

Poisson intensity estimation with reproducing kernels*

Seth Flaxman, Yee Whye Teh and Dino Sejdinovic

*Department of Statistics
24-29 St Giles'
Oxford OX1 3LB
United Kingdom
e-mail:*

flaxman@stats.ox.ac.uk; y.w.teh@stats.ox.ac.uk; dino.sejdinovic@stats.ox.ac.uk

Abstract: Despite the fundamental nature of the inhomogeneous Poisson process in the theory and application of stochastic processes, and its attractive generalizations (e.g. Cox process), few tractable nonparametric modeling approaches of intensity functions exist, especially when observed points lie in a high-dimensional space. In this paper we develop a new, computationally tractable Reproducing Kernel Hilbert Space (RKHS) formulation for the inhomogeneous Poisson process. We model the square root of the intensity as an RKHS function. Whereas RKHS models used in supervised learning rely on the so-called representer theorem, the form of the inhomogeneous Poisson process likelihood means that the representer theorem does not apply. However, we prove that the representer theorem does hold in an appropriately transformed RKHS, guaranteeing that the optimization of the penalized likelihood can be cast as a tractable finite-dimensional problem. The resulting approach is simple to implement, and readily scales to high dimensions and large-scale datasets.

MSC 2010 subject classifications: Primary 62G05, 60G55, 46E22.

Keywords and phrases: Nonparametric statistics, computational statistics, spatial statistics, intensity estimation, reproducing kernel Hilbert space, inhomogeneous Poisson processes.

Received June 2017.

Contents

1	Introduction	5082
2	Background and related work	5083
	2.1 Poisson process	5083
	2.2 Reproducing kernel Hilbert spaces	5083
	2.3 Related work	5084
3	Proposed method and kernel transformation	5085
4	Computation of \tilde{k}	5087
	4.1 Explicit Mercer expansion	5088
	4.1.1 Sobolev space on $[0, 1]$ with a periodic boundary condition	5089

*This work was supported by ERC (FP7/617071) and EPSRC (EP/K009362/1). It is an extended version of [14].

4.1.2	Squared exponential kernel	5090
4.1.3	Extending the Mercer expansion to multiple dimensions	5091
4.2	Numerical approximation when Mercer expansions are not available	5091
5	Inference	5093
5.1	Hyperparameter selection	5094
6	Naïve RKHS model	5094
7	Experiments	5095
7.1	1-d synthetic example	5095
7.2	Environmental datasets	5095
7.3	High dimensional synthetic examples	5096
7.4	Computational complexity	5099
7.5	Spatiotemporal point pattern of crimes	5101
8	Conclusion	5101
	References	5102

1. Introduction

Poisson processes are ubiquitous in statistical science, with a long history spanning both theory (e.g. [20]) and applications (e.g. [12]), especially in the spatial statistics and time series literature. Despite their ubiquity, fundamental questions in their application to real datasets remain open. Namely, scalable nonparametric models for intensity functions of inhomogeneous Poisson processes are not well understood, especially in multiple dimensions since the standard approaches, based on kernel smoothing, are akin to density estimation and hence scale poorly with dimension. In this contribution, we propose a step towards such scalable nonparametric modeling and introduce a new Reproducing Kernel Hilbert Space (RKHS) formulation for inhomogeneous Poisson process modeling, which is based on the Empirical Risk Minimization (ERM) framework. We model the square root of the intensity as an RKHS function and consider a risk functional given by a penalized version of the inhomogeneous Poisson process likelihood. However, standard representer theorem arguments do not apply directly due to the form of the likelihood. Namely, the fundamental difference arises since the observation that *no points* occur in some region is just as important as the locations of the points that do occur. Thus, the likelihood depends not only on the evaluations of the intensity at the observed points, but also on its integral across the domain of interest. As we will see, this difficulty can be overcome by appropriately adjusting the RKHS under consideration. We prove a version of the representer theorem in this adjusted RKHS, which coincides with the original RKHS as a space of functions but has a different inner product structure. This allows us to cast the estimation problem as an optimization over a finite-dimensional subspace of the adjusted RKHS. The derived method is demonstrated to give better performance than a naïve unadjusted RKHS method which resorts to an optimization over a subspace without representer theorem guarantees. We describe cases where adjusted RKHS can be described

with explicit Mercer expansions and propose numerical approximations where Mercer expansions are not available. We observe strong performance of the proposed method on a variety of synthetic, environmental, crime and bioinformatics data.

2. Background and related work

2.1. Poisson process

We briefly state relevant definitions for point processes over domains $S \subset \mathbb{R}^D$, following [8]. For Lebesgue measurable subsets $T \subset S$, $N(T)$ denotes the number of events in $T \subset S$. $N(\cdot)$ is a stochastic process characterizing the point process. Our focus is on providing a nonparametric estimator for the first-order intensity of a point process, which is defined as:

$$\lambda(s) = \lim_{|ds| \rightarrow 0} \mathbb{E}[N(ds)]/|ds|. \quad (2.1)$$

The inhomogeneous Poisson process is driven solely by the intensity function $\lambda(\cdot)$:

$$N(T) \sim \text{Poisson}\left(\int_T \lambda(x) dx\right). \quad (2.2)$$

In the homogeneous Poisson process, $\lambda(x) = \lambda$ is constant, so the number of points in any region T simply depends on the volume of T , which we denote $|T|$:

$$N(T) \sim \text{Poisson}(\lambda|T|). \quad (2.3)$$

For a given intensity function $\lambda(\cdot)$, the likelihood of a set of $N = N(S)$ points x_1, \dots, x_N observed over a domain S is given by:

$$\mathcal{L}(x_1, \dots, x_N | \lambda(\cdot)) = \prod_{i=1}^N \lambda(x_i) e^{-\int_S \lambda(x) dx} \quad (2.4)$$

2.2. Reproducing kernel Hilbert spaces

Given a non-empty domain S and a positive definite kernel function $k : S \times S \rightarrow \mathbb{R}$, there exists a unique reproducing kernel Hilbert space (RKHS) \mathcal{H}_k . An RKHS is a space of functions $f : S \rightarrow \mathbb{R}$, in which evaluation is a continuous functional, meaning it can be represented by an inner product $f(x) = \langle f, k(x, \cdot) \rangle_{\mathcal{H}_k}$ for all $f \in \mathcal{H}_k, x \in S$ (this is known as the reproducing property), cf. Berlinet and Thomas-Agnan [5]. While \mathcal{H}_k is in most interesting cases an infinite-dimensional space of functions, due to the classical representer theorem [19], [29, Section 4.2], optimization over \mathcal{H}_k is typically a tractable finite-dimensional problem. In particular, if we have a set of N observations $x_1, \dots, x_N, x_i \in S$ and consider the problem:

$$\min_{f \in \mathcal{H}_k} \{R(f(x_1), \dots, f(x_N)) + \Omega(\|f\|_{\mathcal{H}_k})\}. \quad (2.5)$$

where $R(f(x_1), \dots, f(x_N))$ depends on f through its evaluations on the set of observations only, and Ω is a non-decreasing function of the RKHS norm of f , there exists a solution to Eq. (2.5) of the form $f^*(\cdot) = \sum_{i=1}^N \alpha_i k(x_i, \cdot)$, and the optimization can thus be cast in terms of the so-called *dual coefficients* $\alpha \in \mathbb{R}^N$. This formulation is widely used in the framework of regularized Empirical Risk Minimization (ERM) for supervised learning, where $R(f(x_1), \dots, f(x_N)) = \frac{1}{N} \sum_{i=1}^N L(f(x_i), y_i)$ is the empirical risk corresponding to a loss function L , e.g. squared loss for regression, logistic or hinge loss for classification.

If domain S is compact and kernel k is continuous, one can assign to k its integral kernel operator $T_k : L_2(S) \rightarrow L_2(S)$, given by $T_k g = \int_S k(x, \cdot) g(x) dx$, which is positive, self-adjoint and compact. There thus exists an orthonormal set of eigenfunctions $\{e_j\}_{j=1}^\infty$ of T_k and the corresponding eigenvalues $\{\eta_j\}_{j=1}^\infty$, with $\eta_j \rightarrow 0$ as $j \rightarrow \infty$. This spectral decomposition of T_k leads to Mercer's representation of kernel function k [29, Section 2.2]:

$$k(x, x') = \sum_{j=1}^{\infty} \eta_j e_j(x) e_j(x'), \quad x, x' \in S \quad (2.6)$$

with uniform convergence on $S \times S$. Any function $f \in \mathcal{H}_k$ can then be written as $f = \sum_j b_j e_j$ where $\|f\|_{\mathcal{H}_k}^2 = \sum_j b_j^2 / \eta_j < \infty$.

Note that above we have focused on Mercer expansion with respect to the Lebesgue measure, but other base measures are also often considered in literature, e.g. [28, section 4.3.1].

2.3. Related work

The classic approach to nonparametric intensity estimation is based on smoothing kernels [27, 11] and has a form closely related to the kernel density estimator:

$$\hat{\lambda}(x) = \sum_{i=1}^N \kappa(x_i - x) \quad (2.7)$$

where κ is a smoothing kernel (related to but distinct from the RKHS kernels described in the previous section), that is, any bounded function integrating to 1. Early work in this area focused on edge-corrections and methods for choosing the bandwidth [11, 6, 7]. Connections with RKHS have been considered by, for example, Bartoszyński et al. [4] who use a maximum penalized likelihood approach based on Hilbert spaces to estimate the intensity of a Poisson process. There is long literature on maximum penalized likelihood approaches to density estimation, which also contain interesting connections with RKHS, e.g. [30].

Much recent work on estimating intensities for point processes has focused on Bayesian approaches to modeling Cox processes. The log Gaussian Cox Process [24] and related parameterizations of Cox (doubly stochastic) Poisson processes in terms of Gaussian processes have been proposed, along with Monte Carlo

[1, 12, 31], Laplace approximation [17, 10, 15] and variational [22, 21] inference schemes.

Another related body of literature concerns Cox processes with intensities parameterized as the sum of squares of k Gaussian processes, called the permanent process [23]. Interestingly, calculating the density of the permanent process relies on a kernel transformation similar to the one we propose below. Unlike these approaches, however, we are not working in a doubly stochastic (Cox process) framework; rather we are taking a penalized maximum likelihood estimation perspective to estimate the intensity of an inhomogeneous Poisson process. As future work, it would be worthwhile to explore deeper connections between our formulation and the permanent process, e.g. by considering an RKHS formulation of Cox processes or by considering an inhomogeneous Poisson process whose intensity is the sum of squares of functions in an RKHS.

3. Proposed method and kernel transformation

Let S be a compact domain of observations. Let $k : S \times S \rightarrow \mathbb{R}$ be a continuous positive definite kernel, and \mathcal{H}_k its corresponding RKHS of functions $f : S \rightarrow \mathbb{R}$. We model the intensity function $\lambda(\cdot)$ of an inhomogeneous Poisson process as:

$$\lambda(x) := af^2(x), \quad x \in S, \quad (3.1)$$

which is parameterized by $f \in \mathcal{H}_k$ and an additional scale parameter $a > 0$. The flexibility of choosing k means that we can encode structural assumptions of our domain, e.g. periodicity in time or periodic boundary conditions (see Section 4.1.1). Note that we have squared f to ensure that the intensity is non-negative on S , a pragmatic choice that has previously appeared in the literature (e.g. [22]). While we lose identifiability (since f and $-f$ are equivalent), as shown below we end up with a finite dimensional, and thus tractable, optimization problem.

The rationale for including a is that it allows us to decouple the overall scale and units of the intensity (e.g. number of points per hour versus number of points per year) from the penalty on the complexity of f which arises from the classical regularized Empirical Risk Minimization framework (and which should depend only on how complex, i.e. “wiggly” f is).

We use the inhomogeneous Poisson process likelihood from Eq. (2.4) to write the log-likelihood of a Poisson process corresponding to the observations $\{x_1, \dots, x_N\}$, for $x_i \in S$, and intensity $\lambda(\cdot)$:

$$\ell(x_1, \dots, x_N | \lambda) = \sum_{i=1}^N \log(\lambda(x_i)) - \int_S \lambda(x) dx. \quad (3.2)$$

We will consider the problem of minimization of the penalized negative log likelihood, where the regularization term corresponds to the squared Hilbert space norm of f in parametrization Eq. (3.1):

$$\min_{f \in \mathcal{H}_k} \left\{ - \sum_{i=1}^N \log(af^2(x_i)) + a \int_S f^2(x) dx + \gamma \|f\|_{\mathcal{H}_k}^2 \right\}. \quad (3.3)$$

This objective is akin to a classical regularized empirical risk minimization framework over RKHS: there is a term that depends on evaluations of f at the observed points x_1, \dots, x_N as well as a term corresponding to the RKHS norm. However, the representer theorem does not apply directly to Eq. (3.3): since there is also a term given by the L_2 -norm of f , there is no guarantee that there is a solution of Eq. (3.3) that lies in $\text{span}\{k(x_i, \cdot)\}_{i=1}^N$. We will show that Eq. (3.3) fortunately still reduces to a finite-dimensional optimization problem corresponding to a different kernel function \tilde{k} which we define below.

Using the Mercer expansion of k in Eq. (2.6), we can write the objective Eq. (3.3) as follows:

$$J[f] = - \sum_{i=1}^N \log(af^2(x_i)) + a \|f\|_{L_2(S)}^2 + \gamma \|f\|_{\mathcal{H}_k}^2 \quad (3.4)$$

$$= - \sum_{i=1}^N \log(af^2(x_i)) + a \sum_{j=1}^{\infty} b_j^2 + \gamma \sum_{j=1}^{\infty} \frac{b_j^2}{\eta_j}. \quad (3.5)$$

The last two terms can now be merged together, giving

$$a \sum_{j=1}^{\infty} b_j^2 + \gamma \sum_{j=1}^{\infty} \frac{b_j^2}{\eta_j} = \sum_{j=1}^{\infty} b_j^2 \frac{a\eta_j + \gamma}{\eta_j} = \sum_{j=1}^{\infty} \frac{b_j^2}{\eta_j(a\eta_j + \gamma)^{-1}}.$$

Now, if we define kernel \tilde{k} to be the kernel corresponding to the integral operator $T_{\tilde{k}} := T_k(aT_k + \gamma I)^{-1}$, i.e., \tilde{k} is given by:

$$\tilde{k}(x, x') = \sum_{j=1}^{\infty} \frac{\eta_j}{a\eta_j + \gamma} e_j(x) e_j(x'), \quad x, x' \in S,$$

we see that:

$$J[f] = - \sum_{i=1}^N \log(af^2(x_i)) + \|f\|_{\mathcal{H}_{\tilde{k}}}^2. \quad (3.6)$$

Thus, we have merged the two squared norm terms into a squared norm in a new RKHS. We note that a similar idea has previously been used to modify Gaussian process priors in [9], albeit in a different context, and that a similar transformation appears in the expression for the distribution of a permanent process [23]. We are now ready to state the representer theorem in terms of kernel \tilde{k} .

Theorem 1. *There exists a solution of Eq. (3.3) for observations x_1, \dots, x_N , which takes the form $f^*(\cdot) = \sum_{i=1}^N \alpha_i \tilde{k}(x_i, \cdot)$.*

Proof. Since $\sum_j \frac{b_j^2}{\eta_j} < \infty$ if and only if $\sum_j \frac{b_j^2}{\eta_j(a\eta_j + \gamma)^{-1}} < \infty$, i.e. $f \in \mathcal{H}_k \iff f \in \mathcal{H}_{\tilde{k}}$, we have that the two spaces correspond to exactly the same set of functions. Optimization over \mathcal{H}_k is therefore equivalent to optimization over $\mathcal{H}_{\tilde{k}}$.

The proof now follows by applying the classical representer theorem in \tilde{k} to the representation of the objective function in Eq. (3.6). We decompose $f \in \mathcal{H}_{\tilde{k}}$ as the sum of two functions:

$$f(\cdot) = \sum_{j=1}^N \alpha_j \tilde{k}(x_j, \cdot) + v \tag{3.7}$$

where v is orthogonal in $\mathcal{H}_{\tilde{k}}$ to the span of $\{\tilde{k}(x_j, \cdot)\}_j$. We prove that the first term in the objective $J[f]$ given in Eq. (3.6), $-\sum_{i=1}^N \log(af^2(x_i))$, is independent of v . It depends on f only through the evaluations $f(x_i)$ for all i . Using the reproducing property we have:

$$f(x_i) = \langle f, \tilde{k}(x_i, \cdot) \rangle_{\mathcal{H}_{\tilde{k}}} = \sum_j \alpha_j \tilde{k}(x_j, x_i) + \langle v, \tilde{k}(x_i, \cdot) \rangle_{\mathcal{H}_{\tilde{k}}} = \sum_j \alpha_j \tilde{k}(x_j, x_i) \tag{3.8}$$

where the last step is by orthogonality. Next we substitute into the regularization term:

$$\gamma \left\| \sum_j \alpha_j \tilde{k}(x_j, \cdot) + v \right\|_{\mathcal{H}_{\tilde{k}}}^2 = \gamma \left\| \sum_j \alpha_j \tilde{k}(x_j, \cdot) \right\|_{\mathcal{H}_{\tilde{k}}}^2 + \|v\|_{\mathcal{H}_{\tilde{k}}}^2 \geq \gamma \left\| \sum_j \alpha_j \tilde{k}(x_j, \cdot) \right\|_{\mathcal{H}_{\tilde{k}}}^2. \tag{3.9}$$

Thus, the choice of v has no effect on the first term in $J[f]$ and a non-zero v can only increase the second term $\|f\|_{\mathcal{H}_{\tilde{k}}}^2$, so we conclude that $v = 0$ and that $f^* = \sum_{j=1}^N \alpha_j \tilde{k}(x_j, \cdot)$ is the minimizer. \square

Remark 1. The notions of the inner product in \mathcal{H}_k and $\mathcal{H}_{\tilde{k}}$ are different and thus in general $\text{span}\{k(x_i, \cdot)\} \neq \text{span}\{\tilde{k}(x_i, \cdot)\}$.

Remark 2. Notice that unlike in a standard ERM setting, $\gamma = 0$ does not recover the unpenalized risk, because γ appears in \tilde{k} . Notice further that the overall scale parameter a also appears in \tilde{k} . This is important in practice, because it allows us to decouple the scale of the intensity (which is controlled by a) from its complexity (which is controlled by γ).

Illustration. The eigenspectrum of \tilde{k} where k is a squared exponential kernel is shown in Figure 1 for various settings of a and γ . Reminiscent of spectral filtering studied by Muandet, Sriperumbudur and Schölkopf [25], in the top plot we see that depending on the settings of a and γ , eigenvalues of \tilde{k} are shrunk or inflated as compared to $k(x, x')$ which is shown in black. In the bottom plot, the values of $k(0, x)$ are shown for the same set of kernels.

4. Computation of \tilde{k}

In this section, we consider first the case in which an explicit Mercer expansion is known, and then we consider the more commonly encountered situation in which we only have access to the parametric form of the kernel $k(x, x')$, so we must approximate \tilde{k} . We show experimentally that our approximation is very accurate by considering the Sobolev kernel, which can be expressed in both ways.




FIG 1. Eigenspectrum of \bar{k} (top) and values of \bar{k} (bottom) for various settings of a and γ . The original kernel k is shown in solid black in both figures.

4.1. Explicit Mercer expansion

We start by assuming that we have a kernel k with an explicit Mercer expansion with respect to a base measure of interest (usually the Lebesgue measure on S), so we have eigenvectors $\{e_j(x)\}_{j \in J}$ and eigenvalues $\{\eta_j\}_{j \in J}$:

$$k(x, x') = \sum_{j \in J} \eta_j e_j(x) e_j(x'), \tag{4.1}$$

with an at most countable index set J . Given a and γ we can calculate:

$$\tilde{k}(x, x') = \sum_{j \in J} \frac{\eta_j}{a\eta_j + \gamma} e_j(x) e_j(x') \tag{4.2}$$

up to a desired precision as informed by the spectral decay in $\{\eta_j\}_{j \in J}$. Below we consider kernels for which explicit Mercer expansions are known: a kernel on the Sobolev space $[0, 1]$ with a periodic boundary condition, and the squared exponential kernel. We also show how our formulation can be extended to multiple dimensions using a tensor product formulation. Although not practical for large datasets, the Mercer expansions given below, summing terms up to $j = 50$ (for which the error is less than 10^{-5}), can be used to evaluate approximations for the cases in which Mercer expansions are not available.

4.1.1. Sobolev space on $[0, 1]$ with a periodic boundary condition

We consider a kernel on the Sobolev space on $[0, 1]$ with a periodic boundary condition, proposed by Wahba [32, chapter 2] and recently used in Bach [2]. The kernel is given by:

$$\begin{aligned} k(x, y) &= 1 + \sum_{m=1}^{\infty} \frac{2 \cos(2\pi m(x - y))}{(2\pi m)^{2s}} \\ &= 1 + \sum_{m=1}^{\infty} \frac{2}{(2\pi m)^{2s}} [\cos(2\pi mx) \cos(2\pi my) + \sin(2\pi mx) \sin(2\pi my)], \\ &= 1 + \frac{(-1)^{s-1}}{(2s)!} B_{2s}(\{x - y\}), \end{aligned}$$

where $s = 1, 2, \dots$ denotes the order of the Sobolev space and $B_{2s}(\{x - y\})$ is the Bernoulli polynomial of degree $2s$ applied to the fractional part of $x - y$. The corresponding RKHS is the space of functions on $[0, 1]$ with absolutely continuous $f, f', \dots, f^{(s-1)}$ and square integrable $f^{(s)}$ satisfying a periodic boundary condition $f^{(l)}(0) = f^{(l)}(1)$, $l = 0, \dots, s - 1$. For more details, see [32, Chapter 2].

Bernoulli polynomials admit a simple form for low degrees. In particular,

$$\begin{aligned} B_2(t) &= t^2 - t + \frac{1}{6}, \\ B_4(t) &= t^4 - 2t^3 + t^2 - \frac{1}{30}, \\ B_6(t) &= t^6 - 3t^5 + \frac{5}{2}t^4 - \frac{1}{2}t^2 + \frac{1}{42}. \end{aligned}$$

Moreover, note that:

$$\int_0^1 2 \cos(2\pi mx) \sin(2\pi m'x) dx = 0,$$

$$\int_0^1 2 \cos(2\pi mx) \cos(2\pi m'x) dx = \delta(m - m'),$$

$$\int_0^1 2 \sin(2\pi mx) \sin(2\pi m'x) dx = \delta(m - m').$$

Thus, the desired Mercer expansion (with respect to the Lebesgue measure) is given by $k(x, y) = \sum_{m \in \mathbb{Z}} \eta_m e_m(x) e_m(y)$ with eigenfunctions $e_0(x) = 1$ and for $m = \{1, 2, \dots\}$, $e_m(x) = \sqrt{2} \cos(2\pi mx)$, $e_{-m}(x) = \sqrt{2} \sin(2\pi mx)$ and corresponding eigenvalues $\eta_0 = 1$, $\eta_m = \eta_{-m} = (2\pi m)^{-2s}$.

Now, the adjusted kernel $\tilde{k}(x, y)$ from (4.2) is given by

$$\begin{aligned} \tilde{k}(x, y) &= \sum_{m \in \mathbb{Z}} \frac{\eta_m}{\eta_m + \gamma} e_m(x) e_m(y) \\ &= \frac{1}{1 + \gamma} + \sum_{m=1}^{\infty} \frac{2 \cos(2\pi m(x - y))}{1 + \gamma(2\pi m)^{2s}}. \end{aligned}$$

4.1.2. Squared exponential kernel

A Mercer expansion for the squared exponential kernel was proposed in [36] and refined in [13]. However, this expansion is with respect to a Gaussian measure on \mathbb{R} , i.e., it consists of eigenfunctions which form an orthonormal set in $L^2(\mathbb{R}, \nu)$ where $\nu = \mathcal{N}(0, \ell^2 I)$. The formalism can therefore be used to estimate Poisson intensity functions with respect to such Gaussian measure. In the classical framework, where the intensity is with respect to a Lebesgue measure, numerical approximations of Mercer expansion, as described in Section 4.2 are needed. Following the exposition in [28, section 4.3.1] and the relevant errata¹ we parameterize the kernel as:

$$k(x, x') = \exp\left(-\frac{\|x - x'\|^2}{2\sigma^2}\right) \quad (4.3)$$

The Mercer expansion with respect to $\nu = \mathcal{N}(0, \ell^2 I)$ then has the eigenvalues

$$\eta_i = \sqrt{\frac{2a}{A}} B^i, \quad (4.4)$$

and eigenfunctions

$$e_i(x) = \frac{1}{\sqrt{\sqrt{a/c} 2^i i!}} \exp(-(c - a)x^2) H_i(\sqrt{2c}x) \quad (4.5)$$

where H_i is the i -th order (physicist's) Hermite polynomial, $a = \frac{1}{4\sigma^2}$, $b = \frac{1}{2\ell^2}$, $c = \sqrt{a^2 + 2ab}$, $A = a + b + c$, and $B = b/A$. Thus we have the following eigenvalues for \tilde{k} :

$$\tilde{\eta}_i = \frac{\eta_i}{a\eta_i + \gamma} = \frac{1}{a + \gamma \sqrt{\frac{A}{2a}} B^{-i}} \quad (4.6)$$

while the eigenfunctions remain the same.

¹<http://www.gaussianprocess.org/gpml/errata.html>

4.1.3. Extending the Mercer expansion to multiple dimensions

The extension of any kernel to higher dimensions can be constructed by considering tensor product spaces: $\mathcal{H}_{k_1 \otimes k_2}$ (where k_1 and k_2 could potentially be different kernels with different hyperparameters). If k_1 has eigenvalues η_i and eigenfunctions e_i and k_2 has eigenvalues δ_j and eigenfunctions f_j , then the eigenvalues of the product space are then given by the Cartesian product $\eta_i \delta_j, \forall i, j$, and similarly the eigenfunctions are given by $e_i(x)f_j(y)$. Our regularized kernel has the following Mercer expansion:

$$\widetilde{k_1 \otimes k_2}((x, y), (x', y')) = \sum_{ij} \frac{\eta_i \delta_j}{a\eta_i \delta_j + \gamma} e_i(x)e_i(x')f_j(y)f_j(y') \quad (4.7)$$

Notice that $\widetilde{k_1 \otimes k_2}$ is the kernel corresponding to the integral operator $(T_{k_1} \otimes T_{k_2})(aT_{k_1} \otimes T_{k_2} + \gamma I)^{-1}$ which is different than $\tilde{k}_1 \otimes \tilde{k}_2$.

Notice that this approach does not lead to a method that scales well in high dimensions, which is further motivation for the approximations developed below.

4.2. Numerical approximation when Mercer expansions are not available

We propose an approximation to \tilde{k} given access only to a kernel k for which we do not have an explicit Mercer expansion with respect to Lebesgue measure. We only assume that we can form Gram matrices corresponding to k and calculate their eigenvectors and eigenvalues. As a side benefit, this representation will also enable scalable computations through Toeplitz / Kronecker algebra [10, 16, 15] or primal reduced rank approximations [34].

Let us first consider the one-dimensional case where $S = [0, 1]$ and construct a uniform grid $\mathbf{u} = (u_1, \dots, u_m)$ on $[0, 1]$. We consider the relationship between the eigendecomposition of $K_{\mathbf{uu}}$ and the eigenvalues and eigenfunctions of the integral kernel operator T_k . The integral kernel operator T_k can be approximated with the (scaled) kernel matrix $\frac{1}{m}K_{\mathbf{uu}} : \mathbb{R}^m \rightarrow \mathbb{R}^m$, where $[K_{\mathbf{uu}}]_{ij} = k(u_i, u_j)$.

Let $\lambda_i^u, \mathbf{e}_i^u$ be the eigenvalue/eigenvector pairs of the matrix $K_{\mathbf{uu}}$, i.e., its eigendecomposition is given by $K_{\mathbf{uu}} = Q\Lambda Q^\top = \sum_{i=1}^m \lambda_i^u \mathbf{e}_i^u (\mathbf{e}_i^u)^\top$. Then the estimates of the eigenvalues/eigenfunctions of the integral operator T_k are given by the Nyström method (see Rasmussen and Williams [28, Section 4.3] and references therein, especially Baker [3]):

$$\hat{\eta}_i = \frac{1}{m} \lambda_i^u, \quad \hat{e}_i(x) = \frac{\sqrt{m}}{\lambda_i^u} K_{x\mathbf{u}} \mathbf{e}_i^u, \quad (4.8)$$

with $K_{x\mathbf{u}} = [k(x, u_1), \dots, k(x, u_m)]$. We now combine this with Eq. (4.2) to estimate \tilde{k} :

$$\widehat{\tilde{k}}(x, x') = \sum_{i=1}^m \frac{\hat{\eta}_i}{a\hat{\eta}_i + \gamma} \hat{e}_i(x)\hat{e}_i(x') \quad (4.9)$$

$$\begin{aligned}
&= \sum_{i=1}^m \frac{\frac{1}{m}\lambda_i^u}{\frac{a}{m}\lambda_i^u + \gamma} \cdot \frac{m}{(\lambda_i^u)^2} K_{\mathbf{xu}} \mathbf{e}_i^u (\mathbf{e}_i^u)^\top K_{\mathbf{ux}'} \\
&= K_{\mathbf{xu}} \left\{ \sum_{i=1}^m \frac{1}{\left(\frac{a}{m}\lambda_i^u + \gamma\right) \lambda_i^u} \mathbf{e}_i^u (\mathbf{e}_i^u)^\top \right\} K_{\mathbf{ux}'}.
\end{aligned}$$

For an estimate of the whole matrix $\tilde{K}_{\mathbf{xx}}$ we thus have

$$\begin{aligned}
\hat{K}_{\mathbf{xx}} &= K_{\mathbf{xu}} \left\{ \sum_{i=1}^m \frac{1}{\left(\frac{a}{m}\lambda_i^u + \gamma\right) \lambda_i^u} \mathbf{e}_i^u (\mathbf{e}_i^u)^\top \right\} K_{\mathbf{ux}} \\
&= K_{\mathbf{xu}} Q \left(\frac{a}{m} \Lambda^2 + \gamma \Lambda \right)^{-1} Q^\top K_{\mathbf{ux}}. \tag{4.10}
\end{aligned}$$

The above is reminiscent of the Nyström method [34] proposed for speeding up Gaussian process regression. It has computational cost $O(m^3 + N^2m)$. A reduced rank representation for Eq. (4.10) is straightforward by considering only the top p eigenvalues/eigenvectors of $K_{\mathbf{uu}}$. Furthermore, a primal representation with the features corresponding to kernel \tilde{k} is readily available and is given by

$$\tilde{\phi}(x) = \left(\frac{a}{m} \Lambda^2 + \gamma \Lambda \right)^{-1/2} Q^\top K_{\mathbf{ux}}, \tag{4.11}$$

which allows linear computational cost in the number N of observations.

For $D > 1$ dimensions, one can exploit Kronecker and Toeplitz algebra approaches. Without loss of generality, we assume that $\text{vol}(S) = 1$ (otherwise we would have to multiply the kernel matrices below by $\text{vol}(S)$). Assuming that the $K_{\mathbf{uu}}$ matrix corresponds to a Cartesian product structure of the one-dimensional grids of size m , one can write $K_{\mathbf{uu}} = K_1 \otimes K_2 \cdots \otimes K_D$. Thus, the eigenspectrum can be efficiently calculated by eigendecomposing each of the smaller $m \times m$ matrices K_1, \dots, K_D and then applying standard Kronecker algebra, thereby avoiding ever having to form the prohibitively large $m^D \times m^D$ matrix $K_{\mathbf{uu}}$. For regular grids and stationary kernels, each small matrix will be Toeplitz structured, yielding further efficiency gains [35]. The resulting approach thus scales linearly in dimension D .

An even simpler alternative to the above is to sample the points u_1, \dots, u_m uniformly from the domain S using Monte Carlo or Quasi-Monte Carlo (see [26] for a discussion in the context of RKHS). We found this approach to work well in practice in high-dimensions ($D = 15$), even when m was fixed, meaning that the scaling was effectively independent of the dimension D .

Using the Sobolev kernel in Sec. 4.1.1, we compared the exact calculation of $\tilde{K}_{\mathbf{uu}}$ with $s = 1$, $a = 10$, and $\gamma = .5$ to our approximate calculation. For illustration we compared a coarse grid of size 10 on the unit interval (left) to a finer grid of size 100. The RMSE was 2E-3 for the coarse grid and 1.6E-5 for the fine grid, as shown in Fig. 2. In the same figure we compared the exact calculation of $\tilde{K}_{\mathbf{xx}}$ with $s = 1$, $a = 10$, and $\gamma = .5$ to our Nyström-based approximation, where $x_1, \dots, x_{400} \sim \text{Beta}(.5, .5)$ distribution. The RMSE was 0.98E-3. A low-rank approximation using only the top 5 eigenvalues gives the RMSE of 1.6E-2.

As Figure 2, demonstrates, good approximation is possible with a fairly coarse grid $\mathbf{u} = (u_1, \dots, u_m)$ as well as with a low-rank approximation.




FIG 2. Using the Sobolev kernel in Sec. 4.1.1, we compared the exact calculation of $\tilde{K}_{\mathbf{u}\mathbf{u}}$ with $s = 1$, $a = 10$, and $\gamma = .5$ to our approximate calculation. For illustration we tried a coarse grid of size 10 on the unit interval (top left) to a finer grid of size 100 (top right). The RMSE was $2E-3$ for the coarse grid and $1.6E-5$ for the fine grid. We compare the exact calculation of $\tilde{K}_{\mathbf{x}\mathbf{x}}$ with $s = 1$, $a = 10$, and $\gamma = .5$ to our Nyström-based approximation, where $x_1, \dots, x_{400} \sim \text{Beta}(.5, .5)$ distribution (bottom left). The RMSE was $0.98E-3$. A low-rank approximation using only the top 5 eigenvalues gives the RMSE of $1.6E-2$ (bottom right).

5. Inference

The penalized risk can be readily minimized with gradient descent.² Let $\alpha = [\alpha_1, \dots, \alpha_N]^\top$ and \tilde{K} be the Gram matrix corresponding to \tilde{k} such that $\tilde{K}_{ij} = \tilde{k}(x_i, x_j)$. Then $[f(x_1), \dots, f(x_N)]^\top = \tilde{K}\alpha$ and the gradient of the objective function J from (3.6) is given by

$$\begin{aligned} \nabla_\alpha J &= -\nabla_\alpha \sum_i \log(a f^2(x_i)) + \gamma \nabla_\alpha \|f\|_{\mathcal{H}_{\tilde{k}}}^2 \\ &= -\nabla_\alpha \sum_i \log(a (\sum_j \tilde{k}_{ij} \alpha_j)^2) + \gamma \nabla_\alpha \alpha^\top \tilde{K} \alpha \\ &= -\sum_i \frac{2a (\sum_j \tilde{k}_{ij} \alpha_j) \nabla_\alpha \sum_j \tilde{k}_{ij} \alpha_j}{a (\sum_j \tilde{k}_{ij} \alpha_j)^2} + 2\gamma \tilde{K} \alpha \end{aligned}$$

²While the objective is not convex (we square the function f meaning it is not identified up to sign), in practice we observed very fast convergence, and stable results given random starting points.

$$\begin{aligned}
&= -\sum_i \frac{2\tilde{K}_{\cdot i}}{\sum_j \tilde{k}_{ij}\alpha_j} + 2\gamma\tilde{K}\alpha \\
&= -2\sum_i (\tilde{K}_{\cdot i}/(\tilde{K}\alpha)) + 2\gamma\tilde{K}\alpha
\end{aligned}$$

where $\cdot/$ denotes element-wise division. Computing \tilde{K} requires $\mathcal{O}(N^2)$ time and memory, and each gradient and likelihood computation requires matrix-vector multiplications which are also $\mathcal{O}(N^2)$. Overall, the running time is $\mathcal{O}(qN^2)$ for q iterations of the gradient descent method, where q is usually very small in practice.

5.1. Hyperparameter selection

Analogously to the classical problem of bandwidth selection in kernel intensity estimation (e.g. [11, 6, 7]), some criteria must be adopted in order to select hyperparameters of the kernel k and also γ and a . We suggest crossvalidating on the negative log-likelihood of the inhomogeneous Poisson process (i.e. before we introduced the penalty term) from Eq. (3.2). The difficulty with this approach is that we must deal with the integral $\int_S f^2(u)du$ of the intensity over the domain, which, for our model $f(\cdot) = \sum_{j=1}^N \alpha_j \tilde{k}(x_j, \cdot)$ is generally intractable. As an approximation, we suggest either grid or Monte Carlo integration. Recall that in Section 4.2 we approximated \tilde{k} using a set of locations $\mathbf{u} = (u_1, \dots, u_m)$. We can reuse these points to approximate the integral:

$$\int_S f^2(u)du \approx \frac{1}{m} \sum_i f^2(u_i). \quad (5.1)$$

As $f(u_i) = \tilde{K}_{u_i \mathbf{x}}\alpha$, this approximation is given by $\frac{1}{m}\alpha^\top \tilde{K}_{\mathbf{xu}}\tilde{K}_{\mathbf{ux}}\alpha$.

6. Naïve RKHS model

In this section, we compare the proposed approach, which uses the representer theorem in the transformed kernel \tilde{k} , to the naïve one, where a solution to Eq. (3.3) of the form $f(\cdot) = \sum_{j=1}^N \alpha_j k(x_j, \cdot)$ is sought even though the representer theorem in k need not hold. Despite being theoretically suboptimal, this is a natural model to consider, and it might perform well in practice.

The corresponding optimization problem is:

$$\min_{f \in \mathcal{H}_k} \left\{ -\sum_{i=1}^N \log(af^2(x_i)) + a \int_S f^2(x)dx + \gamma \|f\|_{\mathcal{H}_k}^2 \right\}$$

While the first and the last term are straightforward to calculate for any $f(\cdot) = \sum_j \alpha_j k(x_j, \cdot)$, $\int_S f^2(x)dx$ needs to be estimated. As in the previous

section, we consider a uniform grid or set of sampled points $\mathbf{u} = (u_1, \dots, u_m)$ covering the domain and use approximation

$$\int_S f^2(u) du \approx \frac{1}{m} \sum_i f^2(u_i) = \frac{1}{n} \alpha^\top K_{\mathbf{xu}} K_{\mathbf{ux}} \alpha. \quad (6.1)$$

The optimization problem thus reads:

$$\min_{\alpha \in \mathbb{R}^N} \left\{ - \sum_{i=1}^N \log(a(\alpha^\top K_{\mathbf{x}x_i})^2) + \alpha^\top \left(\frac{a}{n} K_{\mathbf{xu}} K_{\mathbf{ux}} + \gamma K_{\mathbf{xx}} \right) \alpha \right\}. \quad (6.2)$$

As above, the gradient of this objective with respect to α can be readily calculated, and optimized with gradient descent.

7. Experiments

We use cross-validation to choose the hyperparameters in our methods: a , the fixed intensity, γ , the roughness penalty, and the length-scale of the kernel k , minimizing the negative log-likelihood as described in Section 5.1.

To calculate RMSE, we either make predictions at a grid of locations and calculate RMSE compared to the true intensity at that grid or for the high-dimensional synthetic example we pick new uniform samples of locations over the domain and calculate the RMSE at these locations. We used limited memory Broyden–Fletcher–Goldfarb–Shanno (BFGS) in all experiments involving optimization, and found that it converged very quickly and was not sensitive to initial values. Code for our experiments is available at <https://github.com/BigBayes/kernelpoisson>.

7.1. 1-d synthetic example

We generated a synthetic intensity using the Mercer expansion of a squared exponential kernel with lengthscale 0.5, producing a random linear combination of 64 basis functions, weighted with iid draws $\alpha \sim \mathcal{N}(0, 1)$. In Fig. 3 we compare ground truth to estimates made with: our RKHS method with squared exponential kernel, the naïve RKHS approach with squared exponential kernel, and classical kernel intensity estimation with bandwidth selected by crossvalidation. The results are typical of what we observed on 1D and 2D examples: given similar kernel choices, each method performed similarly, and numerically there was not a significant difference in terms of the RMSE compared to the true underlying intensity.

7.2. Environmental datasets

Next we demonstrate our method on a collection of two-dimensional environmental datasets giving the locations of trees. Intensity estimation is a standard




FIG 3. A synthetic dataset, comparing our RKHS method, the naïve model, and kernel smoothing to a synthetic intensity “true”. The rug plot at bottom gives the location of points in the realized point pattern. The RMSE for each method was similar.

first step in both exploratory analysis and modelling of these types of datasets, which were obtained from the R package `spatstat`. We calculated the intensity using various approaches: our proposed RKHS method with \tilde{k} with a squared exponential kernel, the naïve RKHS method with squared exponential kernel, and classical kernel intensity estimation (KIE) with edge correction. Each method used a squared exponential kernel. We report average held-out cross-validated likelihoods in Table 1. With the exception of our method performing better on the Red oak dataset, each method had comparable performance. It is interesting to note, however, that our method does not require any explicit edge correction³, because we are optimizing a likelihood which explicitly takes into account the window. A plot of the resulting intensity surfaces for each method and the effect of edge correction are shown in Fig. 4 for the Black oak dataset.

7.3. High dimensional synthetic examples

We generated random intensity surfaces in the unit hypercube for dimensions $D = 2, \dots, 15$. The intensity was given by a constant multiplied by the square of the sum of 20 multivariate Gaussian pdfs with random means and covariances. The constant was automatically adjusted so that the number of points in the realizations would be held close to constant, in the range 190–210. We expected this to be a relatively simple synthetic example for kernel intensity estimation with a Gaussian kernel in low dimensions, but not in high dimensions. From

³Because no points are observed outside the window S , intensity estimates near the edge are biased downwards [18].




FIG 4. Location of white oak trees in Lansing, Michigan, smoothed with various approaches. Squared exponential kernels are used throughout. Edge correction makes a noticeable difference for classical kernel intensity estimation. Comparing (a) and (c) it is clear that our method is automatically performing edge correction.

each random intensity, we generated two random realizations, and trained our model using 2-fold crossvalidation with these two datasets. We predicted the intensity at a randomly chosen set of points and calculated the mean squared error as compared to the true intensity. For each dimension we repeated this process 100 times comparing kernel intensity estimation, the naive approach, and our approach with \tilde{k} .

Using the same procedure, but a sum of 20 multivariate Student-t distributions with 5 degrees of freedom, random means and covariances, and number of points in the realizations ranging from 10 to 1000, we generated 500 random sur-

TABLE 1
Tree Point Patterns from R Package spatstat

Dataset	Kernel intensity estimation	Naïve approach	Our approach with \tilde{k}
Lansing: Black oak (n = 135)	234	233	227
Hickory (n = 703)	1763	1746	1757
Maple (n = 514)	1239	1228	1233
Misc (n = 105)	179	177	172
New Zealand (n = 86)	119	119	119
Red oak (n = 346)	726	726	739
Redwoods in California (n = 62)	79	84	77
Spruces in Saxonia (n = 134)	215	212	212
Swedish pines (n = 71)	91	89	90
Waka national park (n = 504)	1142	1141	1144
White oak (n = 448)	992	992	996

faces, with dimension $D = 2, \dots, 25$. We expected this to be a difficult synthetic example for all of the methods due to a potential for model misspecification, as we continue to use squared exponential kernels, but the intensity surface is non-Gaussian.

As shown in Fig. 5(a) once we reach dimension 9 and above, our RKHS method with \tilde{k} begins to outperform kernel intensity estimation, where performance is measured as the fraction of times that the MSE is smaller across 100 random datasets for each D . Our method also significantly outperforms the naïve RKHS method as shown in Fig. 5(b). For $D = 15$ the difference between the two RKHS methods is not significant. This could be due to the fact that the number of points in the point pattern remains fixed, so the problem becomes very hard in high dimensions.⁴ As shown in the Fig. 5(c), kernel intensity estimation is almost always better than the naïve RKHS approach, although the difference is not significant in high dimensions.

For the Student-t experiment, as shown in Fig. 5(d)-(f), our RKHS method always outperforms kernel intensity estimation and is better than the naïve method in dimensions below $D = 20$. To assess the amount of improvement, rather than just its statistical significance, we compared the percent improvement in terms of MSE gained by our method versus the competitors, just focusing on $D = 10$ in Fig. 6. On this metric (intuitively, “how much do you expect to improve on average”) our method shows reasonably stable results as compared to KIE, while the performance of the naïve method is revealed to be very variable. Indeed, the standard deviation across the random surfaces for $D = 10$ of the MSE was 56 for both our method and KDE but 166 for the naïve method, perhaps due to overfitting.

⁴Note that our experiments are sensitive to the overall number of points in the synthetic point patterns; since kernel density estimation is a consistent method [33], we should expect kernel intensity estimation to become more accurate as the number of points grows. However, consistency in the sense of classical statistics is not necessarily useful in point processes, because our observations are not iid; the number of points that we observe is in fact part of the dataset since it reflects the underlying intensity.




FIG 5. Three methods were compared: our RKHS method, the naïve RKHS method, and kernel intensity estimation, based on 100 random surfaces for each dimension D in two experimental setups. In (a)-(c), the intensity surface was the squared sum of skewed multivariate Gaussians. In (d)-(f) the surface was a mixture of skewed multivariate Student- t distributions, with 5 degrees of freedom. In (a) and (d): comparison of our RKHS method versus KIE. In (b) and (e): our RKHS method versus the naïve RKHS method. In (c) and (f): comparison of KIE and the naïve RKHS approach. We used squared exponential kernels for all methods. In the Gaussian case (a)-(c), our method significantly outperforms kernel intensity estimation as the dimension increases, and outperforms the naïve method throughout. Kernel intensity estimation almost always outperforms the naïve approach. In the Student- t case (d)-(f), our method always outperforms kernel intensity estimation, and outperforms the naïve approach until very high dimensions. Neither kernel intensity estimation nor the naïve approach are consistently better than each other.

7.4. Computational complexity

Using the synthetic data experimental setup, we evaluated the time complexity of our method with respect to dimensionality d , number of points in the point pattern dataset n , and number of points s used to estimate \hat{k} (Fig. 7b), confirming our theoretical analysis.




FIG 6. To understand the practical (as opposed to statistical) significance of the results in Fig. 5(d)-(f), where we generated random surfaces by squaring the sums of multivariate Student- t distributions, we considered dimension $D = 10$, in which our RKHS method was better than both the naïve method and kernel intensity estimation (KIE) but there was not a significant difference between KIE and the naïve method, and calculated the percent improvement in terms of MSE comparing our RKHS method to KIE (a) and our RKHS method to naïve method (b). The improvement of our method over KIE is apparent, albeit perhaps only modest in this example. Meanwhile, our method is sometimes quite a bit better than the naïve method, which is often very inaccurate.




FIG 7. (a): Run-time of our method versus number of points in the point pattern dataset. (b): Run-time of our method versus number of sample points used to calculate \tilde{k} .

The effect of the dimensionality d was negligible in practice, because the main calculations rely only on an $n \times n$ Gram matrix whose calculation is relatively fast even for high dimensions. Our method's time complexity scales as $\mathcal{O}(n^2)$ as shown in Fig. 7a (but as discussed in Section 4.2, a primal representation is available which would give linear scaling.) where we used $s = 200$ sample points to estimate \tilde{k} . While a small s worked well in practice, we investigated much larger values of s . As shown in Fig. 7b the time complexity scaled as $\mathcal{O}(s^2)$ where the number of points was fixed to be 150; note that we fixed the rank of the eigendecomposition to be 20.




FIG 8. Log-likelihood for various frequencies of a periodic spatiotemporal kernel in a dataset of 18,441 geocoded, date-stamped theft events from Chicago, using our RKHS model. The dataset is for 12 weeks starting January 1, 2004, and the maximum log-likelihood is attained when the frequency is 12, meaning that there is a weekly cycle in the data. Results using the naïve model were less sensible, with a maximum at period 1 (indicating no periodicity), with periods 5 and 2 (corresponding to a 16.8 day cycle and a 42 day cycle) also having high likelihoods.

7.5. Spatiotemporal point pattern of crimes

To demonstrate the ability to use domain specific kernels and learn interpretable hyperparameters, we used 12 weeks (84 days) of geocoded, date-stamped reports of theft obtained from Chicago’s data portal (data.cityofchicago.org) starting January 1, 2004, a relatively large spatiotemporal point pattern consisting of 18,441 events. We used the following kernel: $\exp(-.5s^2/\lambda_s^2)(\exp(-2\sin^2(t\pi p)) + 1)(\exp(-.5t^2/\lambda_t^2))$ which is the product of a separable squared exponential space and decaying periodic time kernel (with frequency p in a time domain normalized to range from 0 to 1) plus a separable squared exponential space and time kernel. After finding reasonable values for the lengthscales and other hyperparameters of \tilde{k} through exploratory data analysis, we used 2-fold cross-validation and calculated average test log-likelihoods for the number of total cycles p in the 84 weeks = 1, 2, ..., 14 or equivalently a period of length 12 weeks (meaning no cycle), 6 weeks, ..., 6 days. These log-likelihoods are shown in Fig. 8; we found that the most likely frequency is 12, or equivalently a period lasting 1 week. This makes sense given known day-of-week effects on crime.

8. Conclusion

We presented a novel approach to inhomogeneous Poisson process intensity estimation using a representer theorem formulation in an appropriately transformed RKHS, providing a scalable approach giving strong performance on synthetic

and real-world datasets. Our approach outperformed the classical baseline of kernel intensity estimation and a naïve approach for which the representer theorem guarantees did not hold. In future work, we will consider marked Poisson processes and other more complex point process models, as well as Bayesian extensions akin to Cox process modeling.

References

- [1] ADAMS, R. P., MURRAY, I. and MACKAY, D. J. (2009). Tractable non-parametric Bayesian inference in Poisson processes with Gaussian process intensities. In *Proceedings of the 26th Annual International Conference on Machine Learning* 9–16. ACM.
- [2] BACH, F. (2015). On the Equivalence between Quadrature Rules and Random Features. *arXiv:1502.06800*. [MR3634888](#)
- [3] BAKER, C. T. H. (1977). *The Numerical Treatment of Integral Equations. Monographs on Numerical Analysis Series*. Oxford: Clarendon Press.
- [4] BARTOSZYNSKI, R., BROWN, B. W., MCBRIDE, C. M. and THOMPSON, J. R. (1981). Some nonparametric techniques for estimating the intensity function of a cancer related nonstationary Poisson process. *The Annals of Statistics* 1050–1060.
- [5] BERLINET, A. and THOMAS-AGNAN, C. (2004). *Reproducing Kernel Hilbert Spaces in Probability and Statistics*. Kluwer. [MR2239907](#)
- [6] BERMAN, M. and DIGGLE, P. (1989). Estimating weighted integrals of the second-order intensity of a spatial point process. *Journal of the Royal Statistical Society. Series B (Methodological)* 81–92.
- [7] BROOKS, M. M. and MARRON, J. S. (1991). Asymptotic optimality of the least-squares cross-validation bandwidth for kernel estimates of intensity functions. *Stochastic Processes and their Applications* **38** 157–165.
- [8] CRESSIE, N. and WIKLE, C. K. (2011). *Statistics for spatio-temporal data* **465**. Wiley. [MR2848400](#)
- [9] CSATÓ, L., OPPER, M. and WINTHER, O. (2001). TAP Gibbs Free Energy, Belief Propagation and Sparsity. In *Advances in Neural Information Processing Systems* 657–663.
- [10] CUNNINGHAM, J. P., SHENOY, K. V. and SAHANI, M. (2008). Fast Gaussian process methods for point process intensity estimation. In *ICML* 192–199. ACM.
- [11] DIGGLE, P. (1985). A kernel method for smoothing point process data. *Applied Statistics* 138–147.
- [12] DIGGLE, P. J., MORAGA, P., ROWLINGSON, B., TAYLOR, B. M. et al. (2013). Spatial and spatio-temporal Log-Gaussian Cox processes: extending the geostatistical paradigm. *Statistical Science* **28** 542–563. [MR3161587](#)
- [13] FASSHAUER, G. E. and MCCOURT, M. J. (2012). Stable evaluation of Gaussian radial basis function interpolants. *SIAM Journal on Scientific Computing* **34** A737–A762. [MR2914302](#)

- [14] FLAXMAN, S., TEH, Y. W. and SEJDINOVIC, D. (2017). Poisson intensity estimation with reproducing kernels. In *Artificial Intelligence and Statistics (AISTATS)* 270–279.
- [15] FLAXMAN, S. R., WILSON, A. G., NEILL, D. B., NICKISCH, H. and SMOLA, A. J. (2015). Fast Kronecker inference in Gaussian processes with non-Gaussian likelihoods. *International Conference on Machine Learning*.
- [16] GILBOA, E., SAATCI, Y. and CUNNINGHAM, J. (2013). Scaling Multidimensional Inference for Structured Gaussian Processes. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **PP** 1–1.
- [17] ILLIAN, J. B., SØRBYE, S. H., RUE, H. et al. (2012). A toolbox for fitting complex spatial point process models using integrated nested Laplace approximation (INLA). *The Annals of Applied Statistics* **6** 1499–1530.
- [18] JONES, M. C. (1993). Simple boundary correction for kernel density estimation. *Statistics and Computing* **3** 135–146.
- [19] KIMELDORF, G. and WAHBA, G. (1971). Some results on Tchebycheffian spline functions. *Journal of Mathematical Analysis and Applications* **33** 82–95.
- [20] KINGMAN, J. F. C. (1993). *Poisson processes*. *Oxford Studies in Probability* **3**. The Clarendon Press Oxford University Press, New York. Oxford Science Publications. [MR1207584 \(94a:60052\)](#)
- [21] KOM SAMO, Y. L. and ROBERTS, S. (2015). Scalable Nonparametric Bayesian Inference on Point Processes with Gaussian Processes. In *ICML* 2227–2236.
- [22] LLOYD, C., GUNTER, T., OSBORNE, M. and ROBERTS, S. (2015). Variational Inference for Gaussian Process Modulated Poisson Processes. In *ICML* 1814–1822.
- [23] MCCULLAGH, P. and MØLLER, J. (2006). The permanental process. *Advances in applied probability* 873–888.
- [24] MØLLER, J., SYVERSVEEN, A. R. and WAAGEPETERSEN, R. P. (1998). Log Gaussian Cox processes. *Scandinavian Journal of Statistics* **25** 451–482.
- [25] MUANDET, K., SRIPERUMBUDUR, B. and SCHÖLKOPF, B. (2014). Kernel Mean Estimation via Spectral Filtering. In *Advances in Neural Information Processing Systems*.
- [26] OATES, C. J. and GIROLAMI, M. A. (2016). Control Functionals for Quasi-Monte Carlo Integration. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics, AISTATS 2016, Cadiz, Spain, May 9–11, 2016* (A. GRETTON and C. C. ROBERT, eds.). *JMLR Workshop and Conference Proceedings* **51** 56–65. JMLR.org.
- [27] RAMLAU-HANSEN, H. (1983). Smoothing Counting Process Intensities by Means of Kernel Functions. *Ann. Statist.* **11** 453–466.
- [28] RASMUSSEN, C. E. and WILLIAMS, C. K. (2006). *Gaussian processes for machine learning*. MIT Press. [MR2514435](#)
- [29] SCHÖLKOPF, B. and SMOLA, A. J. (2002). *Learning with kernels: support vector machines, regularization, optimization and beyond*. MIT Press.

- [30] SILVERMAN, B. W. (1982). On the Estimation of a Probability Density Function by the Maximum Penalized Likelihood Method. *Ann. Statist.* **10** 795–810.
- [31] TEH, Y. W. and RAO, V. (2011). Gaussian process modulated renewal processes. In *Advances in Neural Information Processing Systems* 2474–2482.
- [32] WAHBA, G. (1990). *Spline models for observational data* **59**. Siam.
- [33] WIED, D. and WEISSBACH, R. (2012). Consistency of the kernel density estimator: a survey. *Statistical Papers* **53** 1–21.
- [34] WILLIAMS, C. and SEEGER, M. (2001). Using the Nyström method to speed up kernel machines. In *Advances in Neural Information Processing Systems* 682–688.
- [35] WILSON, A. G., DANN, C. and NICKISCH, H. (2015). Thoughts on Massively Scalable Gaussian Processes. *arXiv:1511.01870*.
- [36] ZHU, H., WILLIAMS, C. K., ROHWER, R. and MORCINIEC, M. (1997). Gaussian regression and optimal finite dimensional linear models. *Technical report*.