
Large-Scale Cox Process Inference using Variational Fourier Features

ST John¹ James Hensman¹

Abstract

Gaussian process modulated Poisson processes provide a flexible framework for modeling spatiotemporal point patterns. So far this was restricted to one dimension, binning to a predetermined grid, or small data sets of up to a few thousand data points. Here we introduce Cox process inference based on Fourier features. This sparse representation induces global rather than local constraints on the function space and is computationally efficient. This allows us to formulate a grid-free approximation that scales well with the number of data points and the size of the domain. We demonstrate that this allows MCMC approximations to the non-Gaussian posterior. In practice, we find that Fourier features have more consistent optimization behavior than previous approaches. Our approximate Bayesian method can fit over 100 000 events with complex spatiotemporal patterns in three dimensions on a single GPU.

1. Introduction

Modeling spatiotemporal point patterns is a common task in geostatistics, for example in ecology and epidemiology (Diggle et al., 2013; Vanhatalo & Vehtari, 2007). We are further motivated by modeling events occurring in a city, e.g., taxi pickups or crime incidents (Flaxman et al., 2018).

Gaussian process (GP) modulated Poisson processes provide a flexible Bayesian model for such data. The model includes a GP prior over a latent function $f(\cdot)$, which is related to the rate $\lambda(\cdot)$ of an inhomogeneous Poisson process through a link function. This is usually taken to be $f(\cdot) = \log(\lambda(\cdot))$, resulting in the Log Gaussian Cox process (Møller et al., 1998). These models are computationally challenging as they are doubly intractable (Murray, 2007). The likelihood involves an integral of the process over the

spatiotemporal domain, which cannot be computed in general. There are three potential remedies to this issue. First, the classic approach is to discretize the input domain (see, e.g., Taylor et al., 2015), assuming that the rate is constant over each grid cell. However, this discretized structure scales poorly with the number of grid cells. Second, we can use a thinning strategy to construct an exact MCMC sampler (Adams et al., 2009; Gunter et al., 2014), which demands the inverse link function to be bounded, e.g., sigmoidal. However, this is prohibitively expensive for more than a few thousand points. Finally, we can use a square-root link function $\lambda(\cdot) = f(\cdot)^2$, as proposed by Lloyd et al. (2015). Combined with variational inference for the latent function, this can make the integral tractable. Although their approach scales linearly with the number of events, it scales poorly with the size of the domain, especially if the rate has high variance. Flaxman et al. (2017) make use of the same link function with a frequentist approach, and Walder & Bishop (2017) use a Laplace approximation.

In this work we build on Lloyd et al. (2015)’s proposal. We similarly make use of the square-root link function for tractability, but introduce several innovations that greatly improve the practical applicability of the method.

First, we extend the derivation to be able to use the Fourier representation of the GP proposed by Hensman et al. (2018). This allows for a choice of kernels in the Matérn family. Additionally, we show how the posterior process can be approximated using standard MCMC methods (Hensman et al., 2015). This allows for flexible representations of the posterior over functions $f(\cdot)$, and a Bayesian treatment of the hyperparameters, whereas Lloyd et al. (2015)’s approach was restricted to a Gaussian approximation for $f(\cdot)$ and point estimates of the hyperparameters.

Second, we deliver insights and improvements to the model. The use of a square-root link function can result in *nodal lines*: where the latent function $f(\cdot)$ crosses zero, the rate function $\lambda(\cdot)$ must approach zero on both sides, resulting in regions of the posterior that have unsatisfactory artifacts. In addition, these nodal lines create multiple modes in the posterior that are hard to sample and difficult to approximate well with a Gaussian approach. We investigate these effects and demonstrate how they can be mitigated by suitable specification of the prior parameters.

¹PROWLER.io, 66-68 Hills Road, Cambridge CB2 1LA, United Kingdom. Correspondence to: ST John <st@prowler.io>, James Hensman <james@prowler.io>.

Third, motivated by spatiotemporal problems where a periodic function is expected *a priori*, we construct a periodic kernel that can be combined with the Fourier features approach to Matérn kernels under our inference scheme. We also show how to obtain a closed-form confidence interval for the rate λ under a Gaussian approximation.

The key result of our work is that we can now scale Bayesian point process inference on a single GPU to hundreds of thousands of events from multiple observations with complex spatiotemporal patterns. We demonstrate this on the Porto taxi data set (Moreira-Matias et al., 2013).

2. Cox Process Inference

A Cox process is an inhomogeneous Poisson process, where the rate λ is itself a stochastic process. The probability of the data $\mathcal{D} = \{\mathbf{x}_n\}_{n=1}^N$ under an inhomogeneous Poisson process with known rate function $\lambda(\cdot)$ is given by

$$p(\mathcal{D} | \lambda(\cdot)) = \exp\left(-\int_{\mathcal{T}} \lambda(\mathbf{x}) d\mathbf{x}\right) \prod_{n=1}^N \frac{\lambda(\mathbf{x}_n)^{n_{x_n}}}{n_{x_n}!}, \quad (1)$$

where \mathcal{T} is the domain of the observation and n_{x_n} is the multiplicity of events at location \mathbf{x}_n . In the following, we assume that all events are distinct and $n_{x_n} = 1$.

We want to infer the posterior distribution of the rate function given an observation \mathcal{D} , which is given by $p(\lambda(\cdot) | \mathcal{D}) = \frac{p(\lambda(\cdot))p(\mathcal{D} | \lambda(\cdot))}{\int p(\lambda(\cdot))p(\mathcal{D} | \lambda(\cdot)) d\lambda(\cdot)}$. The marginal likelihood in the denominator involves a double integral, which gives rise to the so-called ‘‘double intractability’’ (Murray, 2007).

2.1. Multiple Observations

We can consider multiple observations corresponding to several draws from the same distribution. For example, when considering the rate of taxi pickups in a city, we can take into account observations $\mathcal{D}_o = \{\mathbf{x}_n^{(o)}\}_{n=1}^{N_o}$ from several days, $o = 1, \dots, O$, each containing N_o events. Instead of modeling each day independently, we model the pickup rate of the average day, while making use of all available data. Assuming that observations are i.i.d. with the same rate function $\lambda(\cdot)$, the likelihood factorizes across observations,

$$p(\mathcal{D}_1, \dots, \mathcal{D}_O | \lambda(\cdot)) = \prod_{o=1}^O p(\mathcal{D}_o | \lambda(\cdot)), \quad (2)$$

where $p(\mathcal{D}_o | \lambda(\cdot))$ is given by eq. (1).

2.2. Cox Process Constructions with Gaussian Processes

We model the intensity function $\lambda(\cdot)$ using a GP, $f(\cdot) \sim \mathcal{GP}$. While $f(\cdot)$ is unbounded between $-\infty$ and $+\infty$, the rate $\lambda(\cdot)$ must be non-negative everywhere. This is achieved by an inverse link function $\rho(\cdot)$ that ensures $\lambda(\cdot) = \rho(f(\cdot)) \geq$

0 for any value of f . Lloyd et al. (2015) used $\lambda(\cdot) = f(\cdot)^2$, which has the advantage of tractable analytic derivations and, with a sparse GP approximation, linear scaling in the number of observed events. However, the lack of injectivity of f^2 can lead to nodal lines (see section 4.5). To mitigate this disadvantage, we extend the link function to include a constant offset¹ β , so that $\lambda(\cdot) = (f(\cdot) + \beta)^2$. Defining Θ to contain β and the GP hyperparameters, our model is

$$p(\mathcal{D}, f(\cdot), \Theta) = p(\mathcal{D} | \lambda = (f(\cdot) + \beta)^2) p(f(\cdot) | \Theta) p(\Theta). \quad (3)$$

For tractable inference we approximate the exact GP posterior $p(f(\cdot), \Theta | \mathcal{D})$ by $q(f(\cdot) | \Theta) q(\Theta)$, where the approximate GP posterior is $q(f(\cdot) | \Theta) = \mathcal{GP}(f; \tilde{\mu}(\cdot), \tilde{\Sigma}(\cdot, \cdot))$. Then $\mathbb{E}_{q(f(\cdot))}[\lambda(\cdot)] = \tilde{\mu}(\cdot)^2 + \tilde{\Sigma}(\cdot)$ is the expected rate under the posterior. We now derive the inference objective, and later discuss our choices for $q(f(\cdot))$.

2.3. Inference Objective

We rederive the Evidence Lower Bound Objective (ELBO) from Lloyd et al. (2015) to account for our extended model specification. We optimize $q(f(\cdot))$ by minimizing the KL divergence to the true posterior, where we leave the dependence on the hyperparameters implicit:²

$$\begin{aligned} \mathcal{K} &= \text{KL}[q(f(\cdot)) \| p(f(\cdot) | \mathcal{D})] \\ &= -\mathbb{E}_{q(f(\cdot))} \left[\log \frac{p(f(\cdot), \mathcal{D}) / p(\mathcal{D})}{q(f(\cdot))} \right] \\ &= -\mathbb{E}_q \left[\log \frac{p(f(\cdot))}{q(f(\cdot))} \right] - \mathbb{E}_q [\log p(\mathcal{D} | f(\cdot))] + \log p(\mathcal{D}) \\ &= \text{KL}[q(f(\cdot)) \| p(f(\cdot))] - \mathcal{L}_D + \log p(\mathcal{D}) \\ &= \log p(\mathcal{D}) - \mathcal{L}, \end{aligned} \quad (4)$$

where \mathcal{L} is the ELBO from Lloyd et al. (2015). Minimizing \mathcal{K} with respect to a variational distribution $q(f(\cdot))$ is equivalent to maximizing the ELBO \mathcal{L} . We use the result from Matthews et al. (2016) that the KL between the approximate posterior and prior processes is a KL divergence at the inducing points. For details, see Matthews (2016).

For our Cox process model, the likelihood term \mathcal{L}_D is

$$\begin{aligned} \mathcal{L}_D &= \mathbb{E}_{q(f(\cdot))} [\log p(\{\mathcal{D}_o\}_{o=1}^O | f(\cdot))] \\ &= -O \mathbb{E}_{q(f(\cdot))} \left[\int_{\mathcal{T}} \lambda(\mathbf{x}) d\mathbf{x} \right] \\ &\quad + \sum_{o=1}^O \sum_{n=1}^{N_o} \mathbb{E}_{q(f(\cdot))} [\log \lambda(\mathbf{x}_n^{(o)})] \\ &= -O \mathcal{L}_{f_x} + \mathcal{L}_{f_n}, \end{aligned} \quad (5)$$

¹This is equivalent to a prior mean on the GP. Our method could also be extended to include covariates whose coefficients can be learnt within our inference frameworks.

²The full derivation is given in the Supplementary Material.

where we allow for multiple observations, and the sum in \mathcal{L}_{f_n} extends over the $N = \sum_{o=1}^O N_o$ events in all observations. We now consider these terms separately for our choice of link function, $\lambda(\mathbf{x}) = (f(\mathbf{x}) + \beta)^2$.

The data term \mathcal{L}_{f_n} aims to increase the rate λ at the event locations:

$$\mathcal{L}_{f_n} = \sum_{n=1}^N \mathbb{E}_{q(f(\cdot))} [\log[(f(\mathbf{x}_n) + \beta)^2]]. \quad (6)$$

Each summand in the \mathcal{L}_{f_n} term corresponds to an integral $\int_{-\infty}^{\infty} \log[(f_n + \beta)^2] \mathcal{N}(f_n; \tilde{\mu}(\mathbf{x}_n), \tilde{\Sigma}(\mathbf{x}_n, \mathbf{x}_n)) df_n$, where $f_n = f(\mathbf{x}_n)$. With a change of variables $f = f_n + \beta$, this becomes equivalent to the one-dimensional integral that Lloyd et al. (2015) showed can be evaluated in closed form: $\int_{-\infty}^{\infty} \log(f^2) \mathcal{N}(f; \tilde{\mu}, \tilde{\sigma}^2) df = -\tilde{G}(-\frac{\tilde{\mu}^2}{2\tilde{\sigma}^2}) + \log(\frac{\tilde{\sigma}^2}{2}) - C$, where $\tilde{G}(\cdot)$ can be represented by a lookup table.

The area term \mathcal{L}_{f_x} aims to minimize the overall rate, which ensures that λ becomes small where there are fewer events:

$$\begin{aligned} \mathcal{L}_{f_x} &= \mathbb{E}_{q(f(\cdot))} \left[\int_{\mathcal{T}} (f(\mathbf{x}) + \beta)^2 d\mathbf{x} \right] \\ &= \int_{\mathcal{T}} \mathbb{E}_q[f(\mathbf{x})^2] d\mathbf{x} + 2\beta \int_{\mathcal{T}} \mathbb{E}_q[f(\mathbf{x})] d\mathbf{x} + \beta^2 |\mathcal{T}|. \end{aligned} \quad (7)$$

These integrals over the moments of $f(\cdot)$ are similar to the ‘‘ Ψ statistics’’ that show up in GPs with uncertain input (Titsias & Lawrence, 2010), but over a uniform rather than Gaussian distribution. They depend on the form of $q(f(\cdot))$, which we discuss in the following for a Gaussian variational approximation and variationally sparse MCMC.

3. Approximations

3.1. Variational Gaussian Approximation

To keep the GP tractable, we use a sparse approximation to the full posterior. We consider a set of inducing points $Z = \{\mathbf{z}_m\}_{m=1}^M$, and we collect the inducing variables $u_m = f(\mathbf{z}_m)$ in a vector $\mathbf{u} = f(Z)$. Our approximation matches the prior conditioned on the values at the inducing points: $q(f(\cdot) | \mathbf{u}) = p(f(\cdot) | \mathbf{u})$. We define $\mathbf{k}_u(\cdot) = \text{cov}(\mathbf{u}, f(\cdot)) = k(Z, \cdot)$ and $\mathbf{K}_{\mathbf{u}\mathbf{u}} = \text{cov}(\mathbf{u}, \mathbf{u}) = k(Z, Z)$.

Assuming a Gaussian form for the variational approximate distribution, $q(\mathbf{u}) = \mathcal{N}(\mathbf{m}, \mathbf{S})$, then $q(f(\cdot)) = \int p(f(\cdot) | \mathbf{u}) q(\mathbf{u}) d\mathbf{u} = \mathcal{GP}(f; \tilde{\mu}(\cdot), \tilde{\Sigma}(\cdot, \cdot))$, where

$$\begin{aligned} \tilde{\mu}(\cdot) &= \mathbf{k}_u(\cdot)^\top \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1} \mathbf{m}, \\ \tilde{\Sigma}(\cdot, \cdot) &= k(\cdot, \cdot) - \mathbf{k}_u(\cdot)^\top (\mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1} - \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1} \mathbf{S} \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1}) \mathbf{k}_u(\cdot). \end{aligned} \quad (8)$$

Similar to the derivation by Lloyd et al. (2015), we have

$$\begin{aligned} \mathcal{L}_{f_x} &= \mathbf{m}^\top \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1} \Psi \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1} \mathbf{m} + \sigma^2 |\mathcal{T}| - \text{tr}(\mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1} \Psi) \\ &\quad + \text{tr}(\mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1} \mathbf{S} \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1} \Psi) + 2\beta \Phi^\top \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1} \mathbf{m} + \beta^2 |\mathcal{T}|, \end{aligned} \quad (9)$$

where $\sigma^2 = k(x, x)$ is the variance of the kernel (assuming stationarity), $\Psi = \int_{\mathcal{T}} \mathbf{k}_u(\mathbf{x}) \mathbf{k}_u(\mathbf{x})^\top d\mathbf{x}$, and $\Phi = \int_{\mathcal{T}} \mathbf{k}_u(\mathbf{x}) d\mathbf{x}$. Note that the terms involving β in eq. (9) are missing in Lloyd et al. (2015). The expectation integrals Ψ and Φ for Fourier features will be derived in section 3.4. We can now optimize \mathcal{L} with respect to \mathbf{m} , \mathbf{S} , and Θ , which takes $\mathcal{O}(NM^2)$ operations per gradient step.

3.2. Sparse MCMC

Though the variational Gaussian approximation often works well in practice, in some cases it may be important to determine the posterior distribution of the hyperparameters instead of using point estimates, and not restrict the form of $q(f(\cdot))$ to be Gaussian. Hensman et al. (2015) introduced MCMC for variationally sparse GPs. This approach still relies on a sparse set of inducing features to describe the function, but does not restrict $q(\mathbf{u})$ to be Gaussian. They demonstrate that the optimal variational distribution is

$$\log \hat{q}(\mathbf{u}, \Theta) = \mathcal{L}'_D + \log p(\mathbf{u} | \Theta) + \log p(\Theta) - \log Z, \quad (10)$$

where $\mathcal{L}'_D = \mathbb{E}_{p(f(\cdot) | \mathbf{u}, \Theta)} [\log p(\mathcal{D} | f(\cdot))]$, and the constant Z normalizes the distribution. We can sample from this distribution using MCMC. Here, we use Hamiltonian MC.

In each step, we need to evaluate $\log \hat{q}(\mathbf{u}, \Theta)$ up to a constant: $\log \pi(\mathbf{u}, \Theta) = \mathcal{L}'_D + \log p(\mathbf{u} | \Theta) + \log p(\Theta)$, where the expectation in \mathcal{L}'_D is under the conditional distribution

$$\begin{aligned} p(f(\cdot) | \mathbf{u}, \Theta) &= \mathcal{GP}(f; \mathbf{k}_u(\cdot)^\top \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1} \mathbf{u}, k(\cdot, \cdot) - \mathbf{k}_u(\cdot)^\top \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1} \mathbf{k}_u(\cdot)). \end{aligned} \quad (11)$$

For our Cox process model, the \mathcal{L}'_{f_x} term is now a function of \mathbf{u} , as opposed to \mathbf{m} and \mathbf{S} in eq. (9):

$$\begin{aligned} \mathcal{L}'_{f_x} &= \mathbf{u}^\top \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1} \Psi \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1} \mathbf{u} + \sigma^2 |\mathcal{T}| - \text{tr}(\mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1} \Psi) \\ &\quad + 2\beta \Phi^\top \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1} \mathbf{u} + \beta^2 |\mathcal{T}|. \end{aligned} \quad (12)$$

The \mathcal{L}'_{f_x} term is similar to eq. (6), with the mean and variance of the Gaussian distribution now defined through eq. (11). This MCMC approach requires $\mathcal{O}(NM)$ computations in each step for the evaluation of $\log \pi(\mathbf{u}, \Theta)$.

3.3. Fourier Features

Point process models need to accurately describe the rate function across the entire domain, not just where events occurred. Not having observed any events in a region is also informative (in that it suggests a lower rate). For this reason, it is beneficial to use an approximation that has support everywhere, instead of at a fixed number of inducing points.

The variational Fourier features (VFF) introduced by Hensman et al. (2018) have support across the entire domain. They are based on a spectral representation of the GP, related to random Fourier features (Rahimi & Recht, 2008), but with

fixed frequencies. The random variables corresponding to these features are given by³ $u_m = \langle \phi_m(\cdot), f(\cdot) \rangle_{\mathcal{H}}$, where

$$\phi(x) = [1, \cos(\omega_1(x-a)), \dots, \cos(\omega_M(x-a)), \sin(\omega_1(x-a)), \dots, \sin(\omega_M(x-a))]^\top. \quad (13)$$

The features are parametrized by the frequencies (we choose $\omega_m = 2\pi m/(b-a)$ for $m = 1, \dots, M$) and the bounding box $[a, b]$. The bounding box is part of the approximation and should be chosen somewhat larger than the domain \mathcal{T} .

This choice of \mathbf{u} leads to $\mathbf{k}_{\mathbf{u}}(\cdot) = \phi(\cdot)$ (Hensman et al., 2018). Note that the features do not depend on the kernel parameters. This allows us to calculate the Ψ matrix very efficiently, as it does not change within each iteration of the optimization (see section 3.4). The covariance between features, $\text{cov}(u_m, u_{m'}) = \langle \phi_m, \phi_{m'} \rangle_{\mathcal{H}}$, turns out to have low rank: $\mathbf{K}_{\mathbf{uu}} = \text{diag}(\boldsymbol{\alpha}) + \mathbf{W}\mathbf{W}^\top$, where the rank of \mathbf{W} depends on the kernel (one for Matérn-1/2, two for Matérn-3/2, and three for Matérn-5/2). This means that we can evaluate $\mathbf{K}_{\mathbf{uu}}^{-1}\mathbf{k}_{\mathbf{u}}(\cdot)$ in $\mathcal{O}(NM)$ operations rather than $\mathcal{O}(NM^2)$.

The Fourier features can only be applied to one-dimensional kernels. However, we can apply this framework to higher-dimensional domains by using additive (sum) and/or separable (product) kernels. For a sum kernel, $k(\mathbf{x}, \mathbf{x}') = \sum_{d=1}^D k_d(x_d, x'_d)$, the feature matrix is given by the stacking of the 1D feature matrices. Features in different dimensions do not interact, and $\mathbf{K}_{\mathbf{uu}}$ is block-diagonal. For a product kernel, $k(\mathbf{x}, \mathbf{x}') = \prod_{d=1}^D k_d(x_d, x'_d)$, the features are given by $\phi(\mathbf{x}) = \otimes_{d=1}^D [\phi_1(x_d), \dots, \phi_M(x_d)]^\top$. Again, this is independent of the kernel parameters. $\mathbf{K}_{\mathbf{uu}}$ has Kronecker structure: $\text{cov}(\mathbf{u}, \mathbf{u}) = \mathbf{K}_{\mathbf{uu}} = \otimes_{d=1}^D \mathbf{K}_{\mathbf{uu}}^d$. We follow Hensman et al. (2018)'s recommendation in using a sum-of-Kronecker structure for the variational covariance \mathbf{S} .

3.4. Ψ Statistics for Fourier Features

To implement Fourier features for the Cox process model discussed in this paper, we also need Ψ and Φ (cf. eq. (9)), which we derive in this section.

For Fourier features, $\mathbf{k}_{\mathbf{u}}(\cdot)$ only consists of cosine and sine functions, independent of the kernel, so Ψ and Φ are just integrals over (products of) cosine functions. As they do not depend on any hyperparameters, we can simply precompute $\mathbf{k}_{\mathbf{u}}(\cdot)$, Ψ , and Φ outside the optimization loop. In 1D, with $\mathcal{T} = [c, d]$, we derive (cf. Supplementary Material)

$$\Psi = \int_c^d \mathbf{k}_{\mathbf{u}}(x)\mathbf{k}_{\mathbf{u}}(x)^\top dx = \begin{bmatrix} |\mathcal{T}| & \Psi_{(1,c)} & \Psi_{(1,s)} \\ \Psi_{(1,c)}^\top & \Psi_{(c,c)} & \Psi_{(c,s)} \\ \Psi_{(1,s)}^\top & \Psi_{(c,s)}^\top & \Psi_{(s,s)} \end{bmatrix}, \quad (14)$$

with $\mathcal{C}_m^x = \cos(\omega_m(x-a))$ and $\mathcal{S}_m^x = \sin(\omega_m(x-a))$ and

$$\Psi_m^{(1,c)} = \omega_m^{-1}[\mathcal{S}_m^d - \mathcal{S}_m^c],$$

³Where $\langle g(\cdot), h(\cdot) \rangle_{\mathcal{H}}$ represents the inner product in the Reproducing Kernel Hilbert Space \mathcal{H} associated with the kernel $k(\cdot, \cdot)$.

$$\begin{aligned} \Psi_m^{(1,s)} &= -\omega_m^{-1}[\mathcal{C}_m^d - \mathcal{C}_m^c], \\ \Psi_{m=n}^{(c,c)} &= |\mathcal{T}|/2 + \omega_{4m}^{-1}(\mathcal{S}_{2m}^d - \mathcal{S}_{2m}^c), \\ \Psi_{m \neq n}^{(c,c)} &= \frac{n(\mathcal{C}_m^d \mathcal{S}_n^d - \mathcal{C}_m^c \mathcal{S}_n^c) - m(\mathcal{S}_m^d \mathcal{C}_n^d - \mathcal{S}_m^c \mathcal{C}_n^c)}{2\pi(n^2 - m^2)/(b-a)}, \\ \Psi_{m=n}^{(s,s)} &= |\mathcal{T}|/2 - \omega_{4m}^{-1}(\mathcal{S}_{2m}^d - \mathcal{S}_{2m}^c), \\ \Psi_{m \neq n}^{(s,s)} &= \frac{m(\mathcal{C}_m^d \mathcal{S}_n^d - \mathcal{C}_m^c \mathcal{S}_n^c) - n(\mathcal{S}_m^d \mathcal{C}_n^d - \mathcal{S}_m^c \mathcal{C}_n^c)}{2\pi(n^2 - m^2)/(b-a)}, \\ \Psi_{m=n}^{(c,s)} &= \omega_{4m}^{-1}[\mathcal{C}_{2m}^d - \mathcal{C}_{2m}^c], \\ \Psi_{m \neq n}^{(c,s)} &= \frac{n(\mathcal{S}_m^d \mathcal{S}_n^d - \mathcal{S}_m^c \mathcal{S}_n^c) + m(\mathcal{C}_m^d \mathcal{C}_n^d - \mathcal{C}_m^c \mathcal{C}_n^c)}{2\pi(n^2 - m^2)/(b-a)}. \end{aligned}$$

For multi-dimensional sum and product kernels, the Ψ matrix can be constructed from the 1D cases. For a product kernel, Ψ is the Kronecker product of the 1D Ψ matrices for each dimension. For a sum kernel, Ψ is block-diagonal; the diagonal blocks are equivalent to the 1D Ψ matrices scaled by the total volume divided by the length of that dimension. The off-diagonal blocks are the outer product of the first rows of the 1D Ψ matrices scaled by the total volume divided by the lengths of the two involved dimensions.

The first Fourier feature is the constant 1, so in 1D the Φ vector is equal to the first row of Ψ . For multiple dimensions using a product kernel, Φ is the Kronecker-vector-stack of the first rows of the component matrices of Ψ . For a sum kernel, Φ in each dimension is equivalent to the corresponding 1D case, scaled by the total volume divided by the length of that dimension.

4. Methods

4.1. Parametric Periodic Kernels

When we have more explicit prior knowledge about the behavior of our data, we can encode this in parametric kernels. For example, in spatiotemporal point processes, we may want to model periodic components such as time-of-day, where we want an explicitly periodic behavior on the time dimension to avoid discontinuities at the ‘‘roll-over point’’.

Given a vector of features $\phi(\cdot)$, we can define a parametric kernel as $k(\cdot, \cdot) = \phi(\cdot)^\top \mathbf{G}^{-1} \phi(\cdot)$ using a positive-definite Gram matrix \mathbf{G} . With $u_m = \langle \phi_m(\cdot), f(\cdot) \rangle_{\mathcal{H}}$, this results in $\mathbf{K}_{\mathbf{uu}} = \mathbf{G}$ and $\mathbf{k}_{\mathbf{u}}(\cdot) = \phi(\cdot)$.⁴

To obtain a periodic kernel, we can use sines and cosines with the right periodicity, i.e., $\sin(\omega_m \cdot)$ and $\cos(\omega_m \cdot)$ with $\omega_m = 2\pi m/T$, where $m \in \mathbb{N}$ and $T = d - c$ is the period. Note that these features fit neatly into the Fourier feature framework, which means that the calculation of Ψ and Φ remains the same as before (now with $[a, b] = [c, d]$).

⁴This is similar to the VFF method, except that the process is completely determined by \mathbf{u} , in the sense that the conditional variance $f | \mathbf{u}$ is zero.

If we choose a matrix G with non-zero off-diagonal elements, the correlations between different basis functions lead to a dependence on the start of the domain (i.e., the results are no longer invariant under translation of the data along the periodic dimension). We avoid this by setting $G = \text{diag}(s(\omega_m)^{-1})$, where $s(\omega)$ is the spectral density of a Matérn kernel. This results in the same kernel as the *sparse spectrum GP* (Lázaro-Gredilla et al., 2010), but whilst periodicity was avoided in that work by randomizing the frequencies, we have deliberately selected the frequencies to capture a suitable prior over periodic functions.

The effect of the spectrum $s(\omega)$ is similar to non-periodic kernels. For example, the spectrum of Matérn- $1/2$ decays more slowly than that of the higher-order Matérn kernels. This means they contain a relatively stronger contribution of higher frequencies, which describes less regular functions.

There is a non-identifiability issue between variance and lengthscale: increasing the lengthscale reduces the overall mass $\sum_{m=1}^M s(\omega_m)$ and mimics a reduction in variance. This makes it hard to optimize or do inference on the hyperparameters. This issue can be avoided by normalizing the discrete spectrum and only introducing the variance σ^2 at the end as a multiplicative scaling factor, i.e., $s'(\omega_i) = \sigma^2 s(\omega_i) / (\sum_{m=1}^M s(\omega_m))$.

4.2. Uncertainty Prediction in Variational Inference

Beyond inferring a mean intensity function, we can also compute an uncertainty interval for the intensity function, given the data.

Due to the highly skewed likelihood of the Poisson process, it does not make sense to consider the variance of λ directly. However, we can evaluate the percentiles numerically. We have $f(\mathbf{x}^*) \sim \mathcal{N}(\tilde{\mu}, \tilde{\sigma}^2)$ (cf. eq. (8)). First, we normalize, $g = f/\tilde{\sigma}$, so that $g \sim \mathcal{N}(\tilde{\mu}/\tilde{\sigma}, 1)$. Then $g^2 = f^2/\tilde{\sigma}^2$ is distributed according to a non-central χ^2 distribution, $g^2 \sim \chi^2(k = 1, \lambda = \tilde{\mu}^2/\tilde{\sigma}^2)$, for which there exist standard library functions to evaluate the percentiles. We can backtransform these by $f^2 = g^2\tilde{\sigma}^2$. We evaluate this using the conditional mean and variance of the posterior GP at each point for which we want to compute the percentiles.

4.3. Test Set Likelihood

To compare models and evaluate their performance, we need to calculate the test set likelihood. This is difficult in a point process model due to the double intractability. The probability density for a test set $\mathcal{D}^* = \{\mathbf{x}_n^*\}_{n=1}^{N^*}$ given our model trained on the training set \mathcal{D} is exactly given by

$$p(\mathcal{D}^* | \mathcal{D}) = \int p(\mathcal{D}^* | f(\cdot)) p(f(\cdot) | \mathcal{D}) df(\cdot), \quad (15)$$

where $p(\mathcal{D}^* | f(\cdot)) = \exp(-\int_{\mathcal{T}} \lambda(\mathbf{x}) d\mathbf{x}) \prod_{n=1}^{N^*} \lambda(\mathbf{x}_n^*)$ (assuming $O^* = 1$) with $\lambda(\cdot) = (f(\cdot) + \beta)^2$. We generally

want to calculate log densities:

$$\mathcal{L}_{\text{exact}}^* = \log p(\mathcal{D}^* | \mathcal{D}) \approx \log \int q(f(\cdot)) p(\mathcal{D}^* | f(\cdot)) df(\cdot), \quad (16)$$

where we approximated the true posterior by $q(f)$. Lloyd et al. (2015) further approximate this using Jensen’s inequality, similar to the ELBO:

$$\mathcal{L}_{\text{ELBO}}^* = \int q(f(\cdot)) \log p(\mathcal{D}^* | f(\cdot)) df(\cdot) \leq \mathcal{L}_{\text{exact}}^*. \quad (17)$$

However, there is no guarantee that $\mathcal{L}_{\text{ELBO}}^* \approx \mathcal{L}_{\text{exact}}^*$. Instead, we consider the mean likelihood that assumes an inhomogeneous Poisson process with a deterministic rate function that is given by the mean $\hat{\lambda}(\cdot) = \tilde{\mu}(\cdot)^2 + \tilde{\sigma}^2(\cdot)$:

$$\mathcal{L}_{\text{mean}}^* = \log p(\mathcal{D}^* | \hat{\lambda}) = -\int_{\mathcal{T}} \hat{\lambda}(\mathbf{x}) d\mathbf{x} + \sum_{n=1}^{N^*} \hat{\lambda}(\mathbf{x}_n^*). \quad (18)$$

We can approximate the integral in the exact test set log density by drawing samples from the posterior, $f_s \sim q(f(\cdot))$:

$$\mathcal{L}_{\text{exact}}^* \approx \mathcal{L}_{\text{sample}}^* = \log \frac{1}{N_s} \sum_{s=1}^{N_s} p(\mathcal{D}^* | f_s(\cdot)). \quad (19)$$

In practice, the samples need to be evaluated on a fine grid, which is infeasible in multiple dimensions. We show for the 1D case that eq. (18) and eq. (19) generally correlate well (cf. Figure S1 in the Supplementary Material). If not specified otherwise, in this paper we use eq. (18) to calculate test set likelihoods. We evaluate the integrals numerically.

4.4. Initialization

When optimizing non-convex objective functions, a reasonable starting value for the unknown variables can help to avoid bad local optima. Here we discuss our approaches to initializing hyperparameters and the variational distribution.

Constant offset β . We can initialize $\beta = \sqrt{\bar{\lambda}}$ from the mean rate $\bar{\lambda} = \langle N_o \rangle / |\mathcal{T}|$, where $\langle N_o \rangle = \sum_{o=1}^O N_o / O$ is the average events count. If the rate is completely uniform, this is exact. In general, it is an overestimate; due to the square-root link function, positive f will influence the mean more strongly than negative f . We can account for this empirically by adjusting β downward, e.g., $\beta = 2/3\sqrt{\bar{\lambda}}$.

Kernel lengthscale. We can obtain a good starting point for the lengthscale hyperparameters by visually inspecting the spatial distribution of data or using the estimated bandwidth from kernel intensity smoothing (KIS).

Kernel variance. For a “wiggly” rate, the magnitude scale of the GP $f(\cdot)$ should be on the order of β to be able to describe regions of low intensity. In this case, we can initialize the variance $\sigma^2 \approx \beta^2 \approx \bar{\lambda}$. Again, this is generally an overestimate, and for more even distributions of events the variance will be lower. This can be determined through visual inspection, based on KIS or otherwise.

For MCMC, we put Gamma priors on the hyperparameters, with shape and scale chosen such that the mean of the prior corresponds to the point estimate discussed before, and the standard deviation of the prior is sufficiently smaller than the mean. This gives the MCMC the freedom to explore but ensures that there is zero mass at zero.

Variational approximate distribution. In variational inference, we also need to initialize the mean \mathbf{m} and covariance S of the approximating Gaussian distribution. When using an appropriate initial value for β , in practice, it works well to simply initialize with a zero mean, $\mathbf{m} = 0$. To obtain different starting conditions, we can draw $\mathbf{m} \sim p(\mathbf{u}) = \mathcal{N}(\mathbf{0}, K_{\mathbf{u}\mathbf{u}})$.

In order for the KL term in the ELBO not to swamp the other terms (which would make optimization difficult), it is advantageous to initialize S with the covariance matrix of the prior, $S = K_{\mathbf{u}\mathbf{u}}$. When using product kernels with Fourier features in two or more dimensions, we want S to be a sum of two Kronecker-structured matrices; in this case, it is not possible to simply set $S = K_{\mathbf{u}\mathbf{u}}$, but we can approximate it. Note that the two summands need to be different to break the symmetry. Ensuring that the summands are not related by a simple scale factor is crucial to avoid duplicate eigenvalues. We choose $S_1^i = 0.2K_{\mathbf{u}\mathbf{u}}^i$ and $S_2^i = K_{\mathbf{u}\mathbf{u}}^i + 0.2$.

4.5. Nodal Lines

Since the inverse link function $\lambda(\cdot) = f(\cdot)^2$ is not injective, different $f(\cdot)$ s can lead to similar rates. For example, $\pm f(\cdot)$ have the same associated rate. A pathological case can arise when a rate is low but not zero around some location x_0 of the input space. Consider a 1D example, in which a function $f_1(\cdot)$ is positive in some regions and negative in others, with a zero crossing at x_0 . This may be a good initial guess early on in the optimization, especially if there are regions with low intensity, and, hence, no direct penalty at the zero crossing. However, this leads to a nodal line⁵ at x_0 , where the rate is forced to be zero. A better candidate $f_2(\cdot)$ would be positive everywhere and low around x_0 . Both f_1 and f_2 constitute local optima, but stepping from one to the other is almost impossible because of the high energy barrier that separates the two (half of $f_1(\cdot)$ would need to change sign, and the intermediate steps would be highly penalized). Once a zero crossing exists, it is unlikely to ever disappear again, which severely restricts the optimization. One example is shown in Figure 1. This is particularly an issue in higher dimensions, and when regions with large numbers of events alternate with low-event regions. When using MCMC, this leads to many local modes, with no mixing between them. For the chains to be able to explore, we need to end up in

⁵We refer to these artifacts as *nodal lines* because of their resemblance to that effect in wave interference.




Figure 1. A close-up of the result for λ_2 and $O = 10$ in Figure 2, showing the effect of nodal lines in a model fit (orange) and a better model fit without zero crossings (blue). Where f crosses zero, λ also collapses to zero (black circles). This is even the case at a point where the actual rate is high. The nodal lines lead to a short lengthscale estimation and oscillations throughout the domain.

the mode with $f(\cdot)$ positive everywhere.

We mitigate this problem by including the constant offset term β , as $(f + \beta)$ is less likely to cross zero. It is also important to choose a sufficiently large number of inducing features (frequencies or points). If there are not enough inducing features to represent small-scale variations, we are more likely to end up with nodal lines, as they make it easier for the intensity to quickly go to zero and back up again.

5. Empirical Results

We compare VFF with the inducing point approach using the Gaussian kernel⁶ (denoted Gauß+IP). For VFF we use the Matérn- $5/2$ kernel, as its behavior is similar to the Gaussian kernel. As the inducing features need to provide support across the entire domain, we put inducing points on a regular D -dimensional grid, without optimizing their positions. As a baseline, we include KIS with edge correction (Diggle, 1985), optimized using the leave-one-out objective. For MCMC we make use of whitening, representing $\mathbf{u} = R\mathbf{v}$, where $RR^T = K_{\mathbf{u}\mathbf{u}}$ and $\mathbf{v} \sim \mathcal{N}(\mathbf{0}, I)$.

5.1. 1D Comparison

We use synthetic 1D examples to show that Fourier features have the same expressive power as Gauß+IP and to demonstrate our methods: uncertainty prediction, MCMC, and sampled vs. mean test set likelihood. We use the three example intensity functions from Adams et al. (2009) as shown in Figure 2 and described in the Supplementary Material.

For variational inference, hyperparameters are initialized as in section 4.4. We train VFF, Gauß+IP, and KIS on training sets containing 1, 10, and 100 draws from the ground truth. The mean inferred intensity and the 5/95 percentiles are shown in Figure 2. We calculate test set likelihoods both

⁶Also known as the squared exponential or exponentiated quadratic kernel.

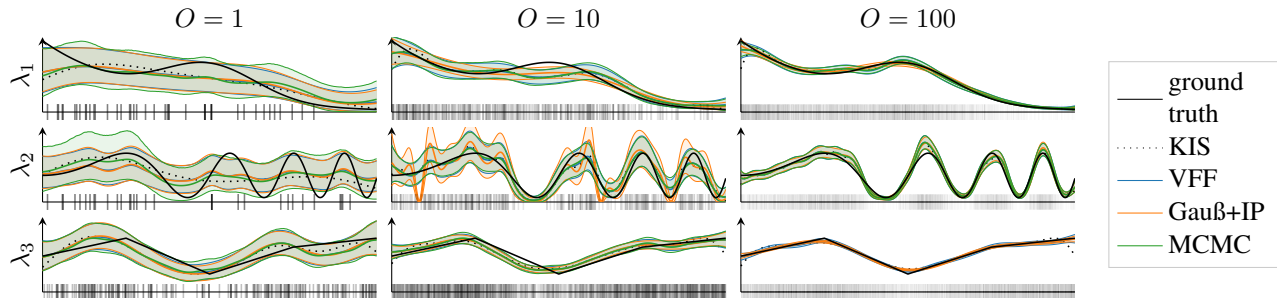


Figure 2. Inferred rate function for the 1D systems λ_1 , λ_2 , and λ_3 from Adams et al. (2009) (rows from top to bottom), based on 1, 10, or 100 draws from the distribution (columns from left to right). We compare the optimized results for VFF (blue) and Gauß+IP (orange) with KIS (dotted) and VFF+MCMC (green). VFF and Gauß+IP give comparable results, and we show the 5/95 percentile interval. With an increased number of observations O , the inferred mean gets closer to the ground truth (black), and the confidence interval becomes tighter. Models are based on $M = 40$ frequencies and $M = 40$ inducing points.

using $\mathcal{L}_{\text{sample}}^*$ and $\mathcal{L}_{\text{mean}}^*$. Figure S1 in the Supplementary Material shows that they correlate well with each other. For each, the likelihoods for VFF and Gauß+IP differ by less than 1%, demonstrating that both approaches have the same capability of modeling these simple data. The resulting point estimates for the hyperparameters are compared with the initial values in Figure S2 in the Supplementary Material. For MCMC, the posterior distributions for the hyperparameters are shown in Figure S3 in the Supplementary Material. For 22 400 data points and 20 frequencies, MCMC takes about 3.5 s per sample; it scales linearly in N and M .

5.2. Real World Data: Porto Taxi Pickups

The real advantage of VFF is in modeling complex, large-scale data in higher dimensions. Here we apply our model to real-world data from the Porto taxi trajectory data set (Moreira-Matias et al., 2013). This contains 1.7×10^6 trajectories covering the entire year from 1 July 2013 to 30 June 2014. We selected only the pickup locations of the trajectories to model. We left the taxi ranks aside (these can be modeled separately, when needed) and focus on the pickups scattered throughout the downtown area.⁷ This results in a data set with 462k pickups across 365 days. We focus on two separate regimes, a spatial model trained on a single day and a spatiotemporal model trained on 100 days.

In each case, we consider multiple optimizations from different initial values for kernel variance, $\sigma^2 \in \{\bar{\lambda}, 1/2\bar{\lambda}\}$, and constant offset, $\beta \in \{\bar{\lambda}^{1/2}, 2/3\bar{\lambda}^{1/2}\}$. We initialize the mean of the variational approximate distribution both to $\mathbf{m} = 0$ and to random draws $\mathbf{m} \sim \mathcal{N}(0, \mathbf{K}_{\mathbf{uu}})$.

2D comparison. Here we show the benefit of Bayesian inference when the data are sparse. The training data are

⁷We manually removed the area around Cais da Estiva and Cais da Ribeira from the analysis; this has a very sharp peak due to the density of bars, and is more appropriately modeled as another “taxi rank”.

the 1089 pickups in the downtown area from Monday 1 July 2013. We compare variational inference using VFF and Gauß+IP and KIS. Figure 3 shows that by using a product kernel and VFF we can resolve features better than Gauß+IP. Moreover, optimization was much faster. VFF took about one second per gradient step, resulting in run times on a single core of 15–30 min for 30×30 frequencies, 1 h for 40×40 , and 2–4 h for 50×50 . Gauß+IP took 50–500 s per gradient step, depending on the size of the grid, and the optimization run times were 12 h (30×30 inducing points), 39 h (40×40), and for 50×50 inducing points, the optimization still had not converged after 48 h. Moreover, we found that for the Gauß+IP model, many optimization runs failed due to numerical instabilities. Without the constant offset (fixing $\beta = 0$), the model ends up with nodal lines. This is shown for one example fit in Figure 3.

The model fit with the highest ELBO does not necessarily describe the best fit overall. For example, for the 50×50 VFF model, the fit with the highest ELBO had larger lengthscales and both lower training and lower test set likelihoods than other fits of the same model with a slightly lower ELBO. For the Gauß+IP model, we found a model fit that has a very low ELBO, but very high likelihood (both training and test set), though with a severe amount of nodal lines. For different hyperparameters, the gap (given by the KL term) may be different. This results in a bias towards a different lengthscales. For this reason, we chose model fits manually from the different runs based on crossvalidation – comparing the training set likelihoods and inferred mean intensity for those model fits that had the highest ELBO values. This indicates that point estimates of the hyperparameters are a bad idea. We need good priors to find the local minimum that describes the data best.

Large-scale demonstration. Here we demonstrate that our method can scale to very large data sets using variational inference. We consider a 3D spatiotemporal model where we include time-of-day. As the training set we took 100

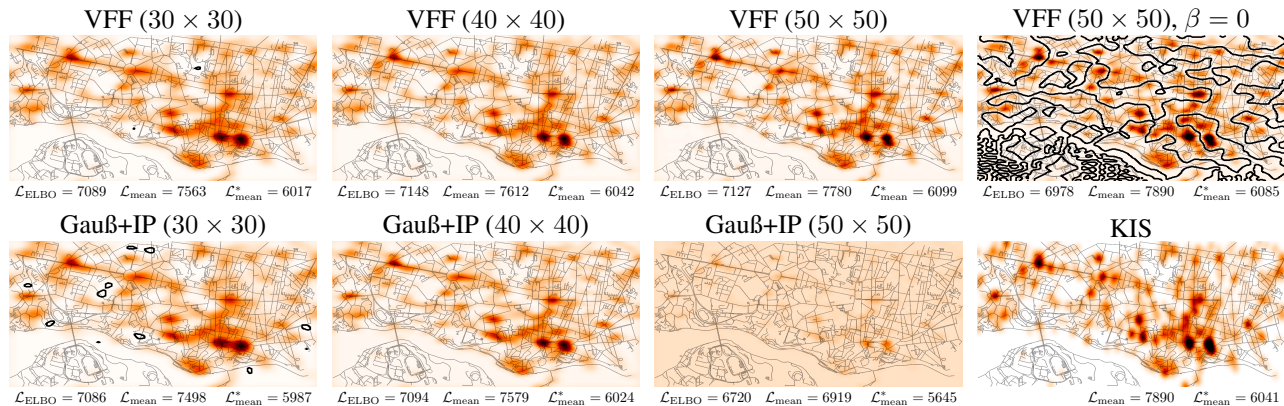


Figure 3. Inferred rate function for a 2D model trained on one day of the Porto taxi data set. We compare KIS, VFF and Gauß+IP for different numbers of inducing frequencies/points (all figures have the same color scale). We give ELBO and the mean likelihoods for training set, $\mathcal{L}_{\text{mean}}$, and test set, $\mathcal{L}_{\text{mean}}^*$. Including more frequencies leads to a better-resolved rate function, and nodal lines (black contours) disappear. Gauß+IP takes a long time and is very hard to optimize; after 48 core-hours the 50×50 model still shows a very flat rate function. At the example of VFF with 50×50 frequencies, we show how fixing $\beta = 0$ leads to many nodal lines. The KIS model overfits.

weekdays, containing 113 020 events in total.

We chose 35×35 frequencies for the spatial dimensions, and a periodic kernel based on the Matérn- $5/2$ spectrum as described in section 4.1 with 25 frequencies for the time dimension. Optimizing the model for 20 000 gradient steps took ca. 24–28 h on a Tesla P100 GPU. We are mainly limited by the amount of GPU memory; the memory requirements grow linearly with the number of data points (and linearly with the number of frequencies in any one dimension, or cubically with the number of frequencies for all three dimensions).

In Figure 4, we show time slices through the resulting 3D model at every three hours. Note that it shows improved spatial resolution compared to the 2D model. Without the periodic kernel on the time dimension, the inferred rate for 00:00 would not match up with that for 24:00, see Figure S4 in the Supplementary Material.

We model the rate of events with large spatial and temporal variation across a large city area. Our model is continuous in space and time, does not rely on a pre-determined grid, and we can encode the inherent temporal periodicity. To the best of our knowledge, this is the first time a principled Bayesian model of this scale has been made possible.

6. Discussion and Conclusions

We have presented two variational approaches to inference in Gaussian process modulated Poisson processes, using Gaussian approximations and MCMC. We have shown that our Fourier-based approximations are effective and robust, and scale to large data sets.

We have discussed the effect of *nodal lines* in the model,

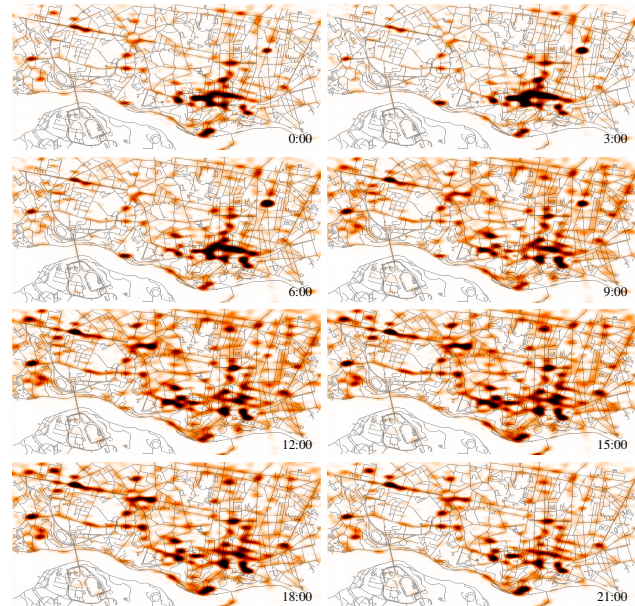


Figure 4. Time slices through the 3D spatiotemporal model of the rate of taxi pickups, trained on 113k events from the Porto taxi data set. A video is available as part of the Supplementary Material.

which are due to the square-root link function: we showed how to mitigate these effects by selecting suitable hyperparameter initial conditions (for optimization) and priors (for MCMC). We discussed that the ELBO is not necessarily good for model selection, and that cross-validation remains important. We found empirically that the choice of optimizer and initial condition can have a huge effect on the resulting model, with Fourier features giving more consistent optimization behavior than inducing points.

References

- Adams, R. P., Murray, I., and MacKay, D. J. C. Tractable nonparametric Bayesian inference in Poisson processes with Gaussian process intensities. In *ICML*, pp. 9–16, 2009.
- Diggle, P. J. A kernel method for smoothing point process data. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 34(2):138–147, 1985.
- Diggle, P. J., Moraga, P., Rowlingson, B., and Taylor, B. M. Spatial and spatio-temporal log-Gaussian Cox processes: Extending the geostatistical paradigm. *Statistical Science*, 28(4):542–563, 2013. URL <https://doi.org/10.1214/13-STS441>.
- Flaxman, S., Whye Teh, Y., and Sejdinovic, D. Poisson intensity estimation with reproducing kernels. *Electronic Journal of Statistics*, 11(2):5081–5104, 2017. URL <https://doi.org/10.1214/17-EJS1339SI>.
- Flaxman, S., Chirico, M., Pereira, P., and Loeffler, C. Scalable high-resolution forecasting of sparse spatiotemporal events with kernel methods: A winning solution to the NIJ “Real-Time Crime Forecasting Challenge”. *ArXiv e-prints, arXiv:1801.02858*, 2018.
- Gunter, T., Lloyd, C., Osborne, M. A., and Roberts, S. J. Efficient Bayesian nonparametric modelling of structured point processes. In *UAI*, pp. 310–319, 2014. URL <http://www.auai.org/uai2014/proceedings/individuals/258.pdf>.
- Hensman, J., Matthews, A. G., Filippone, M., and Ghahramani, Z. MCMC for variationally sparse Gaussian processes. In *NIPS*, pp. 1648–1656, 2015. URL <http://papers.nips.cc/paper/5875-mcmc-for-variationally-sparse-gaussian-processes.pdf>.
- Hensman, J., Durrande, N., and Solin, A. Variational Fourier features for Gaussian processes. *Journal of Machine Learning Research*, 18(151):1–52, 2018. URL <http://jmlr.org/papers/v18/16-579>.
- Lázaro-Gredilla, M., Quiñonero-Candela, J., Rasmussen, C. E., and Figueiras-Vidal, A. R. Sparse spectrum Gaussian process regression. *Journal of Machine Learning Research*, 11(Jun):1865–1881, 2010. URL <http://www.jmlr.org/papers/v11/lazaro-gredilla10a>.
- Lloyd, C., Gunter, T., Osborne, M. A., and Roberts, S. J. Variational inference for Gaussian process modulated Poisson processes. In *ICML*, pp. 1814–1822, 2015. URL <http://proceedings.mlr.press/v37/lloyd15.html>.
- Matthews, A. G. *Scalable Gaussian process inference using variational methods*. PhD thesis, Department of Engineering, University of Cambridge, 2016.
- Matthews, A. G., Hensman, J., Turner, R., and Ghahramani, Z. On sparse variational methods and the Kullback-Leibler divergence between stochastic processes. In *AISTATS*, pp. 231–239, 2016. URL <http://proceedings.mlr.press/v51/matthews16.pdf>.
- Møller, J., Syversveen, A. R., and Waagepetersen, R. P. Log Gaussian Cox processes. *Scandinavian Journal of Statistics*, 25(3):451–482, 1998.
- Moreira-Matias, L., Gama, J., Ferreira, M., Mendes-Moreira, J., and Damas, L. Predicting taxi-passenger demand using streaming data. *IEEE Transactions on Intelligent Transportation Systems*, 14(3):1393–1402, 2013.
- Murray, I. *Advances in Markov chain Monte Carlo methods*. PhD thesis, Gatsby Computational Neuroscience Unit, University College London, 2007.
- Rahimi, A. and Recht, B. Random features for large-scale kernel machines. In *NIPS*, pp. 1177–1184, 2008. URL <https://papers.nips.cc/paper/3182-random-features-for-large-scale-kernel-machines.pdf>.
- Taylor, B. M., Davies, T., Rowlingson, B., and Diggle, P. J. Bayesian inference and data augmentation schemes for spatial, spatiotemporal and multivariate log-Gaussian Cox processes in R. *Journal of Statistical Software*, 63(7):1–48, 2015. URL <https://www.jstatsoft.org/article/view/v063i07>.
- Titsias, M. and Lawrence, N. D. Bayesian Gaussian process latent variable model. In *AISTATS*, pp. 844–851, 2010. URL <http://proceedings.mlr.press/v9/titsias10a/titsias10a.pdf>.
- Vanhatalo, J. and Vehtari, A. Sparse log Gaussian processes via MCMC for spatial epidemiology. In *Gaussian processes in practice*, pp. 73–89, 2007. URL <http://gpss.cc/gpip/abstract/vehtari.pdf>.
- Walder, C. J. and Bishop, A. N. Fast Bayesian intensity estimation for the permanental process. In *ICML*, pp. 3579–3588, 2017. URL <http://proceedings.mlr.press/v70/walder17a.html>.