = \sum_{m=1}^M w_{ij}^m \kappa_m(t), \quad (3)$$

where the basis functions are often exponential or Gaussian kernels [28]. This kind of approach is generally referred to as non-parametric. In the experiments of Section 5, we investigate the performance of both parametric and non-parametric approaches to learning MHPs from small sequences of observations. We denote the set of exogenous rates by $\boldsymbol{\mu} = \{\mu_i\}_{i=1}^D \in \mathbb{R}_+^D$ and the excitation matrix by $\mathbf{W} := \{\{w_{ij}^m\}_{m=1}^M\}_{i,j=1}^D \in \mathbb{R}_+^{D^2 M}$.

3.2 Maximum Likelihood Estimation

Suppose that we observe a sequence of discrete events $\mathcal{S} = \{(t_n, i_n)\}_{n=1}^N$ over an observation period $[0, T)$. The most common approach to learning the parameters of an MHP given \mathcal{S} is to perform a regularized maximum-likelihood estimation [31, 3, 28], which amounts to minimizing an objective function that is the sum of the negative log-likelihood and of a penalty term that induces some desired structural properties. Specifically, the objective is to solve the optimization problem

$$\hat{\boldsymbol{\mu}}, \hat{\mathbf{W}} = \underset{\boldsymbol{\mu} \geq 0, \mathbf{W} \geq 0}{\operatorname{argmin}} -\log p(\mathcal{S} | \boldsymbol{\mu}, \mathbf{W}) + \frac{1}{\alpha} \mathcal{R}(\boldsymbol{\mu}, \mathbf{W}), \quad (4)$$

where the log-likelihood of the parameters is given by

$$\log p(\mathcal{S} | \boldsymbol{\mu}, \mathbf{W}) = \sum_{(t_n, i_n) \in \mathcal{S}} \log \lambda_{i_n}(t_n) - \sum_{i=1}^D \int_0^T \lambda_i(t) dt. \quad (5)$$

The particular choice of penalty $\mathcal{R}(\boldsymbol{\mu}, \mathbf{W})$, along with the single hyper-parameter α controlling its influence, depends on the problem at hand. For example, a necessary condition to ensure that the learned model is stable is that $\lim_{t \rightarrow \infty} \phi_{ij}(t) = 0$ and that the spectral radius of the excitation matrix is less than 1 [10]. Hence, a common penalty used is

$$\mathcal{R}_p(\boldsymbol{\mu}, \mathbf{W}) = \sum_{i,j=1}^d \sum_{m=1}^M |w_{ij}^m|^p, \quad (6)$$

with $p = 1$ or 2 in [32, 31, 28]. Another common assumption is that the graph is sparse. In this case, a Group-Lasso penalty of the form

$$\mathcal{R}_{1,2}(\boldsymbol{\mu}, \mathbf{W}) = \sum_{i,j=1}^d \sqrt{\sum_{m=1}^M (w_{ij}^m)^2} \quad (7)$$

is commonly used to enforce sparsity in the excitation functions [28].

Small data amplifies the danger of overfitting; hence the choice of regularizers and their hyper-parameters becomes essential. Nevertheless, to control the influence of the penalty in (4), all state-of-the-art methods are limited by the use of a single shared hyper-parameter α . Ideally, we would have a different hyper-parameter to independently control the effect of the penalty on each parameter of the model. However, the number of parameters, *i.e.*, $(D^2M + D)$, grows quadratically with the dimension of the problem D . To make matters worse, the most common approaches used to fine-tune the choice of hyper-parameters, *i.e.*, grid search and random search, become computationally prohibitive when the number of hyper-parameters becomes large. Indeed, the search space exponentially increases with the number of hyper-parameters. Another approach is to use Bayesian optimization of hyper-parameters, but the cost of doing this also becomes prohibitive as the number of samples required to learn the landscape of cost function exponentially increases with the number of hyper-parameters [27]. We describe the details of our proposed approach in the next section.

4 Proposed Learning Approach

We now introduce the proposed approach for learning MHPs. The approach enables us to use different hyper-parameters for each model parameter and efficiently tune them all by taking into account parameter uncertainty. It is based on the variational expectation-maximization (EM) algorithm and jointly optimizes both the model parameters μ and \mathbf{W} , as well as the hyper-parameters α .

First, we can view regularized MLE as a maximum a posteriori (MAP) estimator of the model where parameters are considered as latent variables. Under this interpretation, regularizers on the model parameters correspond to unnormalized priors on the latent variables. The optimization problem becomes

$$\hat{\mu}, \hat{\mathbf{W}} = \underset{\mu \geq 0, \mathbf{W} \geq 0}{\operatorname{argmax}} \log p_{\alpha}(\mu, \mathbf{W}, \mathcal{S}) = \underset{\mu \geq 0, \mathbf{W} \geq 0}{\operatorname{argmax}} \log p(\mathcal{S} | \mu, \mathbf{W}) + \log p_{\alpha}(\mu, \mathbf{W}). \quad (8)$$

Therefore, having a better regularizer means having a better prior. In the presence of a long sequence of observations, we want the prior to be as uninformative as possible (*i.e.*, a smaller regularization) as we have access to enough information for the MLE to accurately estimate the parameters of the model. But in the case where we only observe short sequences, we want to use more informative priors to avoid overfitting (*i.e.*, a larger regularization).

Unfortunately, the MAP estimator cannot adjust the influence of the prior by optimizing over α . Indeed, the cost function in (8) is unbounded from above and solving Equation (8) with respect to α leads trivially to a divergent solution $\frac{1}{\alpha} \rightarrow \infty$. To address this issue, we can take a Bayesian approach, integrate out parameters and optimize the evidence (or marginal likelihood) $p_{\alpha}(\mathcal{S})$ instead of the log-likelihood. Such an approach changes the optimization problem of Equation (8) into

$$\hat{\alpha} = \underset{\alpha \geq 0}{\operatorname{argmax}} \iint p(\mathcal{S} | \mu, \mathbf{W}) p_{\alpha}(\mu, \mathbf{W}) d\mu d\mathbf{W} = \underset{\alpha \geq 0}{\operatorname{argmax}} p_{\alpha}(\mathcal{S}). \quad (9)$$

Unlike the MAP objective function, maximizing the evidence over α does not lead to a degenerate solution because it is upper bounded by the likelihood. However, this optimization problem can be solved only for simple models where the integral has a closed form, which requires a conjugate prior to the likelihood. Therefore, we use variational inference to estimate the evidence and develop a variational EM algorithm to optimize our objective with respect to α .

4.1 Variational Expectation-Maximization for Multivariate Hawkes Processes

Variational inference. The derivation of the variational objective is as follows. First postulate a variational distribution $q_{\gamma}(\mu, \mathbf{W})$, parameterized by the variational parameters γ , approximating the posterior $p(\mu, \mathbf{W} | \mathcal{S})$. The variational parameters γ are chosen such that the Kullback–Leibler divergence between the true posterior $p(\mu, \mathbf{W} | \mathcal{S})$ and the variational distribution $q_{\gamma}(\mu, \mathbf{W})$ is minimized. It is known that minimizing $KL[q_{\gamma}(\mu, \mathbf{W}) || p(\mu, \mathbf{W} | \mathcal{S})]$ is equivalent to maximizing the *evidence lower-bound* (ELBO) [9, 30] defined as

$$\text{ELBO}(q_{\gamma}, \alpha) := \mathbb{E}_{q_{\gamma}}[\log p_{\alpha}(\mu, \mathbf{W}, \mathcal{S})] - \mathbb{E}_{q_{\gamma}}[\log q_{\gamma}(\mu, \mathbf{W})]. \quad (10)$$

By invoking Jensen’s inequality on the integral $p_{\alpha}(\mathcal{S}) = \iint p_{\alpha}(\mu, \mathbf{W}, \mathcal{S}) d\mu d\mathbf{W}$, we obtain the desired lower bound on the evidence $p_{\alpha}(\mathcal{S}) \geq \text{ELBO}(q_{\gamma}, \alpha)$ where, by maximizing $\text{ELBO}(q_{\gamma}, \alpha)$ with respect to γ , the bound becomes tighter.

For simplicity, we adopt the mean-field assumption by choosing a variational distribution $q_\gamma(\boldsymbol{\mu}, \mathbf{W})$ that factorizes over the latent variables². As the parameters $\boldsymbol{\mu}$ and \mathbf{W} of an MHP are non-negative, a good candidate to approximate the posterior is a log-normal distribution. We define the variational parameters $\gamma = \{\boldsymbol{\nu}, e^\sigma\}$ as the mean and the standard deviation of q_γ . We denote the standard deviation by e^σ because, we optimize its log to naturally ensure its positivity and the stability of the optimization procedure. Although we present our learning approach for the log-normal distribution, it is easily generalizable to other distributions.

Variational EM algorithm. In order to efficiently optimize the ELBO with respect to both the variational parameters γ and the hyper-parameters α , we use the variational EM algorithm that iterates over the two following steps: The E-step maximizes the ELBO with respect to the variational parameters γ in order to get a tighter lower-bound on the evidence; and the M-step updates the hyper-parameters α with a closed form update. Details of the two steps are as follows.

The E-step maximizes the ELBO with respect to the variational parameters γ to make the variational distribution $q_\gamma(\boldsymbol{\mu}, \mathbf{W})$ close to the exact posterior $p(\boldsymbol{\mu}, \mathbf{W}|\mathcal{S})$ and to ensure that the ELBO is a good proxy for the evidence. To evaluate the ELBO, we use the black-box variational-inference optimization in [18, 25]. Re-parameterize the model as

$$\begin{aligned}\boldsymbol{\mu} &= g_\gamma(\boldsymbol{\varepsilon}^\mu) = \exp(\boldsymbol{\nu}^\mu + e^{\sigma^\mu} \odot \boldsymbol{\varepsilon}^\mu), \\ \mathbf{W} &= g_\gamma(\boldsymbol{\varepsilon}^\mathbf{W}) = \exp(\boldsymbol{\nu}^\mathbf{W} + e^{\sigma^\mathbf{W}} \odot \boldsymbol{\varepsilon}^\mathbf{W}),\end{aligned}$$

where $\boldsymbol{\varepsilon}^\mu$ (resp. $\boldsymbol{\varepsilon}^\mathbf{W}$) has the same shape as $\boldsymbol{\mu}$ (resp. \mathbf{W}), with each element following a normal distribution $\mathcal{N}(0, 1)$. \odot denotes the element-wise product. This trick enables us to rewrite the first intractable expectation term of the ELBO in (10) as

$$\mathbb{E}_{q_\gamma} [\log p_C(\boldsymbol{\mu}, \mathbf{W}, \mathcal{S})] = \mathbb{E}_{\boldsymbol{\varepsilon} \sim \mathcal{N}(0, I)} [\log p_\alpha(g_\gamma(\boldsymbol{\varepsilon}^\mu), g_\gamma(\boldsymbol{\varepsilon}^\mathbf{W}), \mathcal{S})]. \quad (11)$$

The second term of the ELBO in (10) is the entropy of the log-normal distribution that can be expressed, up to a constant, as $\sum_{\mu_i, \sigma_i} (\mu_i + \sigma_i)$. The ELBO then can be estimated by Monte-Carlo integration as

$$\text{ELBO}(q_\gamma, \alpha) \approx \frac{1}{L} \sum_{\ell=1}^L \log p_\alpha(g_\gamma(\boldsymbol{\varepsilon}_\ell^\mu), g_\gamma(\boldsymbol{\varepsilon}_\ell^\mathbf{W}), \mathcal{S}) + \sum_{\mu_i, \sigma_i} (\mu_i + \sigma_i), \quad (12)$$

where L is the number of Monte-Carlo samples $\boldsymbol{\varepsilon}_1, \dots, \boldsymbol{\varepsilon}_L$. Note that the first term of (12) is the cost function for the MAP problem (8) evaluated at $\{\boldsymbol{\mu}, \mathbf{W}\} = \{g_\gamma(\boldsymbol{\varepsilon}_\ell^\mu), g_\gamma(\boldsymbol{\varepsilon}_\ell^\mathbf{W})\}$ for $\ell \in [n]$. Hence, the E-step summarizes into maximizing the right-hand side of (12) with respect to γ using gradient descent.

In the M-step, the ELBO is used as a proxy for the evidence $p_\alpha(\mathcal{S})$ and is maximized with respect to the hyper-parameters α . Again, we rely on the re-parameterization technique and compute the unbiased estimate of the ELBO in (12). The maximum of the estimate (12) with respect to α has a closed form that depends on the choice of prior. We provide the closed-form solutions in Appendix A for a few proposed priors that emulate common regularizers. To avoid fast changes in α due to the variance of the Monte-Carlo integration, we take an update similar to the one in [6] and take a weighted average between the current estimate and the minimizer of the current Monte-Carlo estimate of the ELBO as

$$\boldsymbol{\alpha} := \beta \cdot \boldsymbol{\alpha} + (1 - \beta) \cdot \underset{\tilde{\boldsymbol{\alpha}}}{\operatorname{argmin}} \frac{1}{L} \sum_{\ell=1}^L \log p_{\tilde{\boldsymbol{\alpha}}}(g_\gamma(\boldsymbol{\varepsilon}_\ell^\mu), g_\gamma(\boldsymbol{\varepsilon}_\ell^\mathbf{W}), \mathcal{S}), \quad (13)$$

where $\beta \in [0, 1]$ is the momentum term.

Algorithm 1 summarizes the proposed variational EM approach.³ The computational complexity of the inner-most loop of Algorithm 1 is L times the complexity of an iteration of gradient descent on the log-likelihood. However, as observed by recent studies in variational inference, using $L = 1$ is usually sufficient in many applications [18]. Hence, we use $L = 1$ in all our experiments, leading to the same computational complexity per-iteration as MLE using gradient descent.

²This assumption can be relaxed using more advanced techniques at the cost of having a higher computational complexity.

³Source code is available publicly.

Algorithm 1 Variational EM algorithm for Multivariate Hawkes Processes

Input: Sequence of observations $\mathcal{S} = \{(t_n, i_n)\}_{n=1}^N$. Initial values for α and γ . Momentum term $0 \leq \beta < 1$. Sample size L of Monte-Carlo integrations. Number of iterations T_E and T_{EM} of E-steps and EM-steps. Learning rate η .

- 1: **for** $t \leftarrow 1, \dots, T_{EM}$ **do**
- 2: **for** $t \leftarrow 1, \dots, T_E$ **do** ▷ E step
- 3: Sample Gaussian noise $\varepsilon_1, \dots, \varepsilon_L \sim \mathcal{N}(0, I)$.
- 4: Evaluate the ELBO using Equation (12)
- 5: Update $\nu \leftarrow \nu + \eta(\nabla_\nu f(\nu, \sigma, \varepsilon; \alpha) + \mathbf{1})$.
- 6: Update $\sigma \leftarrow \sigma + \eta(\nabla_\sigma f(\nu, \sigma, \varepsilon; \alpha) + \mathbf{1})$.
- 7: **end for**
- 8: Sample L Gaussian noise $\varepsilon_1, \dots, \varepsilon_L$. ▷ M step
- 9: Update α using Equation (13).
- 10: **end for**

Output: α, γ

5 Experimental Results

We carry out two sets of experiments. First, we perform a link-prediction task on synthetic data to show that our approach can accurately recover the support of the excitation matrix of the MHP under short sequences. Second, we perform an event-prediction task on real datasets of short sequences to show that our approach outperforms state-of-the-art methods in terms of predictive log-likelihood.

We run our experiments in two different settings. First, in a *parametric* setting where the exponential form of the excitation function is known, we compare our approach (VI-EXP) to the state-of-the-art MLE-based method (MLE-ADM4) from Zhou et al. [31]. Second, we use a *non-parametric* setting where no assumption is made on the shape of the excitation function. We then set the excitation function as a mixture of $M = 10$ Gaussian kernels defined as

$$\kappa_m(t) = (2\pi b^2)^{-1} \exp(-(t - \tau_m)^2 / (2b^2)), \forall m = 1, \dots, M, \quad (14)$$

where τ_m and b are the known location and scale of the kernel. In this setting, we compare our approach (VI-SG) to the state-of-the-art MLE-based methods (MLE-SGLP) of Xu et al. [28] with the same $\{\kappa_m(t)\}$ [4]. Let us stress that the parametric methods have a strong advantage over the non-parametric ones because they are given the true value of the exponential decay ζ .

As our VI approach returns a posterior on the parameters, rather than a point estimate, we use the mode of the approximate log-normal posterior as the inferred edges $\{\hat{w}_{ij}\}$. For the non-parametric setting, we use $\hat{w}_{ij} = \sum_{m=1}^M \hat{w}_{ij}^m$. To mimic the regularization schemes of the baselines, we use a Laplacian prior for the edge weights $\{w_{ij}\}$ to enforce sparsity, and we use a Gaussian prior for the baselines $\{\mu_i\}$. We tune the hyper-parameters of the baselines using grid search [5].

5.1 Synthetic Data

We first evaluate the performance of our VI approach on simulated data. We generate random Erdős–Rényi graphs with $D = 50$ nodes and edge probability $p = \log(D)/D$. Then, a sequence of observations is generated from an MHP with exponential excitation functions defined in (2) with exponential decay $\zeta = 1$. The baselines $\{\mu_i^*\}$ are sampled independently in $\text{Unif}[0, 0.02]$, and the edge weights $\{w_{ij}^*\}$ are sampled independently in $\text{Unif}[0.1, 0.2]$. Results are averaged over 30 graphs with 10 simulations each. For reproducibility, a detailed description of the experimental setup is provided in Appendix E.

To investigate if the support of the excitation matrix can be accurately recovered under small data, we evaluate the performance of each approach on three metrics [32, 28, 14].

⁴ We also performed the experiments with other approaches designed for large-scale datasets, but their performance was below that of the reported baselines [1, 20, 21].

⁵ More details are provided in Appendix E.

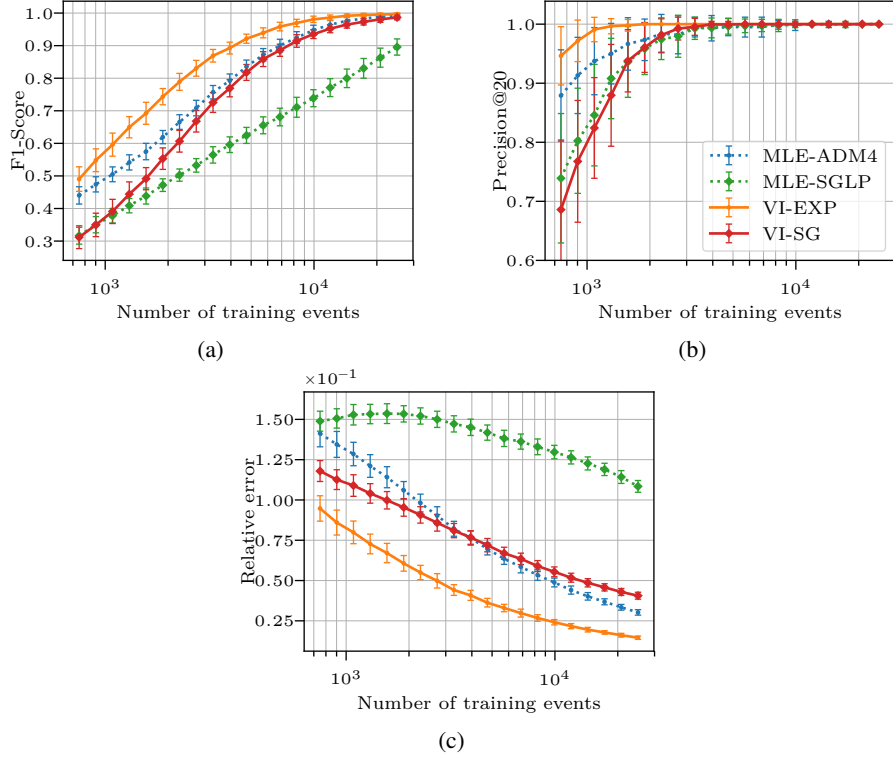


Figure 1: Performance measured by (a) F1-Score, (b) Precision@20, and (c) Relative error with respect to the number of training samples. Our VI approaches are shown in solid lines. The non-parametric methods are highlighted with square markers. Results are averaged over 30 random graphs with 10 simulations each (\pm standard deviation).

- **F1-score.** We zero-out small weights using a threshold $\eta = 0.04$ and measure the F1-score of the resulting binary edge classification problem⁶
- **Precision@k.** Instead of thresholding, we also report the precision@k defined by the average fraction of correctly identified edges in the top k largest estimated weights. Since the proposed VI approach gives an estimate of uncertainty via the variance of the posterior, we select the edges with high weights \hat{w}_{ij} and low uncertainty, *i.e.*, the edges with ratio of lowest standard deviation over weight \hat{w}_{ij} .
- **Relative error.** To evaluate the distance of the estimated weights to the ground truth ones, we use the averaged relative error defined as $|\hat{w}_{ij} - w_{ij}^*|/w_{ij}^*$ when $w_{ij}^* \neq 0$, and $\hat{w}_{ij}/(\min_{w_{kl}^* > 0} w_{kl}^*)$ when $w_{ij}^* = 0$. This metric is more sensitive to errors in small weights w_{ij}^* , and therefore penalizes false positive over false negative errors.

We investigate the sensitivity of each approach to the amount of data available for training by varying the size of the training set from $N = 750$ to $N = 25000$ events, *i.e.*, 15 to 500 events per node. Results are shown in Figure 1. Our approach improves the results in both parametric and non-parametric settings for all metrics. The improvements are more substantial in the non-parametric setting. If the accuracy of the top edges is similar for both VI-SG and MLE-SGLP in terms of precision@20, VI-SG improves the F1-score by about 20% with $N = 5000$ training events. The reason for this improvement is that MLE-SGLP has a much higher false positive rate, which is hurting the F1-score but does not affect the precision@20. VI-SG is also able to reach the same F1-score as the parametric baseline MLE-ADM4 with only $N = 4000$ training events⁷. Note that VI-SG is optimizing $D^2M + D = 25050$ hyper-parameters with minimal additional cost.

⁶Additional results with varying thresholds η are provided in Appendix D.

⁷We present additional results with various thresholds η and k in Appendix D.

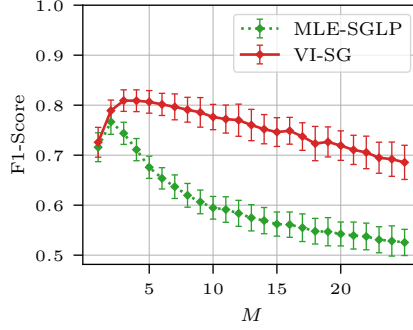


Figure 2: Analysis of the robustness of non-parametric approaches to the number of bases M of excitation functions (for fixed $N = 2000$).

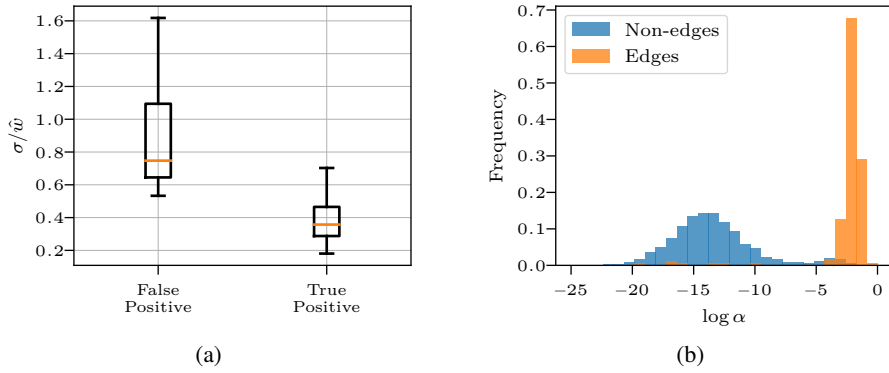


Figure 3: Analysis of the uncertainty of the parameters learned by VI-EXP (for fixed $N = 5000$). (a) Uncertainty of the inferred edges and (b) histogram of learned α . The learned α are smaller for non-edges, and false positive edges have higher uncertainty than the true positive ones.

In the next experiment, we focus on the non-parametric setting, we fix the length of observation to $N = 5000$ and study the effect of increasing M on the performance of the algorithms. The results are shown in Figure 2. We see that our approach is more robust to the choice of M than MLE-SGLP. A possible explanation for this behavior is that MLE-SGLP overfits due to the increasing number of model parameters.

Finally, we investigate the parameters of the model learned by our VI-EXP approach. In Figure 3a, we use the variance of the approximated posterior q_γ as a measure of confidence for edge identification, and we report the distribution of ratio of standard deviation over weight \hat{w}_{ij} for both the true and false positive edges. Similar results hold between the true and false negative edges. The false positive edges have a higher uncertainty than the true positive ones. This is relevant when we cannot identify all edges due to lack of data, even though we still wish to identify a subset of edges with high confidence. In addition, Figure 3b confirms that, as expected, the optimized weight priors α are much larger for true edges in the ground-truth excitation matrix than for non-edges.

5.2 Real Data

We also evaluate the performance of our approach on the following three small datasets:

1. **Epidemics.** This dataset contains records of infection of individuals, along with their corresponding district of residence, during the last Ebola epidemic in West Africa in 2014-2015 [15]. To learn the propagation network of the epidemics, we consider the 54 districts as processes and define infection records as events.

2. **Stock market.** This dataset contains the stock prices of 12 high-tech companies sampled every 2 minutes on the New York Stock Exchange for 20 days in April 2008 [13]. We consider each stock as a process and record an event every time a stock price changes by 0.15% from its current value.
3. **Enron email.** This dataset contains emails between employees of Enron from the Enron corpus. We consider all employees with more than 10 received emails as processes and record an event every time an employee receives an email.

We perform an event-prediction task to show that our approach outperforms the state-of-the-art methods in terms of predictive log-likelihood. To do so, we use the first 70% events as training set, and we compute the held-out averaged log-likelihood on the remaining 30%. We present the results in Table II.

We first see that the non-parametric methods outperform the parametric ones on both the Epidemic dataset and the Stock market dataset. This suggests that the exponential excitation function might be too restrictive to fit their excitation patterns. In addition, our non-parametric approach VI-SG significantly outperforms MLE-SGLP on all datasets. The improvement is particularly clear for the Epidemic dataset, which has the smallest number of events per dimension. Indeed, the top edges learned by VI-SG correspond to contiguous districts as expected. This is not the case for MLE-SGLP, for which the top learned edges correspond to districts that are far from each other.

Table 1: Predictive log-likelihood for the models learned on various real datasets.

| Dataset | Statistics | | Averaged predictive log-likelihood | | | |
|--------------|--------------|-----------------|------------------------------------|----------|--------|----------|
| | #dim (D) | #events (N) | VI-SG | MLE-SGLP | VI-EXP | MLE-ADM4 |
| Epidemics | 54 | 5 349 | −2,06 | −3,03 | −4,31 | −4,61 |
| Stock market | 12 | 7 089 | −1,00 | −2,45 | −2,82 | −2,81 |
| Enron email | 143 | 74 294 | −0,42 | −1,01 | −0,23 | −0,40 |

6 Conclusion

We proposed a novel approach to learn the excitation matrix of a multivariate Hawkes process in the presence of short observation sequences. We observed that state-of-the-art methods are sensitive to the amount of data used for training and showed that the proposed approach outperforms these methods when only short training sequences are available. The common tool to tackle this problem is to design smarter regularization schemes. However, all maximum likelihood-based methods suffer from a common problem: all the model parameters are regularized equally with a few hyper-parameters. We developed a variational expectation maximization algorithm that is able to (1) optimize over an extended set of hyper-parameters, with almost no additional cost and (2) take into account the uncertainty of the learned model parameters by fitting a posterior distribution over them. We performed experiments on both synthetic and real datasets and showed that our approach outperforms state-of-the-art methods under small-data regimes.

Acknowledgments

We would like to thank Negar Kiyavash and Jalal Etesami for the valuable discussions and insightful feedback at the early stage of this work. The work presented in this paper was supported in part by the Swiss National Science Foundation under grant number 200021-182407.

References

- [1] Massil Achab, Emmanuel Bacry, Stéphane Gaïffas, Iacopo Mastromatteo, and Jean-François Muzy. Uncovering causality from multivariate Hawkes integrated cumulants. *The Journal of Machine Learning Research*, 18(1):6998–7025, 2017.
- [2] Andrew Arnold, Yan Liu, and Naoki Abe. Temporal causal modeling with graphical granger methods. In *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '07, pages 66–75, New York, NY, USA, 2007. ACM. ISBN

- 978-1-59593-609-7. doi: 10.1145/1281192.1281203. URL <http://doi.acm.org/10.1145/1281192.1281203>.
- [3] Emmanuel Bacry, Stéphane Gaïffas, and Jean-François Muzy. A generalization error bound for sparse and low-rank multivariate Hawkes processes. *arXiv preprint arXiv:1501.00725*, 2015.
- [4] Emmanuel Bacry, Iacopo Mastromatteo, and Jean-François Muzy. Hawkes processes in finance. *Market Microstructure and Liquidity*, 1(01):1550005, 2015.
- [5] Robert Bamler and Stephan Mandt. Dynamic word embeddings. In *Proceedings of the 34th international conference on Machine learning*, 2017.
- [6] Robert Bamler, Farnood Salehi, and Stephan Mandt. Augmenting and tuning knowledge graph embeddings. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, 2019.
- [7] Oren Barkan. Bayesian neural word embedding. In *AAAI*, pages 3135–3143, 2017.
- [8] JM Bernardo, MJ Bayarri, JO Berger, AP Dawid, D Heckerman, AFM Smith, M West, et al. The variational bayesian em algorithm for incomplete data: with application to scoring graphical model structures. *Bayesian statistics*, 7:453–464, 2003.
- [9] David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.
- [10] D. J. Daley and D. Vere-Jones. *An introduction to the theory of point processes. Vol. I. Probability and its Applications* (New York). Springer-Verlag, New York, second edition, 2003. ISBN 0-387-95541-0. Elementary theory and methods.
- [11] Michael Eichler. Graphical modelling of multivariate time series. *Probability Theory and Related Fields*, 153(1):233–268, Jun 2012. ISSN 1432-2064. doi: 10.1007/s00440-011-0345-8. URL <https://doi.org/10.1007/s00440-011-0345-8>.
- [12] Michael Eichler, Rainer Dahlhaus, and Johannes Dueck. Graphical modeling for multivariate Hawkes processes with nonparametric link functions. *Journal of Time Series Analysis*, 38(2): 225–242, 2017.
- [13] Jalal Etesami, Negar Kiyavash, Kun Zhang, and Kushagra Singhal. Learning network of multivariate Hawkes processes: a time series approach. In *Proceedings of the Thirty-Second Conference on Uncertainty in Artificial Intelligence, UAI’16*, pages 162–171, Arlington, Virginia, United States, 2016. AUAI Press. ISBN 978-0-9966431-1-5. URL <http://dl.acm.org/citation.cfm?id=3020948.3020966>.
- [14] Flavio Figueiredo, Guilherme Borges, Pedro O. S. Vaz de Melo, and Renato Assunção. Fast estimation of causal interactions using Wold processes. In *Proceedings of the 32Nd International Conference on Neural Information Processing Systems, NIPS’18*, pages 2975–2986, USA, 2018. Curran Associates Inc. URL <http://dl.acm.org/citation.cfm?id=3327144.3327220>.
- [15] Tini Garske, Anne Cori, Archchun Ariyaratnam, Isobel M. Blake, Ilaria Dorigatti, Tim Eckmanns, Christophe Fraser, Wes Hinsley, Thibaut Jombart, Harriet L. Mills, Gemma Nedjati-Gilani, Emily Newton, Pierre Nouvellet, Devin Perkins, Steven Riley, Dirk Schumacher, Anita Shah, Maria D. Van Kerkhove, Christopher Dye, Neil M. Ferguson, and Christl A. Donnelly. Heterogeneities in the case fatality ratio in the West African ebola outbreak 2013-2016. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372(1721):20160308, 2017. doi: 10.1098/rstb.2016.0308. URL <https://royalsocietypublishing.org/doi/abs/10.1098/rstb.2016.0308>.
- [16] Niels Richard Hansen, Patricia Reynaud-Bouret, and Vincent Rivoirard. Lasso and probabilistic inequalities for multivariate point processes. *Bernoulli*, 21(1):83–143, 2015. doi: 10.3150/13-BEJ562. URL <https://hal.archives-ouvertes.fr/hal-00722668>. 61 pages.
- [17] Geng Ji, Robert Bamler, Erik B Sudderth, and Stephan Mandt. Bayesian paragraph vectors. *Symposium on Advances in Approximate Bayesian Inference*, 2017.

- [18] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *International Conference on Learning Representations (ICLR)*, 2014.
- [19] Remi Lemonnier and Nicolas Vayatis. Nonparametric markovian learning of triggering kernels for mutually exciting and mutually inhibiting multivariate Hawkes processes. In Toon Calders, Flavia Esposito, Eyke Hüllermeier, and Rosa Meo, editors, *Machine Learning and Knowledge Discovery in Databases*, pages 161–176, Berlin, Heidelberg, 2014. Springer Berlin Heidelberg.
- [20] Scott Linderman and Ryan Adams. Discovering latent network structure in point process data. In *International Conference on Machine Learning*, pages 1413–1421, 2014.
- [21] Scott W Linderman and Ryan P Adams. Scalable Bayesian inference for excitatory point process networks. *arXiv preprint arXiv:1507.03228*, 2015.
- [22] Scott W. Linderman, Yixin Wang, and David M Blei. Bayesian inference for latent Hawkes processes. *NeurIPS Symposium on Advances in Approximate Bayesian Inference Probabilistic*, 2017.
- [23] C. J. Quinn, N. Kiyavash, and T. P. Coleman. Directed information graphs. *IEEE Transactions on Information Theory*, 61(12):6887–6909, Dec 2015. ISSN 0018-9448. doi: 10.1109/TIT.2015.2478440.
- [24] Arvind Rao, Alfred O. Hero, David J. States, and James Douglas Engel. Using directed information to build biologically relevant influence networks. *Journal of Bioinformatics and Computational Biology*, 06(03):493–519, 2008. doi: 10.1142/S0219720008003515.
- [25] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *Proceedings of the 31st International Conference on Machine Learning*, 2014.
- [26] Thomas Schreiber. Measuring information transfer. *Phys. Rev. Lett.*, 85:461–464, Jul 2000. doi: 10.1103/PhysRevLett.85.461. URL <https://link.aps.org/doi/10.1103/PhysRevLett.85.461>.
- [27] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems*, pages 2951–2959, 2012.
- [28] Hongteng Xu, Mehrdad Farajtabar, and Hongyuan Zha. Learning granger causality for Hawkes processes. In *Proceedings of The 33rd International Conference on Machine Learning*, pages 1717–1726. PMLR, 2016. URL <http://proceedings.mlr.press/v48/xuc16.html>.
- [29] Hongteng Xu, Dixin Luo, and Hongyuan Zha. Learning Hawkes processes from short doubly-censored event sequences. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML’17*, pages 3831–3840. JMLR.org, 2017. URL <http://dl.acm.org/citation.cfm?id=3305890.3306077>.
- [30] Cheng Zhang, Judith Butepage, Hedvig Kjellstrom, and Stephan Mandt. Advances in variational inference. *IEEE transactions on pattern analysis and machine intelligence*, 2018.
- [31] Ke Zhou, Hongyuan Zha, and Le Song. Learning social infectivity in sparse low-rank networks using multi-dimensional Hawkes processes. In *AISTATS*, volume 31 of *JMLR Workshop and Conference Proceedings*, pages 641–649. JMLR.org, 2013.
- [32] Ke Zhou, Hongyuan Zha, and Le Song. Learning triggering kernels for multi-dimensional Hawkes processes. In *International Conference on Machine Learning*, volume 28, pages 1301–1309, 2013.