

Variational Inference for Sparse Gaussian Process Modulated Hawkes Process

Rui Zhang^{1,2} Christian Walder^{1,2} Marian-Andrei Rizoiu³

¹The Australian National University, ²Data61 CSIRO, ³University of Technology Sydney
 Rui.Zhang@anu.edu.au, Christian.Walder@data61.csiro.au, Marian-Andrei.Rizoiu@uts.edu.au

Abstract

The Hawkes process (HP) has been widely applied to modeling self-exciting events including neuron spikes, earthquakes and tweets. To avoid designing parametric triggering kernel and to be able to quantify the prediction confidence, the non-parametric Bayesian HP has been proposed. However, the inference of such models suffers from unscalability or slow convergence. In this paper, we aim to solve both problems. Specifically, first, we propose a new non-parametric Bayesian HP in which the triggering kernel is modeled as a squared sparse Gaussian process. Then, we propose a novel variational inference schema for model optimization. We employ the branching structure of the HP so that maximization of evidence lower bound (ELBO) is tractable by the expectation-maximization algorithm. We propose a tighter ELBO which improves the fitting performance. Further, we accelerate the novel variational inference schema to linear time complexity by leveraging the stationarity of the triggering kernel. Different from prior acceleration methods, ours enjoys higher efficiency. Finally, we exploit synthetic data and two large social media datasets to evaluate our method. We show that our approach outperforms state-of-the-art non-parametric frequentist and Bayesian methods. We validate the efficiency of our accelerated variational inference schema and practical utility of our tighter ELBO for model selection. We observe that the tighter ELBO exceeds the common one in model selection.

1 Introduction

The Hawkes process (HP) (Hawkes 1971) is particularly useful to model self-exciting point data – i.e., when the occurrence of a point increases the likelihood of occurrence of new points. The process is parameterized using a background intensity μ , and a triggering kernel ϕ . The Hawkes process can be alternatively represented as a cluster of Poisson processes (PPes) (Hawkes and Oakes 1974). In the cluster, a PP with an intensity μ (denoted as $PP(\mu)$) generates immigrant points which are considered to arrive in the system from the outside, and every existing point triggers offspring points, which are generated internally through the self-excitement, following a $PP(\phi)$. Points can therefore be structured into clusters where each cluster contains either a point and its direct offspring or the background process (an example is shown in Fig. 1a). Connecting all points using the triggering relations yields

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

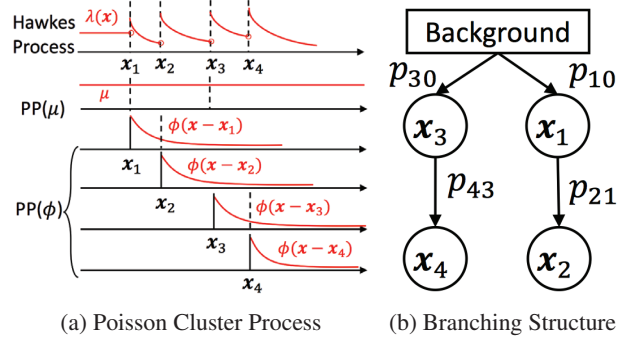


Figure 1: The Cluster Representation of a HP. (a) A HP with a decaying triggering kernel $\phi(\cdot)$ has intensity $\lambda(x)$ which increases after each new point (dash line) is generated. It can be represented as a cluster of PPes: $PP(\mu)$ and $PP(\phi(x - x_i))$ associated with each x_i . (b) The branching structure corresponding to the triggering relationships shown in (a), where an edge $x_i \rightarrow x_j$ means that x_i triggers x_j , and its probability is denoted as p_{ji} .

a tree structure, which is called the branching structure (an example is shown in Fig. 1b corresponding to Fig. 1a). With the branching structure, we can decompose the HP into a cluster of PPes. The triggering kernel ϕ is shared among all cluster Poisson processes relating to a HP, and it determines the overall behavior of the process. Consequently, designing the kernel functions is of utmost importance for employing the HP to a new application, and its study has attracted much attention.

Prior non-parametric frequentist solutions. In the case in which the optimal triggering kernel for a particular application is unknown, a typical solution is to express it using a non-parametric form, such as the work of Lewis and Mohler (2011); Zhou, Zha, and Song (2013); Bacry and Muzy (2014); Eichler, Dahlhaus, and Dueck (2017). These are all frequentist methods and among them, the Wiener-Hoef equation based method (Bacry and Muzy 2014) enjoys linear time complexity. Our method has a similar advantage of linear time complexity per iteration. The Euler-Lagrange equation based solutions (Lewis and Mohler 2011; Zhou, Zha, and Song 2013) require discretizing the input

domain so they face a problem of poorly scaling with the dimension of the domain. The same problem is also faced by Eichler, Dahlhaus, and Dueck (2017)’s discretization based method. In contrast, our method requires no discretization so enjoys scalability with the dimension of the domain.

Prior Bayesian solutions. The Bayesian inference for the HP has also been studied, including the work of Rasmussen (2013); Linderman and Adams (2014); Linderman and Adams (2015). These work require either constructing a parametric triggering kernel (Rasmussen 2013; Linderman and Adams 2014) or discretizing the input domain to scale with the data size (Linderman and Adams 2015). The shortcoming of discretization is just mentioned and to overcome it, Donnet, Rivoirard, and Rousseau (2018) propose a continuous non-parametric Bayesian HP and resort to an unscalable Markov chain Monte Carlo (MCMC) estimator to the posterior distribution. A more recent solution based on Gibbs sampling (Zhang et al. 2019) obtains linear time complexity per iteration by exploiting the HP’s branching structure and the stationarity of the triggering kernel similar to the work of Halpin (2012). This approach considers only high-probability triggering relationships in computations for acceleration. However, those relationships are updated in each iteration, which is less efficient than our pre-computing them. Besides, our variational inference schema enjoys faster convergence than that of the Gibbs sampling based method.

Contributions. In this paper, we propose the first sparse Gaussian process modulated HP which employs a novel variational inference schema, enjoys a linear time complexity per iteration and scales to large real world data. Our method is inspired by the variational Bayesian PP (VBPP) (Lloyd et al. 2015) which provides the Bayesian non-parametric inference only for the whole intensity of the HP without for its components: the background intensity μ and the triggering kernel ϕ . Thus, the VBPP loses the internal reactions between points, and developing the variational Bayesian non-parametric inference for the HP is non-trivial and more challenging than the VBPP. In this paper, we adapt the VBPP for the HP and term the new approach the variational Bayesian Hawkes process (VBHP). The contributions are summarized:

(1) (Sec.3) We introduce a new Bayesian non-parametric HP which employs a sparse Gaussian process modulated triggering kernel and a Gamma distributed background intensity.

(2) (Sec.3&4) We propose a new variational inference schema for such a model. Specifically, we employ the branching structure of the HP so that maximization of the evidence lower bound (ELBO) is tractable by the expectation-maximization algorithm, and we contribute a tighter ELBO which improves the fitting performance of our model.

(3) (Sec.5) We propose a new acceleration trick based on the stationarity of the triggering kernel. The new trick enjoys higher efficiency than prior methods and accelerates the variational inference schema to linear time complexity per iteration.

(4) (Sec.6) We empirically show that VBHP provides more accurate predictions than state-of-the-art methods on synthetic data and on two large online diffusion datasets. We validate the linear time complexity and faster convergence of our accelerated variational inference schema compared to

the Gibbs sampling method, and the practical utility of our tighter ELBO for model selection, which outperforms the common one in model selection.

2 Prerequisite

In this section, we review the Hawkes process, the variational inference and its application to the Poisson process.

2.1 Hawkes Process (HP)

The HP (Hawkes 1971) is a self-exciting point process, in which the occurrence of a point increases the arrival rate $\lambda(\cdot)$, a.k.a. the (conditional) intensity, of new points. Given a set of time-ordered points $\mathcal{D} = \{\mathbf{x}_i\}_{i=1}^N$, $\mathbf{x}_i \in \mathbb{R}^R$, the intensity at \mathbf{x} conditioned on given points is written as:

$$\lambda(\mathbf{x}) = \mu + \sum_{\mathbf{x}_i < \mathbf{x}} \phi(\mathbf{x} - \mathbf{x}_i),$$

where $\mu > 0$ is a constant background intensity, and $\phi : \mathbb{R}^R \rightarrow [0, \infty)$ is the triggering kernel.

We are particularly interested in the branching structure of the HP. As introduced in Sec.1, each point \mathbf{x}_i has a parent that we represent by a one-hot vector $\mathbf{b}_i = [b_{i0}, b_{i1}, \dots, b_{i,i-1}]^T$: each element b_{ij} is binary, $b_{ij} = 1$ represents that \mathbf{x}_i is triggered by \mathbf{x}_j ($0 \leq j \leq i-1$, \mathbf{x}_0 : the background), and $\sum_{j=0}^{i-1} b_{ij} = 1$. A branching structure B is a set of \mathbf{b}_i , namely $B = \{\mathbf{b}_i\}_{i=1}^N$, and if the probability of $b_{ij} = 1$ is defined as $p_{ij} \equiv p(b_{ij} = 1)$, the probability of B can be expressed as $p(B) = \prod_{i=1}^N \prod_{j=0}^{i-1} p_{ij}^{b_{ij}}$. Since \mathbf{b}_i is a one-hot vector, there is $\sum_{j=0}^{i-1} p_{ij} = 1$ for all i .

Given both \mathcal{D} and a branching structure B , the log likelihood of μ and ϕ becomes a sum of log likelihoods of PPs. With B , the HP \mathcal{D} can be decomposed into $\text{PP}(\mu)$ and $\{\text{PP}(\phi(\mathbf{x} - \mathbf{x}_i))\}_{i=1}^N$. The data domain of $\text{PP}(\mu)$ equals that of the HP, which we denote \mathcal{T} , and we denote the data domain of $\text{PP}(\phi(\mathbf{x} - \mathbf{x}_i))$ by $\mathcal{T}_i \subset \mathcal{T}$. As a result, the log likelihood $\log p(\mathcal{D}, B | \mu, \phi)$ is expressed as:

$$\log p(\mathcal{D}, B | \mu, \phi) = \sum_{i=1}^N \left(\sum_{j=1}^{i-1} b_{ij} \log \phi_{ij} + b_{i0} \log \mu \right) - \sum_{i=1}^N \int_{\mathcal{T}_i} \phi - \mu |\mathcal{T}|, \quad (1)$$

where $|\mathcal{T}| \equiv \int_{\mathcal{T}} 1 \, d\mathbf{x}$, $\phi_{ij} \equiv \phi(\mathbf{x}_i - \mathbf{x}_j)$ and $\int_{\mathcal{T}_i} \phi \equiv \int_{\mathcal{T}_i} \phi(\mathbf{x}) \, d\mathbf{x}$. Throughout this paper, we simplify the integration by omitting the integration variable, which is \mathbf{x} unless otherwise specified. This work is developed based on the univariate HP. However, it can be extended to be multivariate following the same procedure.

2.2 Variational Inference (VI)

Consider the latent variable model $p(\mathbf{x}, \mathbf{z} | \theta)$ where \mathbf{x} and \mathbf{z} are the data and the latent variables respectively. The variational approach introduces a variational distribution to approximate the posterior distribution $q(\mathbf{z} | \theta') \approx p(\mathbf{z} | \mathbf{x}, \theta)$ and maximizes a lower bound of the log-likelihood, which can be derived from the non-negative gap perspective:

$$\log p(\mathbf{x} | \theta)$$

Table 1: Notations

VBHP	VBPP	VI
$\mathcal{D} \equiv \{\mathbf{x}_n\}_{n=1}^N$	$\mathcal{D} \equiv \{\mathbf{x}_n\}_{n=1}^N$	\mathbf{x}
B, μ, f, \mathbf{u}	f, \mathbf{u}	\mathbf{z}
$k_0, c_0, \{\alpha_i\}_{i=1}^R, \gamma$	$\{\alpha_i\}_{i=1}^R, \gamma$	$\boldsymbol{\theta}$
$k, c, \mathbf{m}, \mathbf{S}, \{\alpha_i\}_{i=1}^R, \gamma, \{\{q_{ij}\}_{j=0}^{i-1}\}_{i=1}^N$	$\mathbf{m}, \mathbf{S}, \{\alpha_i\}_{i=1}^R, \gamma$	$\boldsymbol{\theta}'$

$$\begin{aligned}
 &= \log \frac{p(\mathbf{x}, \mathbf{z}|\boldsymbol{\theta})}{q(\mathbf{z}|\boldsymbol{\theta}')} - \log \frac{p(\mathbf{z}|\boldsymbol{\theta})}{q(\mathbf{z}|\boldsymbol{\theta}')} \\
 &= \underbrace{\mathbb{E}_{q(\mathbf{z}|\boldsymbol{\theta}')} [\log p(\mathbf{x}|\mathbf{z}, \boldsymbol{\theta})]}_{\text{reconstruction term}} - \underbrace{\text{KL}(q(\mathbf{z}|\boldsymbol{\theta}')||p(\mathbf{z}|\boldsymbol{\theta}))}_{\text{regularisation term}} \\
 &\quad \underbrace{+ \text{KL}(q(\mathbf{z}|\boldsymbol{\theta}')||p(\mathbf{z}|\boldsymbol{\theta}))}_{\text{intractable (non-negative) gap}} \\
 &\geq \text{CELBO}(q(\mathbf{z}), p(\mathbf{x}|\mathbf{z}), p(\mathbf{z})), \tag{2}
 \end{aligned}$$

where we omit $\boldsymbol{\theta}$ and $\boldsymbol{\theta}'$ in conditions. We term this the Common ELBO (CELBO) to differentiate with our tighter ELBO of Sec.4. For notational convenience, we will often omit conditioning on $\boldsymbol{\theta}$ and $\boldsymbol{\theta}'$ hereinafter. Optimizing the CELBO w.r.t. $\boldsymbol{\theta}'$ balances between the reconstruction error and the Kullback-Leibler (KL) divergence from the prior. Generally, the conditional $p(\mathbf{x}|\mathbf{z})$ is known, so is the prior. Thus, for an appropriate choice of q , it is easier to work with this lower bound than with the intractable posterior $p(\mathbf{z}|\mathbf{x})$. We also see that, due to the form of the intractable gap, if q is from a distribution family containing elements close to the true unknown posterior, then q will be close to the true posterior when the CELBO is close to the true likelihood. An alternative derivation applies Jensen’s inequality (Jordan et al. 1999).

2.3 Variational Bayesian Poisson Process (VBPP)

VBPP (Lloyd et al. 2015) applies the VI to the Bayesian Poisson process, which exploits the sparse Gaussian process (GP) to model the Poisson intensity. Specifically, VBPP uses a squared link function to map a sparse GP distributed function f to the Poisson intensity $\lambda(\mathbf{x}) = f^2(\mathbf{x})$. The sparse GP employs the ARD kernel:

$$K(\mathbf{x}, \mathbf{x}') \equiv \gamma \prod_{r=1}^R \exp\left(-\frac{(x_r - x'_r)^2}{2\alpha_r}\right).$$

where γ and $\{\alpha_r\}_{r=1}^R$ are GP hyper-parameters. Let $\mathbf{u} \equiv (f(\mathbf{z}_1), f(\mathbf{z}_2), \dots, f(\mathbf{z}_M))$ where \mathbf{z}_i are inducing points. The prior and the approximate posterior distributions of \mathbf{u} are Gaussian distributions $p(\mathbf{u}) = \mathcal{N}(\mathbf{u}|\mathbf{0}, \mathbf{K}_{zz})$ and $q(\mathbf{u}) = \mathcal{N}(\mathbf{u}|\mathbf{m}, \mathbf{S})$ where \mathbf{m} and \mathbf{S} are the mean vector and the covariance matrix respectively. Note both \mathbf{u} and f employ zero mean priors. Notations of VBPP are connected with those of VI (Sec.2.2) in Table1.

Importantly, the variational joint distribution of f and \mathbf{u} uses the exact conditional distribution $p(f|\mathbf{u})$, i.e.,

$$q(f, \mathbf{u}) \equiv p(f|\mathbf{u})q(\mathbf{u}) \tag{3}$$

which in turn leads to the posterior GP:

$$\begin{aligned}
 q(f) &= \mathcal{N}(f|\nu, \Sigma), \tag{4} \\
 \nu(\mathbf{x}) &\equiv \mathbf{K}_{xz} \mathbf{K}_{zz}^{-1} \mathbf{m}, \\
 \Sigma(\mathbf{x}, \mathbf{x}') &\equiv \mathbf{K}_{xx'} + \mathbf{K}_{xz} \mathbf{K}_{zz}^{-1} (\mathbf{S} \mathbf{K}_{zz}^{-1} - I) \mathbf{K}_{zx'}.
 \end{aligned}$$

Then, the CELBO is obtained by using Eqn.(2):

$$\begin{aligned}
 &\text{CELBO}(q(f, \mathbf{u}), p(\mathcal{D}|f, \mathbf{u}), p(f, \mathbf{u})) \\
 &= \mathbb{E}_{q(f)} [\log p(\mathcal{D}|f)] - \text{KL}(q(\mathbf{u})||p(\mathbf{u})).
 \end{aligned}$$

Note that the second term is the KL divergence between two multivariate Gaussian distributions, so is available in closed form. The first term turns out to be the expectation w.r.t. $q(f)$ of the log-likelihood $\log p(\mathcal{D}|f) = \sum_{i=1}^N \log f^2(\mathbf{x}_i) - \int_{\mathcal{T}} f^2$. The expectation of the integral part is relatively straightforward to compute and the expectation of the other (data-dependent) part is available in almost closed-form with a hyper-geometric function.

3 Variational Bayesian Hawkes Process

3.1 Notations

To extend VBPP to HP, we introduce two more variables: the background intensity μ and the branching structure B , defined in Sec.2.1. We assume that the prior distribution of μ is a Gamma distribution $p(\mu) = \text{Gamma}(\mu|k_0, c_0)$ ¹ and the posterior distribution is approximated by another Gamma distribution $q(\mu) = \text{Gamma}(\mu|k, c)$. For $B = \{\mathbf{b}_i\}_{i=1}^N$ given $\mathcal{D} = \{\mathbf{x}_i\}_{i=1}^N$, we assume the variational posterior $b_{ij}, j = 0, \dots, i-1$, have a categorical distribution: $q_{ij} = q(b_{ij} = 1)$ and $\sum_{j=0}^{i-1} q_{ij} = 1$, and thus, the variational posterior probability of B is expressed as:

$$q(B) = \prod_{i=1}^N \prod_{j=0}^{i-1} q_{ij}^{b_{ij}}. \tag{5}$$

The same squared link function is adopted for the triggering kernel $\phi(\mathbf{x}) = f^2(\mathbf{x})$, so are the priors for f and \mathbf{u} , namely $\mathcal{N}(f|\mathbf{0}, \mathbf{K}_{xx'})$ and $\mathcal{N}(\mathbf{u}|\mathbf{0}, \mathbf{K}_{zz'})$. More link functions such as $\exp(\cdot)$ are discussed by Lloyd et al. (2015). Moreover, we use the same variational joint posterior on f and \mathbf{u} as Eqn.(3). Consequently, we complete the variational joint distribution on all latent variables as below:

$$q(B, \mu, f, \mathbf{u}) \equiv q(B)q(\mu)p(f|\mathbf{u})q(\mathbf{u}), \tag{6}$$

and notations of VBHP are summarized in Table 1.

3.2 CELBO

Based on Eqn.(2)&(6), we obtain the CELBO for VBHP (see details in A.1 of the appendix (App. 2019)):

$$\begin{aligned}
 &\text{CELBO}(q(B, \mu, f, \mathbf{u}), p(\mathcal{D}|B, \mu, f, \mathbf{u}), p(B, \mu, f, \mathbf{u})) \\
 &= \underbrace{\mathbb{E}_{q(B, \mu, f)} [\log p(\mathcal{D}, B|f, \mu)]}_{\text{Data Dependent Expectation (DDE)}} + H_B
 \end{aligned}$$

¹Gamma($\mu|k_0, c_0$) = $\frac{1}{\Gamma(k_0)c_0^k} \mu^{k_0-1} e^{-\mu/c_0}$

$$- \text{KL}(q(\mu)||p(\mu)) - \text{KL}(q(\mathbf{u})||p(\mathbf{u})). \quad (7)$$

where $H_B = -\sum_{i=1}^N \sum_{j=0}^{i-1} q_{ij} \log q_{ij}$ is the entropy of the variational posterior B . The KL terms are between gamma and Gaussian distributions for which closed forms are provided in A.2 of the appendix (App. 2019).

3.3 Data Dependent Expectation

Now, we are left with the problem of computing the data dependent expectation (DDE) in Eqn.(7). The DDE is w.r.t. the variational posterior probability $q(B, \mu, f)$. From Eqn.(6), $q(B, \mu, f) = \int q(B, \mu, f, \mathbf{u}) d\mathbf{u} = q(B)q(\mu)q(f)$ and $q(f)$ is identical to Eqn.(4). As a result, we can compute the DDE w.r.t. $q(B)$ first, and then w.r.t. $q(\mu)$ and $q(f)$.

Expectation w.r.t. $q(B)$. From Eqn.(1), we easily obtain $\log p(\mathcal{D}, B|f, \mu)$ by replacing ϕ with f^2 , whereupon it is clear that only b_{ij} in $\log p(\mathcal{D}, B|f, \mu)$ is dependent on B . Therefore, $\mathbb{E}_{q(B)}[\log p(\mathcal{D}, B|f, \mu)]$ is computed as:

$$\begin{aligned} & \mathbb{E}_{q(B)}[\log p(\mathcal{D}, B|f, \mu)] \\ &= \sum_{i=1}^N \left(\sum_{j=1}^{i-1} q_{ij} \log f_{ij}^2 + q_{i0} \log \mu \right) - \sum_{i=1}^N \int_{\mathcal{T}_i} f^2 - \mu | \mathcal{T} | \end{aligned}$$

where $f_{ij} \equiv f(\mathbf{x}_i - \mathbf{x}_j)$.

Expectation w.r.t. $q(f)$ and $q(\mu)$. We compute the expectation w.r.t. $q(f)$ and $q(\mu)$ by exploiting the expectation and the log expectation of the Gamma distribution: $\mathbb{E}_{q(\mu)}(\mu) = kc$ and $\mathbb{E}_{q(\mu)}(\log(\mu)) = \psi(k) + \log c$, and also the property $\mathbb{E}(x^2) = \mathbb{E}(x)^2 + \text{Var}(x)$:

$$\begin{aligned} \text{DDE} &= \sum_{i=1}^N \left[\sum_{j=1}^{i-1} q_{ij} \mathbb{E}_{q(f)}(\log f_{ij}^2) + q_{i0} (\psi(k) + \log c) \right. \\ & \quad \left. - \int_{\mathcal{T}_i} \mathbb{E}_{q(f)}^2(f) - \int_{\mathcal{T}_i} \text{Var}_{q(f)}(f) \right] - kc | \mathcal{T} |, \end{aligned}$$

where ψ is the Digamma function. We provide closed form expressions for $\int_{\mathcal{T}_i} \mathbb{E}_{q(f)}^2(f)$ and $\int_{\mathcal{T}_i} \text{Var}_{q(f)}(f)$ in A.3 of the appendix (App. 2019). As in VBPP, $\mathbb{E}_{q(f)}(\log f_{ij}^2) = -\tilde{G}(-\nu_{ij}^2/(2\Sigma_{ij})) + \log(\Sigma_{ij}/2) - C$ is available in the closed form with a hyper-geometric function, where $\nu_{ij} = \nu(\mathbf{x}_i - \mathbf{x}_j)$, $\Sigma_{ij} = \Sigma(\mathbf{x}_i - \mathbf{x}_j, \mathbf{x}_i - \mathbf{x}_j)$ (ν and Σ are defined in Eqn.(4)), $C \approx 0.57721566$ is the Euler-Mascheroni constant and \tilde{G} is defined as $\tilde{G}(z) = {}_1F_1^{(1,0,0)}(0, 1/2, z)$, i.e., the partial derivative of the confluent hyper-geometric function ${}_1F_1$ w.r.t. the first argument. We compute \tilde{G} using the method of Ancarani and Gasaneo (2008) and implement \tilde{G} and \tilde{G}' by linear interpolation of a lookup table (see a demo in the supplementary material).

3.4 Predictive Distribution of ϕ

The predictive distribution of $f(\mathbf{x})$ depends on the posterior \mathbf{u} . We assume that the optimal variational distribution of \mathbf{u} approximates the true posterior distribution, namely $q(\mathbf{u}|\mathcal{D}, \theta^*) = \mathcal{N}(\mathbf{u}|\mathbf{m}^*, \mathbf{S}^*) \approx p(\mathbf{u}|\mathcal{D}, \theta)$. Therefore, there is $q(f|\mathcal{D}, \theta^*) \approx p(f|\mathcal{D}, \theta)$, i.e., the approximate predictive $f(\tilde{\mathbf{x}}) \sim \mathcal{N}(\mathbf{K}_{\tilde{\mathbf{x}}z} \mathbf{K}_{zz}^{-1} \mathbf{m}^*, \mathbf{K}_{\tilde{\mathbf{x}}\tilde{\mathbf{x}}} - \mathbf{K}_{\tilde{\mathbf{x}}z} \mathbf{K}_{zz}^{-1} \mathbf{K}_{z\tilde{\mathbf{x}}})$

$\mathbf{K}_{\tilde{\mathbf{x}}z} \mathbf{K}_{zz}^{-1} \mathbf{S}^* \mathbf{K}_{zz}^{-1} \mathbf{K}_{z\tilde{\mathbf{x}}} \equiv \mathcal{N}(\tilde{\nu}, \tilde{\sigma}^2)$. Given the relation $\phi = f^2$, it is straightforward to derive the corresponding $\phi(\tilde{\mathbf{x}}) \sim \text{Gamma}(\tilde{k}, \tilde{c})$ where the shape $\tilde{k} = (\tilde{\nu}^2 + \tilde{\sigma}^2)^2 / [2\tilde{\sigma}^2(2\tilde{\nu}^2 + \tilde{\sigma}^2)]$ and the scale $\tilde{c} = 2\tilde{\sigma}^2(2\tilde{\nu}^2 + \tilde{\sigma}^2) / (\tilde{\nu}^2 + \tilde{\sigma}^2)$.

4 New Variational Inference Schema

We now propose a new variational inference (VI) schema which uses a tighter ELBO than the common one, i.e.

Theorem 1. For VBHP, there is a tighter ELBO

$$\underbrace{\mathbb{E}_{q(B, \mu, f)} \left[\log p(\mathcal{D}, B|f, \mu) \right]}_{\equiv \text{TELBO}} + H_B \leq \log p(\mathcal{D}).$$

Remark. TELBO is tighter because it is equivalent to the CELBO (Eqn.(7)) except without subtracting non-negative KL divergences over μ and \mathbf{u} . Other graphical models such as the variational Gaussian mixture model (Attias 1999) have a similar TELBO. Later on, we propose a new VI schema based on the TELBO.

Proof. With the variational posterior probability of the branching structure $q(B)$ defined in Eqn.(5) and through the Jensen's inequality, we have:

$$\log p(\mathcal{D}) \geq \sum_B q(B) \log p(\mathcal{D}, B) + H_B, \quad (8)$$

where H_B is the entropy of B defined in Eqn.(7). The term $\sum_B q(B) \log p(\mathcal{D}, B)$ can be understood as follows. Consider that infinite branching structures are drawn from $q(B)$ independently, say $\{B_i\}_{i=1}^\infty$. Given a branching structure B_i , the Hawkes process can be decomposed into a cluster of Poisson processes, denoted as (\mathcal{D}, B_i) , and the corresponding log-likelihood is $\log p(\mathcal{D}, B_i)$. Then, $\sum_B q(B) \log p(\mathcal{D}, B)$ is the mean of all log likelihoods $\{\log p(\mathcal{D}, B_i)\}_{i=1}^\infty$,

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \log p(\mathcal{D}, B_i) &= \lim_{n \rightarrow \infty} \sum_B \frac{n_B}{n} \log p(\mathcal{D}, B) \\ &= \sum_B q(B) \log p(\mathcal{D}, B), \quad (9) \end{aligned}$$

where n_B is the number of occurrences of branching structure B . Since all branching structures $\{B_i\}_{i=1}^\infty$ are i.i.d., the clusters of Poisson processes generated over $\{B_i\}_{i=1}^\infty$ should also be independent, i.e., $\{(\mathcal{D}, B_i)\}_{i=1}^\infty$ are i.i.d.. It follows that

$$\sum_{i=1}^\infty \log p(\mathcal{D}, B_i) = \log p(\{(\mathcal{D}, B_i)\}_{i=1}^\infty). \quad (10)$$

We compute the CELBO of $\log p(\{(\mathcal{D}, B_i)\}_{i=1}^\infty)$ by making $z = (\mu, f, \mathbf{u})$ and $\mathbf{x} = \{(\mathcal{D}, B_i)\}_{i=1}^n$ in Eqn.(2):

$$\begin{aligned} \log p(\{(\mathcal{D}, B_i)\}_{i=1}^n) &\geq E_{q(f, \mu)} \left[\log p(\{(\mathcal{D}, B_i)\}_{i=1}^n | f, \mu) \right] \\ &\quad - \text{KL}(q(\mu)||p(\mu)) - \text{KL}(q(\mathbf{u})||p(\mathbf{u})). \quad (11) \end{aligned}$$

Further, we plug Eqn.(10) and Eqn.(11) into Eqn.(9):

$$\text{Eqn.(9)} = \lim_{n \rightarrow \infty} \frac{1}{n} \log p(\{(\mathcal{D}, B_i)\}_{i=1}^n)$$

$$\begin{aligned}
&\stackrel{(a)}{\geq} \lim_{n \rightarrow \infty} \frac{1}{n} E_{q(f, \mu)} [\log p(\{(\mathcal{D}, B_i)\}_{i=1}^n | f, \mu)] \\
&\stackrel{(b)}{=} \lim_{n \rightarrow \infty} \sum_{n_B} \frac{n_B}{n} E_{q(f, \mu)} [\log p(\mathcal{D}, B | f, \mu)] \\
&= E_{q(f, \mu, B)} [\log p(\mathcal{D}, B | f, \mu)]
\end{aligned}$$

where (a) is because the finite values of KL terms are divided by infinitely large n , and (b) is due to i.i.d. (\mathcal{D}, B_i) and the variational posterior B being independent of f and μ . Finally, we plug the above inequality into Eqn.(8) and obtain the TELBO. \square

New Optimization Schema for VBHP To optimize the model parameters, we employ the expectation-maximization algorithm. Specifically, in the **E step**, all q_{ij} are optimized to maximize the CELBO, and in the **M step**, m , S , k and c are updated to increase the CELBO. We don't use the TELBO to optimize the variational distributions because it doesn't guarantee minimizing the KL divergence between variational and true posterior distributions. Instead, the TELBO is employed to select GP hyper-parameters:

$$\{\alpha_i^*\}_{i=1}^R, \gamma^* = \operatorname{argmax}_{\{\alpha_i\}_{i=1}^R, \gamma} \text{TELBO}.$$

The TELBO bounds the marginal likelihood more tightly than CELBO, and is therefore expected to lead to a better predictive performance — an intuition which we empirically validate in Sec.6.

The updating equations for q_{ij} are derived through maximization of Eqn.(7) under the constraints $\sum_{j=0}^{i-1} q_{ij} = 1$ for all i . This maximization problem is dealt with the Lagrange multiplier method, and yields the below updating equations:

$$q_{ij} = \begin{cases} \exp(\mathbb{E}_{q(f)}(\log f_{ij}^2))/A_i, & j > 0; \\ \theta \exp(\psi(k))/A_i, & j = 0, \end{cases}$$

where $A_i = \theta \exp(\psi(k)) + \sum_{j=1}^{i-1} \exp(\mathbb{E}_{q(f)}(\log f_{ij}^2))$ is the normalizer.

Furthermore, and similarly to VBPP, we fix the inducing points on a regular grid over \mathcal{T} . Despite the observation that more inducing points lead to better fitting accuracy (Lloyd et al. 2015; Snelson and Ghahramani 2006), in the case of our more complex VBHP, more inducing points may cause slow convergence (Fig.5a (App. 2019)) for some hyper-parameters, and therefore lead to poor performance in limited iterations. Generally, more inducing points improve accuracy at the expense of longer fitting time.

5 Acceleration Trick

Time Complexity Without Acceleration In the E step of model optimization, updating q_{ij} requires computing the mean and the variance of all f_{ij} , which both take $O(M^3 + M^2N^2)$ with N points in the HP and M inducing points. Here, we omit the dimension of data R since normally $M > R$ for a regular grid of inducing points. Similarly, in the M step, computing the hyper-geometric term requires the means and variances of all the f_{ij} . Finally, computation of the integral terms takes $O(M^3N)$. Thus, the total time complexity per iteration is $O(M^3N + M^2N^2)$.

Acceleration to Linear Time Complexity To accelerate our VBHP, similarly to Zhang et al. (2019) we exploit the stationarity of the triggering kernel, assuming the kernel has negligible values for sufficiently large inputs. As a result, sufficiently distant pairs of points do not enter into the computations. This trick reduces possible parents of a point from all prior points to a set of neighbors. The number of relevant neighbors is bounded by a constant C and as a result the total time complexity is reduced to $O(CM^3N)$.

Specifically, we introduce a compact region $\mathcal{S} = \times_{r=1}^R [\mathcal{S}_r^{\min}, \mathcal{S}_r^{\max}] \subseteq \mathcal{T}$ so that $\phi(\mathbf{x}_i - \mathbf{x}_j) = 0$ and $q_{ij} = 0$ if $\mathbf{x}_i - \mathbf{x}_j \notin \mathcal{S}$. As a result, all terms related to $\mathbf{x}_i - \mathbf{x}_j \notin \mathcal{S}$ vanish. To choose a suitable \mathcal{S} , we again use the TELBO, taking the smallest \mathcal{S} for which the TELBO doesn't drop significantly; we optimize \mathcal{S}_r^{\min} and \mathcal{S}_r^{\max} by grid search with other dimensions fixed (so that this step is run R times in total) and we optimize \mathcal{S}_r^{\min} after optimizing \mathcal{S}_r^{\max} .

Rather than selecting pairs of points in each iteration in the manner of Halpin's trick (Halpin 2012; Zhang et al. 2019), our method pre-computes those pairs, leading to gains in computational efficiency. The similar aspect is that both tricks have hyper-parameters to select to threshold the triggering kernel value. We employ the TELBO for hyper-parameter selection while frequentist methods use the cross validation.

6 Experiments

Evaluation. We employ two metrics: the first is the L_2 distance (for cases with a known ground truth), which measures the difference between predictive and truth Hawkes kernels, formulated as $L_2(\phi_{\text{pred}}, \phi_{\text{true}}) = (\int_{\mathcal{T}} (\phi_{\text{pred}}(\mathbf{x}) - \phi_{\text{true}}(\mathbf{x}))^2 d\mathbf{x})^{0.5}$ and $L_2(\mu_{\text{pred}}, \mu_{\text{true}}) = |\mu_{\text{pred}} - \mu_{\text{true}}|$; the second is the hold-out log likelihood (**HLL**), which describes how well the predictive model fits the test data, formulated as $\log p(\mathcal{D}_{\text{Test}} = \{\mathbf{x}_i\}_{i=1}^N | \mu, f) = \sum_{i=1}^N \log \lambda(\mathbf{x}_i) - \int_{\mathcal{T}} \lambda$. To calculate the HLL for each process, we generate a number of test sequences by every time randomly assigning each point of the original process to either a training or testing sequence with equal probability; HLLs of test sequences are normalized (by dividing test sequence length) and averaged.

Prediction. We use the pointwise mode of the approximate posterior triggering kernel as the prediction because it is computationally intractable to find the posterior mode at multiple point locations (Zhang et al. 2019). Besides, we exploit the mode of the approximate posterior background intensity as the predictive background intensity.

Baselines. We use the following models as baselines. **(1)** A parametric Hawkes process equipped with the sum of exponential (**SumExp**) triggering kernel $\phi(\mathbf{x}) = \sum_{i=1}^K a_1^i a_2^i \exp(-a_2^i \mathbf{x})$ and the constant background intensity. **(2)** The ordinary differential equation (**ODE**) based non-parametric non-Bayesian Hawkes process (Zhou, Zha, and Song 2013). The code is publicly available (Bacry et al. 2017). **(3)** Wiener-Hopf (**WH**) equation based non-parametric non-Bayesian Hawkes process (Zhou, Zha, and Song 2013). The code is publicly available (Bacry et al. 2017). **(4)** The Gibbs sampling based Bayesian non-parametric Hawkes process (**Gibbs Hawkes**) (Zhang et al. 2019). For fairness, the ARD kernel is used by Gibbs Hawkes and corresponding eigenfunc-




Figure 2: The Relationship between the Log Marginal Likelihood and the L_2 Distance. In (a), the true ϕ_{sin} (dash green) is plotted with the median (solid) and the $[0.1, 0.9]$ interval (filled) of the approximate posterior triggering kernel obtained by VBHP and Gibbs Hawkes (10 inducing points). It uses the maximum point of the TELBO (red star in (b)). In (c), the maximum point of the TELBO is marked. The maximum point overlaps with that of the CELBO. $[0, 1.4]$ is used as the support of the predictive triggering kernel and 10 inducing points are used.

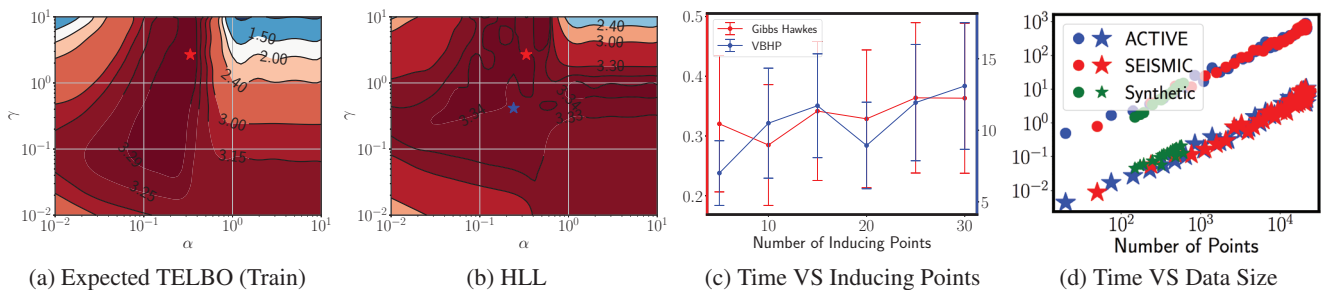


Figure 3: (a)~(b) The Relationship between the TELBO and the HLL; (c)~(d) Average Fitting Time (Seconds) Per Iteration. In (a), the maximum point is marked by the red star. In (b), the maximum points of the TELBO and CELBO are marked by red and blue stars. (c) is plotted on 50 processes. (d) shows the fitting time of Gibbs Hawkes (star) and VBHP (circle) on 120 processes. 10 inducing points are used unless specified.

tions are approximated by Nyström method (Williams and Seeger 2001), where regular grid points are used as VBHP. Different from batch training in (Zhang et al. 2019), all experiments are conducted on single sequences.

6.1 Synthetic Experiments

Synthetic Data. Our synthetic data are generated from three Hawkes processes over $\mathcal{T} = [0, \pi]$, whose triggering kernels are sin, cos and exp functions respectively, shown as below, and whose background intensities are the same $\mu = 10$:

$$\begin{aligned} \phi_{\text{sin}}(x) &= 0.9[\sin(3x) + 1], x \in [0, \pi/2]; \text{ otherwise, } 0; \\ \phi_{\text{cos}}(x) &= \cos(2x) + 1, x \in [0, \pi/2]; \text{ otherwise, } 0; \\ \phi_{\text{exp}}(x) &= 5 \exp(-5x), x \in [0, \infty). \end{aligned}$$

As a result, for any generated sequence, say $\{x_i\}_{i=1}^N$, $\mathcal{T}_i = [0, \pi - x_i]$ is used in the CELBO and the TELBO.

Model Selection. As the marginal likelihood $p(\mathcal{D}|\theta)$ is a key advantage of our method over non-Bayesian approaches (Zhou, Zha, and Song 2013; Bacry and Muzy 2014), we investigate its efficacy for model selection. Fig.2b shows the contour plot of the approximate log marginal likelihood

(the TELBO) of a sequence. It is observed that the contour plot of the TELBO has agreement to the contour plots of $L_2(\phi)$ (Fig.2c) — GP hyper-parameters with relatively high marginal likelihoods have relatively low L_2 errors. Fig.2a plots the posterior triggering kernel corresponding to the maximum approximate marginal likelihood. Similar agreement is also observed between the TELBO and the HLL (Fig.3a, 3b). This demonstrates the practical utility of both the marginal likelihood itself and our approximation of it.

Evaluation. To evaluate VBHP on synthetic data, 20 sequences are drawn from each model and 100 pairs of train and test sequences drawn from each sample to compute the HLL. We select GP hyper-parameters of Gibbs Hawkes and of VBHP by maximizing approximate marginal likelihoods. Table 2 shows evaluations for baselines and VBHP (using 10 inducing points for trade-off between accuracy and time, so does Gibbs Hawkes) in both L_2 and HLL. VBHP achieves the best performance but is two orders of magnitude slower than Gibbs Hawkes per iteration (shown as Fig.3c and 3d). The TELBO performs closely to the CELBO in the L_2 error and this is also reflected in Fig.2c where the maximum

Measure	Data	SumExp	ODE	WH	Gibbs Hawkes	VBHP (C)	VBHP (T)
L ₂	Sin	$\phi:0.693_{\pm 0.028}$	$0.665_{\pm 0.121}$	$2.463_{\pm 0.145}$	$0.408_{\pm 0.198}$	0.152 $_{\pm 0.091}$	$0.183_{\pm 0.076}$
		$\mu:2.968_{\pm 1.640}$	$4.514_{\pm 3.808}$	$6.794_{\pm 5.054}$	$4.108_{\pm 3.949}$	$0.640_{\pm 0.528}$	0.579 $_{\pm 0.523}$
	Cos	$\phi:0.473_{\pm 0.102}$	$0.697_{\pm 0.065}$	$1.743_{\pm 0.083}$	$0.667_{\pm 0.686}$	$0.325_{\pm 0.073}$	0.292 $_{\pm 0.096}$
		$\mu:2.751_{\pm 1.902}$	$7.030_{\pm 5.662}$	$6.099_{\pm 4.613}$	$4.685_{\pm 4.421}$	$0.555_{\pm 0.294}$	0.515 $_{\pm 0.293}$
	Exp	$\phi:$ 0.133 $_{\pm 0.138}$	$1.835_{\pm 0.539}$	$2.254_{\pm 2.042}$	$0.676_{\pm 0.233}$	$0.257_{\pm 0.086}$	$0.235_{\pm 0.102}$
		$\mu:3.290_{\pm 1.991}$	$8.969_{\pm 8.604}$	$16.66_{\pm 20.95}$	$7.648_{\pm 9.647}$	0.471 $_{\pm 0.432}$	$0.486_{\pm 0.418}$
HLL	Sin	$3.490_{\pm 0.400}$	$3.489_{\pm 0.413}$	$3.233_{\pm 0.273}$	$3.492_{\pm 0.406}$	$3.488_{\pm 0.400}$	3.497 $_{\pm 0.406}$
	Cos	$3.874_{\pm 0.544}$	$3.872_{\pm 0.552}$	$3.613_{\pm 0.373}$	$3.871_{\pm 0.562}$	$3.876_{\pm 0.541}$	3.878 $_{\pm 0.548}$
	Exp	$2.825_{\pm 0.481}$	$2.822_{\pm 0.496}$	$2.782_{\pm 0.490}$	$2.826_{\pm 0.492}$	$2.826_{\pm 0.491}$	2.829 $_{\pm 0.487}$
	ACTIVE	$1.692_{\pm 1.371}$	$0.880_{\pm 2.716}$	$0.710_{\pm 0.943}$	$1.323_{\pm 2.160}$	$1.824_{\pm 1.159}$	1.867 $_{\pm 1.181}$
	SEISMIC	$2.943_{\pm 0.959}$	$2.582_{\pm 1.665}$	$1.489_{\pm 1.796}$	$3.110_{\pm 1.251}$	$3.143_{\pm 0.895}$	3.164 $_{\pm 0.843}$

Table 2: Results on Synthetic and Real World Data (Mean \pm One Standard Variance). VBHP (C) and (T) use the CELBO and the TELBO to update the hyper-parameters respectively.

points of the TELBO and the CELBO overlap. In contrast, the TELBO consistently improves the performance of VBHP in the HLL, which is also reflected in Fig.3b where hyper-parameters selected by the TELBO tend to have a higher HLL. Interestingly, when the parametric model SumExp uses the same triggering kernel (a single exponential function) as the ground truth ϕ_{exp} , SumExp fits ϕ_{exp} best in L₂ distance while due to learning on single sequences, the background intensity has relatively high errors. Although our method is not aware of the parametric family of the ground truth, it performs well. Compared with non-parametric frequentist methods which have strong fitting capacity but suffer from noisy data and have difficulties with hyper-parameter selection, our Bayesian solution overcomes these disadvantages and achieves better performance.

6.2 Real World Experiments

Real World Data. We conclude our experiments with two large scale tweet datasets. **ACTIVE** (Rizoiu et al. 2018) is a tweet dataset which was collected in 2014 and contains $\sim 41k$ (re)tweet temporal point processes with links to Youtube videos. Each sequence contains at least 20 (re)tweets. **SEISMIC** (Zhao et al. 2015) is a large scale tweet dataset which was collected from October 7 to November 7, 2011, and contains $\sim 166k$ (re)tweet temporal point processes. Each sequence contains at least 50 (re)tweets.

Evaluation. Similarly to synthetic experiments, we evaluate the fitting performance by averaging HLL of 20 test sequences randomly drawn from each original datum. We scale all original data to $\mathcal{T} = [0, \pi]$ (leading to $\mathcal{T}_i = [0, \pi - x_i]$ used in the CELBO and the TELBO for a sequence $\{x_i\}_{i=1}^N$) and employ 10 inducing points to balance time and accuracy. The model selection is performed by maximizing the approximate marginal likelihood. The obtained results are shown in Table 2. Again, we observe similar predictive performance of VBHP: the TELBO performs better the CELBO; VBHP achieves best scores. This demonstrates our Bayesian model and novel VI schema are useful for flexible real life data.

Fitting Time. We further evaluate the fitting speed² of

VBHP and Gibbs Hawkes on synthetic and real world point processes, which is summarized in Fig.3c and 3d. The fitting time is averaged over iterations and we observe that the increasing trends with the number of inducing points and with the data size are similar between Gibbs Hawkes and VBHP. Although VBHP is significantly slower than Gibbs Hawkes per iteration, VBHP converges faster, in 10~20 iterations (Fig.5 (App. 2019)), giving an average convergence time of 549 seconds for a sequence of 1000 events, compared to 699 seconds for Gibbs Hawkes. The slope of VBHP in Fig.3d is 1.04 (log-scale) and the correlation coefficient is 0.96, so we conclude that the fitting time is linear to the data size.

7 Conclusions

In this paper, we presented a new Bayesian non-parametric Hawkes process whose triggering kernel is modulated by a sparse Gaussian process and background intensity is Gamma distributed. We provided a novel VI schema for such a model: we employed the branching structure so that the common ELBO is maximized by the expectation-maximization algorithm; we contributed a tighter ELBO which performs better in model selection than the common one. To address the difficulty with scaling with the data size, we utilize the stationarity of the triggering kernel to reduce the number of possible parents for each point. Different from prior acceleration methods, ours enjoys higher efficiency. On synthetic data and two large Twitter diffusion datasets, VBHP enjoys linear fitting time with the data size and fast convergence rate, and provides more accurate predictions than those of state-of-the-art approaches. The novel ELBO is also demonstrated to exceed the common one in model selection.

8 Acknowledgements

We would like to thank Dongwoo Kim for helpful discussions. Rui is supported by the Australian National University and Data61 CSIRO PhD scholarships.

²The CPU we use is Intel(R) Core(TM) i7-7700 CPU @

3.60GHz and the language is Python 3.6.5.

References

- Ancarani, L. U., and Gasaneo, G. 2008. Derivatives of any order of the confluent hypergeometric function $f_1^1(a, b, z)$ with respect to the parameter a or b . *Journal of Mathematical Physics* 49(6):063508.
- App. 2019. Appendix: Variational inference for sparse gaussian process modulated hawkes process. In the supplementary material.
- Attias, H. 1999. Inferring parameters and structure of latent variable models by variational bayes. In *Proceedings of the Fifteenth conference on Uncertainty in artificial intelligence*, 21–30. Morgan Kaufmann Publishers Inc.
- Bacry, E., and Muzy, J.-F. 2014. Second order statistics characterization of hawkes processes and non-parametric estimation. *arXiv preprint arXiv:1401.0903*.
- Bacry, E.; Bompain, M.; Gaïffas, S.; and Poulsen, S. 2017. tick: a Python library for statistical learning, with a particular emphasis on time-dependent modeling. *ArXiv e-prints*.
- Donnet, S.; Rivoirard, V.; and Rousseau, J. 2018. Nonparametric bayesian estimation of multivariate hawkes processes. *arXiv preprint arXiv:1802.05975*.
- Eichler, M.; Dahlhaus, R.; and Dueck, J. 2017. Graphical modeling for multivariate hawkes processes with non-parametric link functions. *Journal of Time Series Analysis* 38(2):225–242.
- Halpin, P. F. 2012. An em algorithm for hawkes process. *Psychometrika* 2.
- Hawkes, A. G., and Oakes, D. 1974. A cluster process representation of a self-exciting process. *Journal of Applied Probability* 11(3):493–503.
- Hawkes, A. G. 1971. Spectra of some self-exciting and mutually exciting point processes. *Biometrika* 58(1):83–90.
- Jordan, M. I.; Ghahramani, Z.; Jaakkola, T. S.; and Saul, L. K. 1999. An introduction to variational methods for graphical models. *Machine learning* 37(2):183–233.
- Lewis, E., and Mohler, G. 2011. A nonparametric em algorithm for multiscale hawkes processes. *Journal of Nonparametric Statistics* 1(1):1–20.
- Linderman, S., and Adams, R. 2014. Discovering latent network structure in point process data. In *International Conference on Machine Learning*, 1413–1421.
- Linderman, S. W., and Adams, R. P. 2015. Scalable bayesian inference for excitatory point process networks. *arXiv preprint arXiv:1507.03228*.
- Lloyd, C.; Gunter, T.; Osborne, M.; and Roberts, S. 2015. Variational inference for gaussian process modulated poisson processes. In *ICML*, 1814–1822.
- Rasmussen, J. G. 2013. Bayesian inference for hawkes processes. *Methodology and Computing in Applied Probability* 15(3):623–642.
- Rizoiu, M.-A.; Mishra, S.; Kong, Q.; Carman, M.; and Xie, L. 2018. Sir-hawkes: Linking epidemic models and hawkes processes to model diffusions in finite populations. In *Proceedings of the 27th International Conference on World Wide Web*, 419–428. International World Wide Web Conferences Steering Committee.
- Snelson, E., and Ghahramani, Z. 2006. Sparse gaussian processes using pseudo-inputs. In *Advances in neural information processing systems*, 1257–1264.
- Williams, C. K. I., and Seeger, M. 2001. Using the nyström method to speed up kernel machines. In Leen, T. K.; Dietterich, T. G.; and Tresp, V., eds., *Advances in Neural Information Processing Systems 13*. MIT Press. 682–688.
- Zhang, R.; Walder, C.; Rizoiu, M.-A.; and Xie, L. 2019. Efficient Non-parametric Bayesian Hawkes Processes. In *International Joint Conference on Artificial Intelligence (IJCAI'19)*.
- Zhao, Q.; Erdogdu, M. A.; He, H. Y.; Rajaraman, A.; and Leskovec, J. 2015. Seismic: A self-exciting point process model for predicting tweet popularity. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1513–1522. ACM.
- Zhou, K.; Zha, H.; and Song, L. 2013. Learning triggering kernels for multi-dimensional hawkes processes. In *International Conference on Machine Learning*, 1301–1309.