

Exact, Fast and Expressive Poisson Point Processes via Squared Neural Families

Russell Tsuchida¹, Cheng Soon Ong^{1,2}, Dino Sejdinovic³

¹Data61-CSIRO

²Australian National University

³University of Adelaide

Abstract

We introduce squared neural Poisson point processes (SNEPPPs) by parameterising the intensity function by the squared norm of a two layer neural network. When the hidden layer is fixed and the second layer has a single neuron, our approach resembles previous uses of squared Gaussian process or kernel methods, but allowing the hidden layer to be learnt allows for additional flexibility. In many cases of interest, the integrated intensity function admits a closed form and can be computed in quadratic time in the number of hidden neurons. We enumerate a far more extensive number of such cases than has previously been discussed. Our approach is more memory and time efficient than naive implementations of squared or exponentiated kernel methods or Gaussian processes. Maximum likelihood and maximum a posteriori estimates in a reparameterisation of the final layer of the intensity function can be obtained by solving a (strongly) convex optimisation problem using projected gradient descent. We demonstrate SNEPPPs on real, and synthetic benchmarks, and provide a software implementation.

1 Introduction

Intensity Functions Our goal is to develop machine learning methods for learning intensity functions, which are central building blocks for probabilistic models of points in sets. An intensity function λ quantifies the number of points $N(A)$ in a set $A \subseteq \mathbb{X}$ divided by the size of the set, as the size of the set approaches zero. In other words, λ is a nonnegative function that may be integrated over the set A against some measure μ to give the expected number of points. We give formal definition of intensity functions and inhomogeneous Poisson point processes (PPPs) in § 2.1. PPPs and their extensions find a wide range of applications in, e.g., seismology (Ogata 1988), ecology (Renner et al. 2015) and popularity prediction in social networks (Mishra, Rizoio, and Xie 2016). The domain \mathbb{X} may be of arbitrary dimension d , although much attention in statistics has been given to the one, two, three and four dimensional settings, which are often identified with the qualifiers “queueing” or “temporal” (1D), “spatial” (2D or 3D) or “spatio-temporal” (3D or 4D). Here we consider the case of arbitrary dimension d .

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: A SNEPPP spatial intensity function fit (left) to 200,000 events using NASA wildfire data (right) (NASA FIRMS 2023). In § 5.4, we discuss a video of a spatio-temporal fit using 100 million events.

Desiderata for Intensity Functions We discuss three useful properties that intensity functions should possess: *expressivity*, *tractable integration* and *tractable optimisation*. Firstly, one needs to choose a function class \mathcal{H} from which to represent the intensity function $\lambda \in \mathcal{H}$. This function class should be *expressive* enough to represent the data at hand. Importantly, λ must be nonnegative, which precludes direct application of some of the standard choices of \mathcal{H} such as RKHSs or certain neural networks. Secondly, computing the expected number of points in a set requires integrating the intensity function, as does measuring the likelihood of a given realisation from a PPP. The integrated intensity function therefore needs to be evaluated often in a typical model fitting and testing pipeline. *Computing or approximating the integrated intensity function* can be challenging for all but simple models and observation windows. Finally, even if the integrated intensity function can be computed, finding plausible members of the function class \mathcal{H} can be difficult. Even point estimation (i.e. choosing a single “best” $\lambda \in \mathcal{H}$) according to some optimisation criterion can be hard if the *optimisation criterion is nonconvex*, which is typically the case for non-trivial hypothesis classes \mathcal{H} .

Contributions We introduce a new class of intensity functions that addresses each of the three desiderata mentioned above. We use the squared Euclidean norm of a two layer neural network to parameterise the intensity function. This represents a highly flexible universal approximating func-

tion class, with even more flexibility if the hidden layer is also trained. Under this function class, the integrated intensity function can be computed exactly for a very wide range of neural network activation functions. In order to compute the integrated intensity function, we extend a recently developed framework for modelling probability densities (Tsuchida, Ong, and Sejdinovic 2023). Unlike naive kernel methods which scale cubically in the number of datapoints in computing the integrated intensity function, our method scales only in the dimensionality of the feature mappings. Finally, we show that if the hidden layer is not trained, the classical maximum likelihood estimate (MLE) and maximum a posteriori (MAP) estimate can be found by minimising a (strongly) convex function over a convex set. Such estimates can be found using projected gradient descent. The performance and features of our method are illustrated on some real and synthetic benchmark datasets. We empirically demonstrate the efficacy and efficiency of our open source method on a large scale case study of wildfire data from NASA with about 100 million events.

2 Background

2.1 Poisson Point Processes

We briefly present definitions for PPP likelihoods and intensity functions, referring the reader to classical sources for more details (Cressie 1993, § 8.5) (Baddeley 2007, § 4.3). Let $\mathbb{X} \subseteq \mathbb{R}^d$ with Borel sigma algebra \mathcal{F} . Let ν be a nonnegative measure such that $\nu(A) < \infty$ for any compact $A \in \mathcal{F}$. We call ν the intensity measure. An (inhomogeneous) *Poisson point process* (PPP) N on \mathbb{X} specifies that for any such A , the number $N(A)$ of events within A is governed by a Poisson distribution with a parameter equal to the evaluated intensity measure $\nu(A) > 0$,

$$\mathbb{P}(N(A) = N) = \frac{(\nu(A))^N \exp(-\nu(A))}{N!}. \quad (1)$$

Let \mathbf{x}_i denote one of the $N(A)$ points in A . Conditional on the event $N(A) = N$, the ordered N -tuple $(\mathbf{x}_1, \dots, \mathbf{x}_N)$, is distributed independently and identically with probability distribution $\nu(\cdot)/\nu(A)$. If ν is absolutely continuous with respect to some *base measure* μ and $\lambda = \frac{d\nu}{d\mu} : \mathbb{X} \rightarrow [0, \infty)$ is the Radon-Nikodym derivative called the *intensity function*, the random vector $(\mathbf{x}_1, \dots, \mathbf{x}_N)$ admits a conditional density function with respect to μ

$$p(\mathbf{x}_1, \dots, \mathbf{x}_N | N(A) = N, \lambda) = \frac{\prod_{i=1}^N \lambda(\mathbf{x}_i)}{\left(\int_A \lambda(\mathbf{u})\mu(d\mathbf{u})\right)^N}. \quad (2)$$

Finally, combining (1) and (2), the *joint likelihood* of the events and the number $N(A)$ of the events is

$$p(\mathbf{x}_1, \dots, \mathbf{x}_N, N | \lambda) = \frac{1}{N!} \prod_{i=1}^N \lambda(\mathbf{x}_i) \exp\left(-\int_A \lambda(\mathbf{u})\mu(d\mathbf{u})\right). \quad (3)$$

2.2 Squared Neural Families

One way to construct a probability density function (with respect to some base measure) is to divide a nonnegative

function by its integral (with respect to a base measure). The challenge in this method is computing the integral, also called the normalising constant. Tsuchida, Ong, and Sejdinovic (2023) construct families of probability distributions by normalising the squared Euclidean norm of two layer neural networks. They show that such a procedure often results in a closed-form normalising constant in terms of a neural network kernel (NNK). We will use an analogous result in the context of Poisson point processes instead of probability distributions, see Identity 1.

3 Flexible Intensity with a Tractable Integral

3.1 Squared Neural Poisson Point Process

Our Parameterisation of Intensity Function λ We investigate the use of finite features for the intensity. We begin by modelling the intensity function as the product of a hyperparameter $0 < \alpha < \infty$ and the squared Euclidean norm of a finite feature model with *fixed* features ψ ,

$$\lambda(\mathbf{x}) = \alpha \|\mathbf{V}\psi(\mathbf{x})\|_2^2, \quad (4a)$$

where $\mathbf{V} \in \mathbb{R}^{m \times n}$ and $\psi : \mathbb{R}^d \rightarrow \mathbb{R}^n$. Examples of such ψ include radial basis function networks with fixed hidden layer parameters (Broomhead and Lowe 1988, which are universal approximators), random Fourier features (Rahimi and Recht 2007, which can approximate universal kernels) and finite dimensional kernel feature mappings (Shawe-Taylor and Cristianini 2004, definition 9.1, 9.5, 9.6 and others). Shallow and deep neural networks with fixed (but possibly randomly sampled) hidden parameters also fall under (4a) (Cho and Saul 2009, for example, (8) and (9)). Fixed features are previously used in the context of PPPs when $m = 1$ by Walder and Bishop (2017), however our extension to $m \geq 1$ not only increases flexibility but also allows us to prove optimisation properties (see Proposition 2).

We then set the features ψ to be a hidden neural network layer with learnable hidden parameters. Suppose the hidden activation function is σ (applied element-wise), and input data $\mathbf{x} \in \mathbb{R}^d$ is preprocessed by passing through a warping function $\mathbf{t} : \mathbb{R}^d \rightarrow \mathbb{R}^D$. We set the intensity function λ to be equal to the squared Euclidean norm of the neural network's output multiplied by a hyperparameter $0 < \alpha < \infty$,

$$\lambda(\mathbf{x}) = \alpha \|\mathbf{V}\psi(\mathbf{x})\|_2^2, \quad \psi(\mathbf{x}) = \sigma(\mathbf{W}\mathbf{t}(\mathbf{x}) + \mathbf{b}), \quad (4b)$$

where we introduced additional parameters to be learned: hidden layer weights and biases $\mathbf{W} \in \mathbb{R}^{n \times D}$ and $\mathbf{b} \in \mathbb{R}^n$. For fixed (\mathbf{W}, \mathbf{b}) , (4b) is a special case of (4a).

Extension to Product Spaces We provide one additional extension of (4a) and (4b) in Appendix B, which is particularly useful when the space \mathbb{X} naturally decomposes into a Cartesian product (for example, as one might be interested in when studying spatio-temporal models). For the purposes of notational simplicity, we discuss our central ideas in terms of models (4a) and (4b), noting that these ideas naturally extend to the model Appendix B. Our large scale case study in § 5 uses the setting in Appendix B.

Integrated Intensity Function One challenge is in computing the integrated intensity function,

$$\Lambda = \int_A \lambda(\mathbf{x})\mu(d\mathbf{x}) = \alpha \int_A \|\mathbf{V}\psi(\mathbf{x})\|_2^2 \mu(d\mathbf{x}), \quad A \subseteq \mathbb{X},$$

focusing mostly on the setting where $A = \mathbb{X}$ (i.e. the process is observed over a full window)¹. We note that a setting amounting to full observation window is also considered by other machine learning approaches (Flaxman, Teh, and Sejdinovic 2017; Walder and Bishop 2017). Much of the machinery we develop here mirrors the use of squared neural families for density estimation (Tsuchida, Ong, and Sejdinovic 2023), however we introduce a number of new kernels and results specific to the setting of intensity estimation and aim to give a self-contained presentation.

Regularisation Note that either (4a) or (4b) implies that the conditional likelihood (2) does not depend on α , but the joint likelihood (3) does depend on α . The inclusion of α , borrowed from Flaxman, Teh, and Sejdinovic (2017, see remark 2), allows us to explicitly control the overall scale of the intensity independently from other measures of complexity of the intensity (such as the parameter Euclidean norm, as might typically be done using ℓ^2 regularisation). This can be helpful for model fitting. Intuitively speaking, if $\|\mathbf{V}\psi(\mathbf{x})\|_2^2$ is initialised to have the value 1, and we observe N samples over a volume of $\mu(A)$, α should be set to roughly $N/\mu(A)$ to encourage the initialised model to predict the correct number of training observations.

Bayesian Extensions One may place a Bayesian prior over parameters of λ and compute with the resulting posterior. For example, if one were to employ a Gaussian prior over \mathbf{V} , then $\mathbf{V}\psi(\cdot)$ would be a vector-valued Gaussian process prior, extending previously considered models (Walder and Bishop 2017). While we do not consider such Bayesian models in detail here, we mention related work in § 4.3. We compute MAP estimates in Proposition 2.

3.2 Closed Form Integration

The following identity is proven in Appendix A. It is a variation of the result in Tsuchida, Ong, and Sejdinovic (2023), the difference being that here λ is an intensity function rather than an unnormalised probability density function.

Proposition 1. *Under (4a) (and therefore (4b)), the intensity function and integrated intensity function are*

$$\lambda(\mathbf{x}) = \alpha \text{Tr} \left(\mathbf{V}^\top \mathbf{V} \widetilde{\mathbf{K}}(\mathbf{x}) \right) \quad \text{and} \quad \Lambda = \alpha \text{Tr} \left(\mathbf{V}^\top \mathbf{V} \mathbf{K} \right),$$

where $\widetilde{\mathbf{K}}(\mathbf{x})$ is the PSD matrix with ij th entry $\psi_i(\mathbf{x})\psi_j(\mathbf{x})$ and \mathbf{K} is the PSD matrix with ij th entry $\kappa(i, j)$,

$$\kappa(i, j) = \int_{\mathbb{X}} \psi_i(\mathbf{x})\psi_j(\mathbf{x})\mu(d\mathbf{x}).$$

¹We are also able to handle cases where A is “full” in a subset of dimensions. For example, in the spatio-temporal setting described in § 3.4, we integrate over a full spatial domain and a partially observed temporal domain.

Neural Network Kernels In the special case of (4b), we identify κ with a broad family of kernels called neural network kernels (NNKs). We let $\Theta = (\mathbf{W}, \mathbf{b}) \in \mathbb{R}^{n \times (D+1)}$, and denote the i th row of Θ by $\theta_i = (\mathbf{w}_i, b_i) \in \mathbb{R}^{D+1}$. In this case, it is more natural to consider the kernel as operating over the space of parameters θ_i, θ_j indexed by i, j . We define the NNK $k_{\sigma, \mathbf{t}, \mu}(\theta_i, \theta_j)$ to be such a kernel,

$$\begin{aligned} \kappa(i, j) &= \int_{\mathbb{X}} \sigma(\mathbf{w}_i^\top \mathbf{t}(\mathbf{x}) + b_i) \sigma(\mathbf{w}_j^\top \mathbf{t}(\mathbf{x}) + b_j) \mu(d\mathbf{x}) \\ &=: k_{\sigma, \mathbf{t}, \mu}(\theta_i, \theta_j). \end{aligned} \tag{5}$$

Closed Forms and Computational Complexity If we have a closed form for the kernel κ , we may compute $\Lambda = \Lambda(\mathbf{V}, \Theta)$ in $\mathcal{O}(m^2n + n^2)$. Closed forms for the NNK $k_{\sigma, \mathbf{t}, \mu}(\theta_i, \theta_j)$ for various choices of σ, \mathbf{t} and μ are available; see Han et al. (2022, table 1) and Tsuchida, Ong, and Sejdinovic (2023, table 1) and references therein. These NNKs allow the use of a variety of activation functions, including error functions, ReLU, Leaky ReLU, GELU (Hendrycks and Gimpel 2016), snake (Ziyin, Hartwig, and Ueda 2020), Gaussians, polynomials, and Gabor functions. Each evaluation of such NNK typically incurs a cost of $\mathcal{O}(d)$, depending on \mathbf{w}_j and \mathbf{w}_j only through a Euclidean inner product between θ_i and θ_j . This linear complexity in input dimension d is extremely favourable in light of the exponential complexity that would be implied by a naive numerical integration technique. It is also favourable compared with kernel methods that use the representer theorem, which scale cubically in the number of datapoints N .

3.3 Novel Examples of NNKs

In § 3.2, we reduced calculation of the integrated intensity function to calculation of the NNK $k_{\sigma, \mathbf{t}, \mu}(\theta_i, \theta_j)$, and pointed to a vast literature that provides closed-forms for NNKs. Below we discuss two other examples of NNKs not mentioned in the aforementioned works.

Exponential Family Type Take an identity warping function $\mathbf{t}(\mathbf{x}) = \mathbf{x} = \text{Id}(\mathbf{x})$, a hyperrectangular domain $\mathbb{X} = [a, b]^d$, an exponential activation function $\sigma = \exp$ and the base measure to be Lebesgue on $[a, b]^d$. The NNK can then be computed in closed-form,

$$\begin{aligned} k_{\exp, \text{Id}, d\mathbf{x}}(\theta_i, \theta_j) &= e^{b_i + b_j} \int_{[a, b]^d} \exp(\mathbf{x}^\top (\mathbf{w}_i + \mathbf{w}_j)) d\mathbf{x} \\ &= e^{b_i + b_j} \prod_{r=1}^d \left(\frac{\exp(x_r (w_{ir} + w_{jr}))}{w_{ir} + w_{jr}} \Big|_a^b \right). \end{aligned}$$

In this case, the NNK $k_{\exp, \text{Id}, d\mathbf{x}}(\theta_i, \theta_j)$ also defines a partition function of an exponential family (Wainwright, Jordan et al. 2008; Nielsen and Garcia 2009) with identity sufficient statistic supported on $[a, b]^d$. In a similar vein, other exponential families can define analogous SNEPPPs.

Absolutely Homogeneous Activations on the Sphere A function $\sigma^{(p)}$ is said to be absolutely p -homogeneous if for some $p > 0$, for all a and z , $\sigma^{(p)}(|a|z) = |a|^p \sigma^{(p)}(z)$. Examples of absolutely p -homogeneous functions include

monomials, the Heaviside step function, the ReLU (raised to any positive power), and the Leaky ReLU.

If the bias parameter $\mathbf{b} = \mathbf{0}$, the activation $\sigma = \sigma^{(p)}$ is p -homogeneous, and the warping function $\mathbf{t}(\mathbf{x}) = \mathbf{x}$ is the identity, we may translate between NNKs with uniform base measure on the sphere and NNKs with Gaussian base measure on \mathbb{R}^d , as follows. It is well-known that if \mathbf{x} is zero mean isotropic Gaussian, then $\mathbf{x}/\|\mathbf{x}\|$ and $\|\mathbf{x}\|$ are independent. Using this fact and the absolute homogeneity property,

$$\begin{aligned} k_{\sigma^{(p)}, \text{Id}, \Phi}(\boldsymbol{\theta}_i, \boldsymbol{\theta}_j) &= \mathbb{E}[\sigma^{(p)}(\mathbf{w}_i^\top \mathbf{x})\sigma^{(p)}(\mathbf{w}_j^\top \mathbf{x})] \\ &= \mathbb{E}[\|\mathbf{x}\|^{2p}\sigma^{(p)}\left(\mathbf{w}_i^\top \frac{\mathbf{x}}{\|\mathbf{x}\|}\right)\sigma^{(p)}\left(\mathbf{w}_j^\top \frac{\mathbf{x}}{\|\mathbf{x}\|}\right)] \\ &= \mathbb{E}[\|\mathbf{x}\|^{2p}]k_{\sigma^{(p)}, \text{Id}, U(\mathbb{S}^{d-1})}(\boldsymbol{\theta}_i, \boldsymbol{\theta}_j). \end{aligned} \quad (6)$$

Finally, $\mathbb{E}[\|\mathbf{x}\|^{2p}]$ is just the p th moment of a Chi-squared random variable, which is available in closed-form.

As an example, consider the family of activation functions of the form $\sigma(z) = \Theta(z)z^p$, where Θ is the Heaviside step function. The ReLU is obtained when $p = 1$. The arc-cosine kernel of order p is known in closed form if $\mathbf{b} = \mathbf{0}$ (Cho and Saul 2009). In particular, when $p = 1$,

$$k_{(\cdot)\Theta(\cdot), \text{Id}, \Phi}(\boldsymbol{\theta}_i, \boldsymbol{\theta}_j) = \frac{\|\mathbf{w}_i\|\|\mathbf{w}_j\|}{2\pi} \left(\sin \gamma + (\pi - \gamma) \cos \gamma \right),$$

where $\gamma = \cos^{-1} \frac{\mathbf{w}_i^\top \mathbf{w}_j}{\|\mathbf{w}_i\|\|\mathbf{w}_j\|}$ is the angle between \mathbf{w}_i and \mathbf{w}_j . Using (6), we may write the NNK on the sphere,

$$k_{(\cdot)^p\Theta(\cdot), \text{Id}, U(\mathbb{S}^{d-1})}(\boldsymbol{\theta}_i, \boldsymbol{\theta}_j) = \frac{k_{(\cdot)^p\Theta(\cdot), \text{Id}, \Phi}(\boldsymbol{\theta}_i, \boldsymbol{\theta}_j)\Gamma(d/2)}{2^p\Gamma(p + d/2)}, \quad (7)$$

where Γ is the gamma function. We use this newly derived kernel to build intensity functions on the sphere using neural networks with ReLU activations ($p = 1$) in § 5.4.

3.4 Extension to Product Spaces

It is sometimes helpful to view the domain \mathbb{X} as a Cartesian product over multiple sets. For example, we wish to model an intensity function that varies over the surface of the Earth as well as time, we might consider $\mathbb{X} = \mathbb{S}^2 \times \mathbb{T}$, where \mathbb{S}^2 is the unit sphere and \mathbb{T} is some discrete or continuous set indexing time. Closed-form integrated intensity functions then follow as a special case of the reasoning in earlier sections. A more detailed discussion of how to construct tractable product space intensities is given in Appendix B.1.

Example We write $\mathbb{X} = \mathbb{Y} \times \mathbb{T}$, where \mathbb{Y} might represent a spatial domain and \mathbb{T} might represent a temporal domain. We write $\mathbf{x} = (\mathbf{y}, \boldsymbol{\tau})$ and $\boldsymbol{\tau} \in \mathbb{T}$ and $\mathbf{y} \in \mathbb{Y}$. Let \odot denote the Hadamard product and consider the model

$$\begin{aligned} \lambda(\mathbf{x}) &= \alpha \|\mathbf{V}(\boldsymbol{\psi}_1(\mathbf{y}) \odot \boldsymbol{\psi}_2(\boldsymbol{\tau}))\|_2^2 \\ &= \alpha \text{Tr} \left(\mathbf{V}^\top \mathbf{V}(\widetilde{\mathbf{K}}_1(\mathbf{y}) \odot \widetilde{\mathbf{K}}_2(\boldsymbol{\tau})) \right), \quad \text{where} \end{aligned}$$

$\widetilde{\mathbf{K}}_1(\mathbf{y}) = \boldsymbol{\psi}_1(\mathbf{y})\boldsymbol{\psi}_1(\mathbf{y})^\top$ and $\widetilde{\mathbf{K}}_2(\boldsymbol{\tau}) = \boldsymbol{\psi}_2(\boldsymbol{\tau})\boldsymbol{\psi}_2(\boldsymbol{\tau})^\top$. If the base measure μ decomposes as $\mu(\cdot) = \mu_1(\cdot)\mu_2(\cdot)$, the

integrated intensity function is then

$$\Lambda = \alpha \text{Tr} \left(\mathbf{V}^\top \mathbf{V}(\mathbf{K}_1 \odot \mathbf{K}_2) \right),$$

where the ij th entries of \mathbf{K}_1 and \mathbf{K}_2 are respectively

$$\int_{\mathbb{Y}} \psi_{1i}(\mathbf{y})\psi_{1j}(\mathbf{y})d\mu_1(\mathbf{y}) \quad \text{and} \quad \int_{\mathbb{T}} \psi_{2i}(\boldsymbol{\tau})\psi_{2j}(\boldsymbol{\tau})d\mu_2(\boldsymbol{\tau}).$$

These kernels are tractable under the same settings (4a) and (4b). We consider the special case where $d = 1$, \mathbf{t} is the identity, $\psi_2(\boldsymbol{\tau}) = \text{ReLU}(w\boldsymbol{\tau} + b)$, μ_2 is Lebesgue and $\mathbb{T} = [T_1, T_2]$ is an interval in Appendix C.

4 Properties of SNEPPP

4.1 Optimisation Properties

A common procedure for finding point estimates for λ is via MLE. A single realisation of a PPP observed over a window A is $\{\mathbf{x}_i\}_{i=1}^N$, where N is itself a draw from a Poisson distribution according to (1). One may fit a PPP by observing a single realisation over A and maximising the likelihood (3) with respect to (the parameters of) λ , or alternatively by observing multiple independently drawn realisations over possibly different windows A_1, \dots, A_r and then maximising the product of likelihoods. We focus here on the setting of a single realisation for notational simplicity and to follow previous works (Flaxman, Teh, and Sejdinovic 2017; Walder and Bishop 2017), but our discussion extends to multiple realisations. It is likely that many of our insights extend beyond MLE (e.g. Bayesian inference schemes that utilise gradient-descent like algorithms). We assume that $\mathbf{V}^\top \mathbf{V}$ is positive definite, with smallest eigenvalue greater than or equal to $\epsilon_1 > 0$. That is, we set $\mathbf{V}^\top \mathbf{V} = \mathbf{M} + \epsilon_1 \mathbf{I}$ and reparameterise the model in terms of a PSD matrix \mathbf{M} .

Convexity Compare the two NLLs, under (2) and (3) respectively, and note that the second NLL is strictly convex in \mathbf{M} , being the sum of a linear penalty term and a strictly convex model fit term,

$$-\log p(\mathbf{x}_1, \dots, \mathbf{x}_N, N|\lambda) = \quad (8)$$

$$\alpha \text{Tr} \left((\mathbf{M} + \epsilon_1 \mathbf{I}) \mathbf{K} \right) - \sum_{i=1}^N \log \text{Tr} \left((\mathbf{M} + \epsilon_1 \mathbf{I}) \widetilde{\mathbf{K}}(\mathbf{x}_i) \right)$$

In contrast, the NLL of (2), which is the objective of an analogous density estimation procedure, is in general non-convex. We expect that most easy to implement constrained optimisation routines perform well in optimising (8) when the hidden layer is fixed. In fact, projected gradient descent (PGD) with sufficiently small constant step size η converges. PGD amounts to alternating gradient and projection iterations, beginning with initial guess \mathbf{M}_1 ,

$$\mathbf{M}_{t+1/2} = \mathbf{M}_t - \eta \nabla C(\mathbf{M}_t) \quad \text{Gradient update}$$

$$\mathbf{M}_{t+1} = \mathbf{Q}_{t+1/2} \mathbf{D}_{t+1/2}^+ \mathbf{Q}_{t+1/2}^\top \quad \text{Projection,}$$

where $\mathbf{Q}\mathbf{D}\mathbf{Q}^\top$ is an eigendecomposition and \mathbf{D}^+ is \mathbf{D} with negative entries set to zero². We may optimise the likelihood or the posterior under a particular choice of prior. This

²Frobenius norm projection of a symmetric matrix onto the space of PSD matrices sets the negative eigenvalues of the symmetric matrix to zero (Boyd and Vandenberghe 2004, § 8.1.1).

prior is well-defined and amounts to a Gaussian distribution projected to the space of PSD matrices over M — see Appendix A for more details.

Proposition 2. *Let $\mathbb{M}_+^n = \{M \in \mathbb{R}^{n \times n} \mid M = M^\top, M \succeq 0\}$ denote the convex cone of PSD matrices. Let $\epsilon_1 \geq 0$. The NLL (8) is strictly convex in M over \mathbb{M}_+^n . Let $\epsilon_1 > 0$ and $\epsilon_2 \geq 0, \beta = \epsilon_2 + 1/\epsilon_1^2$ and consider the MLE or MAP estimate*

$$M^* = \operatorname{argmin}_{M \in \mathbb{M}_+^n} \overbrace{-\frac{1}{N} \log p(\mathbf{x}_1, \dots, \mathbf{x}_N, N | \lambda)}^{\triangleq C(M)} + \frac{\epsilon_2}{2} \|M\|_F^2.$$

Choose the learning rate $\eta = 1/\beta$. Then PGD satisfies

$$\frac{C(M_{t+1}) - C(M^*) \leq 3\beta \|M_1 - M^*\|_F^2 + C(M_1) - C(M^*)}{t + 1}.$$

Suppose additionally that $\epsilon_2 > 0$. Then PGD satisfies

$$\|M_{t+1} - M^*\|_F^2 \leq \exp\left(-\frac{t\epsilon_2}{2\beta}\right) \|M_1 - M^*\|_F^2.$$

It is straightforward take a Laplace approximation about the MAP. We leave this for future work.

4.2 Special Cases

Log-linear Models The log-linear model $\lambda(\mathbf{x}) = \exp(\mathbf{w}^\top \mathbf{x})$ is a widely used classical statistical model that is easy to integrate (see § 3.3 for an example). We obtain this model as a special case of SNEPPP when $\sigma = \exp, m = n = 1$ and the bias \mathbf{b} is constrained to be zero. We refer to this model as “log-linear”, and use it as a baseline for experiments in § 5. Our implementation differs slightly to that of a typical statistical implementation; we use Adam to optimise parameters rather than second order methods.

Mixture Models The model $\lambda(\mathbf{x}) = \sum_{i=1}^n v_i^2 \exp(\mathbf{w}_i^\top \mathbf{x})$ is also easy to integrate, and is obtained as a special case of SNEPPP when $\sigma = \exp, m = n \geq 1$, the bias \mathbf{b} is constrained to be zero, and the readout parameter \mathbf{V} is constrained to be diagonal. Borrowing terminology from i.i.d. density modelling, we refer to this model as “log-linear mixture”, and use it as a baseline for experiments in § 5.

4.3 Related Work

Kernel Methods for Densities Our method uses the same machinery for closed-form integration as squared neural families (SNEFYs) (Tsuchida, Ong, and Sejdinovic 2023) do for probability density models. SNEFYs represent the unnormalised probability density as the squared norm of a single layer neural network. SNEFYs in turn can be seen as neural network variants of kernel-based methods for representing probability densities (Marteau-Ferey, Bach, and Rudi 2020; Rudi and Ciliberto 2021). These kernel methods amount to using something analogous to the square of a linear combination of elements of an RKHS for the unnormalised probability density.

Kernel Methods and Gaussian Processes for Intensities Flaxman, Teh, and Sejdinovic (2017) use a squared element of an RKHS as the intensity function. A Bayesian extension is also available (Walder and Bishop 2017), where the intensity function is a squared Gaussian process (GP), and a Laplace approximation to the posterior is obtained. An earlier work uses variational inference to obtain an approximate posterior under a squared GP model (Lloyd et al. 2015). Random Fourier features are exploited by (Sellier and Dellaportas 2023) to build routines involving generalised stationary kernels. We note that it would be relatively straightforward to compute Laplace approximations to the posterior by considering the Hessian about our MAP estimates. We leave empirical evaluation of such models for future work.

Optimisation Properties In all of the PPP works described above, the intensity function is a squared GP with a finite number of features, which is a special case of (4a) when $m = 1$ under a Gaussian prior over \mathbf{V} . To the best of our knowledge, we are the first to note that allowing $m = n$ and constraining $\mathbf{V}^\top \mathbf{V}$ to be positive definite allows for convex optimisation over $\mathbf{V}^\top \mathbf{V}$ for MLE or MAP estimates via Proposition 4.1. In previous works, optimisation over \mathbf{V} with $m = 1$ is nonconvex (see for example Flaxman, Teh, and Sejdinovic (2017, footnote 2)). Our new result is worth highlighting, given that in the density estimation setting for similar models, MLE or MAP estimates are also not convex, but other surrogate losses are used in their place (Rudi and Ciliberto 2021, Theorem 7). In contrast, our method allows for principled MLE and MAP estimation.

Augmented Permanent Point Processes Kim, Asami, and Toda (2022) introduce augmented permanent point processes (APP), in which the intensity function is a squared GP of a function (called the covariate map) of the input variable \mathbf{x} , rather than \mathbf{x} itself. The integrated intensity function is still computed with respect to \mathbf{x} . In the setting of APP, during training, the input \mathbf{x} and the covariate of \mathbf{x} are jointly observed, which differs to our setting in which only \mathbf{x} is observed. Direct comparisons are therefore not possible, however we note that the implementation of Kim, Asami, and Toda (2022) are able to handle the non-augmented setting by choosing an identity covariate map. In this setting, the resulting model is equivalent to that of Walder and Bishop (2017), and the posterior mode is the squared RKHS model of Flaxman, Teh, and Sejdinovic (2017). The MAP optimisation problem involves a so-called equivalent kernel (Williams and Rasmussen 2006; Menon and Williamson 2018), whose calculation can be performed in one of three methods: naive, random Fourier mappings (RFM) (Rahimi and Recht 2007) and Nyström (Williams and Seeger 2000).

Tractable Domains of Integration None of the related works on intensity estimation explicitly describe closed-form integration on interesting domains such as the hypersphere \mathbb{S}^{d-1} . In contrast, we may use the newly derived NNK (7) or previously derived NNKs (Tsuchida, Ong, and Sejdinovic 2023, Kernel 3) to compute integrated intensity functions on the hypersphere.

	NLL (exact)	NLL (MC)	Count percent error	Time (seconds)
Log-linear	-0.03 ± 0.01	-0.02 ± 0.01	0.03 ± 0.03	1.79 ± 0.31
Log-linear mixture	-0.07 ± 0.02	-0.07 ± 0.02	0.02 ± 0.02	2.28 ± 0.20
SNEPPP	$-0.56 \pm 0.01^*$	$-0.62 \pm 0.01^*$	0.04 ± 0.03	6.20 ± 0.18
RFM	n/a	-0.38 ± 0.03	0.07 ± 0.03	7.80 ± 0.73
Nyström	n/a	-0.52 ± 0.02	0.09 ± 0.03	4.74 ± 0.33
Naive	n/a	-0.51 ± 0.03	0.09 ± 0.04	18.48 ± 0.48
Log-linear	-0.04 ± 0.01	-0.04 ± 0.01	0.04 ± 0.02	2.38 ± 0.14
Log-linear mixture	-0.08 ± 0.01	-0.08 ± 0.01	0.03 ± 0.02	2.45 ± 0.18
SNEPPP	$-0.78 \pm 0.03^*$	-0.84 ± 0.03	0.03 ± 0.02	7.07 ± 0.17
RFM	n/a	-0.23 ± 0.02	0.07 ± 0.03	4.17 ± 0.31
Nyström	n/a	-0.35 ± 0.03	0.09 ± 0.04	4.92 ± 0.48
Naive	n/a	-0.36 ± 0.03	0.11 ± 0.05	25.88 ± 1.10

Table 1: Shown are means \pm sample standard deviations over 100 randomly generated synthetic datasets and parameter initialisations. Asterisked and bold values are significantly better than closest competitor according to a two sample t-test. Bold numbers indicate values which are not significantly different to the best. Some methods do not provide exact integrated intensity functions or NLL, so we resort to Monte Carlo (MC) estimation. (Top) Bei real data experiment results. (Bottom) clmfires real data experiment results. Results for the Copper dataset are given in Table 3, Appendix D.

5 Experiments

5.1 Metrics for Intensity Functions

There are various ways in which one may assess the quality of an intensity function λ with respect to held-out test data $\{\tilde{x}_i\}_{i=1}^M$ observed over a potentially partial window $A \subseteq \mathbb{X}$. There are some important distinctions between assessing the quality of density models (2) and PPP models (3). We briefly describe three ways in which PPPs can be evaluated.

Test Log Likelihood We evaluate the test log likelihood on a held out piece of data that is independent of the data used for training in two ways. Firstly, we may exploit the *complete independence of disjoint subsets* of PPPs. Suppose for example that we observe and fit an intensity function to wildfires observed over the entire Earth for years 2003 to 2023, excluding the year 2018. Each of $N(\mathbb{S}^2 \times \{\tau\})$ for distinct $\tau \in \{2003, \dots, 2023\}$ are mutually independent. We may therefore test on data from 2018. The same idea applies to spatially disjoint observations. Secondly, we may artificially but exactly obtain multiple independent realisations from a single realisation of data by a process called *splitting or thinning*, see Appendix D.

Predictions of Counts Suppose we fit a PPP using data from a window that does not include $A \subseteq \mathbb{X}$. We reserve observations over the window A for testing. We predict the number of events within A and compare the result with such test observations. Alternatively, we may predict the number of events in any window and compare the result with an independent realisation over the same window.

Ground Truth Intensities A final option is to compare the fitted intensity function λ against a ground truth intensity function λ_{GT} . Such comparisons can only be made when using synthetic data. We consider the root mean squared error (RMSE) over a random sample $\{z_i\}_{i=1}^P$ taken uniformly from \mathbb{X} , $\text{RMSE}^2 = \frac{1}{P} \sum_{i=1}^P (\lambda(z_i) - \lambda_{\text{GT}}(z_i))^2 \approx \frac{1}{|\mathbb{X}|} \int_{\mathbb{X}} (\lambda(x) - \lambda_{\text{GT}}(x))^2 dx$.

5.2 Baselines

Log-linear and Mixture Models We compare against log-linear and log-mixture models, which are implemented as a special case of SNEPPP (see § 4.2).

Kernel and GP Methods We use an open source TensorFlow library (Kim, Asami, and Toda 2022) for GP baselines that compute the equivalent kernel using Naive, Nyström and Random Fourier Methods (RFM). The computational complexity scales cubically in the number of data points in the Naive case, and cubically in the number of features in the Nyström or RFM case. While exact NLL is in theory possible using the RFM and Nyström methods using the same techniques we discuss here, they are not implemented in the software package that we use. We do not compare with APP which covers a different problem setting (see § 4.3).

5.3 Evaluation on Synthetic and Real Data

We measure the NLL (using a thinning process), RMSE and computation time (in seconds) for an existing synthetic setting (Baddeley et al. 2012; Kim, Asami, and Toda 2022). Full details and results are given in Appendix D. We also perform benchmarks on three real datasets — bei, copper and clmfires — also considered by Kim, Asami, and Toda (2022). See Appendix D for details. Results are summarised in Tables 1 and 3.

5.4 Case Study on Wildfires

We now turn our attention to the spatio-temporal modelling of the incidence of wildfires around planet Earth. We use a massive freely available dataset (NASA FIRMS 2023). In the November and December months of the year 2000, over 200,000 datapoints are observed — we call this restricted dataset the spatial dataset. Between the years 2000 and 2022 over 98 million datapoints are observed — we call this complete dataset the spatio-temporal dataset. In this setting, having an observation domain over the sphere (and potentially the time domain) and a huge number of events, we are not




Figure 2: (Left) Using a naive kernel method or squared GP requires solving a linear system, incurring $\mathcal{O}(N^3)$ time. In contrast, finite feature methods, such as SNEPPP, scale linearly in the number of datapoints. In this setting, the naive method runs out of memory beyond 5.5% of data usage, but even if we had sufficient memory to perform the required calculation, it would take roughly 88 days to compute using 100% data. In contrast, SNEPPP takes less than 8 hours. (Right) Training SNEPPPs using more data decreases the test NLL, after appropriately accounting for thinning.

aware of any publicly available software tools that we can quantitatively compare with. Instead, we highlight attractive features of applying our model to big data, and compare performance against the special cases described in § 4.2.

Finite Feature but not Naive Methods Scale to Big Data

We apply thinning to the process and compare the training time for a SNEPPP model and a Naive permanent point process model on the spatial dataset. We use the squared exponential kernel implemented by Kim, Asami, and Toda (2022), which assumes a domain of $\mathbb{X} = \mathbb{R}^d$ instead of $\mathbb{X} = \mathbb{S}^{d-1}$, noting that we are only interested in computation time. Such a model stands in as a good proxy in terms of computation time for any future developments of naive equivalent kernels for $\mathbb{X} = \mathbb{S}^{d-1}$. We find that the SNEPPP model is able to use all of the data, and its training time scales favourably with the number of training points. In contrast, the naive model is too memory intensive to use all data, and its training time scales poorly with the number of training points. See Figure 2.

Using More Data Improves Predictive Performance We measure the predictive performance of SNEPPP as the number of training examples increases in Figure 2. Note that events are not sampled i.i.d. from some distribution and we do not necessarily expect for standard generalisation error bounds to apply, even for fixed feature models such as Marteau-Ferey, Bach, and Rudi (2020, Theorem 5). While we do not necessarily expect the generalisation gap to decay like $\mathcal{O}(N^{-1/2})$ as in the i.i.d. setting, we empirically find that both the test NLL and the error in the number of predicted events improves as more training data is added. We leave theoretical analysis for future work.

Fit to 100 Million Spatio-temporal Events We use a spatio-temporal model with the temporal component as described in § 3.4 and a spatial component consisting of ReLU activations according to (7). We provide a video of our fitted intensity functions varying over time in the supplementary material and Appendix D. Due to the linear scaling in

N of our closed form integral in SNEPPP, this result only required 8 hours of computation time for fitting. We also find that the general SNEPPP model shows increased performance against (mixtures of) log-linear models.

6 Conclusion

We propose an efficient and flexible model for PPPs based on two-layer squared neural networks. In contrast to recent work (Tsuchida, Ong, and Sejdinovic 2023) which considered density estimation, this paper uses the neural network to model the expected value (i.e. intensity) of a PPP. We prove that the resulting integrated intensity function has a closed form. Based on the link between neural networks and kernel methods, this closed form provides a linear time (with respect to input dimension) computation of the integrated intensity function. We also derive two novel NNKs which are well-suited to fitting a large scale wildfire model. We derive a projected gradient descent approach for optimising the NLL, and show that the problem is convex if we hold the first neural network layer fixed. We illustrate the efficacy of our method with empirical experiments on several different performance metrics. Our empirical results show that our exact integration method has similar computation time requirements, but outperform other baseline methods. As a demonstration of the efficiency of our approach, we apply SNEPPP to a case study of 100 million wildfire events.

Limitations and Future Work In order to simplify our notation, we assumed that a single activation function σ is used. One may straightforwardly extend our description to the case where more than one activation function is used in the same layer, as might be the case for certain random Fourier feature inspired feature mappings with lower variance than their single activation function counterparts (Sutherland and Schneider 2015). Our empirical results indicate that increasing the size of the training data improves test likelihood. While the generalisation error of supervised learning is theoretically well studied, we are unaware of any rates of convergence of generalisation error for PPP.

Acknowledgements

Russell and Cheng Soon would like to acknowledge the support of the Machine Learning and Artificial Intelligence Future Science Platform, CSIRO. The authors would like to thank Jia Liu for early discussions about the idea.

References

- Baddeley, A. 2007. *Spatial Point Processes and their Applications*, 1–75. Springer Berlin Heidelberg.
- Baddeley, A.; Chang, Y.-M.; Song, Y.; and Turner, R. 2012. Nonparametric estimation of the dependence of a spatial point process on spatial covariates. *Statistics and its interface*, 5(2): 221–236.
- Boyd, S. P.; and Vandenberghe, L. 2004. *Convex optimization*. Cambridge university press.
- Broomhead, D. S.; and Lowe, D. 1988. Radial basis functions, multi-variable functional interpolation and adaptive networks. Technical report, Royal Signals and Radar Establishment Malvern (United Kingdom).
- Cho, Y.; and Saul, L. K. 2009. Kernel Methods for Deep Learning. In *Advances in Neural Information Processing Systems*, 342–350. PMLR.
- Cressie, N. 1993. *Statistics for Spatial Data*. New York, NY: Wiley.
- Flaxman, S.; Teh, Y. W.; and Sejdinovic, D. 2017. Poisson intensity estimation with reproducing kernels. In *Artificial Intelligence and Statistics*, 270–279. PMLR.
- Han, I.; Zandieh, A.; Lee, J.; Novak, R.; Xiao, L.; and Karbası, A. 2022. Fast neural kernel embeddings for general activations. *Advances in neural information processing systems*, 35: 35657–35671.
- Hendrycks, D.; and Gimpel, K. 2016. Gaussian error linear units. *arXiv preprint arXiv:1606.08415*.
- Kim, H.; Asami, T.; and Toda, H. 2022. Fast Bayesian Estimation of Point Process Intensity as Function of Covariates. In *Advances in Neural Information Processing Systems*.
- Lloyd, C.; Gunter, T.; Osborne, M.; and Roberts, S. 2015. Variational inference for Gaussian process modulated Poisson processes. In *International Conference on Machine Learning*, 1814–1822. PMLR.
- Marteau-Ferey, U.; Bach, F.; and Rudi, A. 2020. Non-parametric models for non-negative functions. *Advances in neural information processing systems*, 33: 12816–12826.
- Menon, A. K.; and Williamson, R. C. 2018. A loss framework for calibrated anomaly detection. In *Proceedings of the 32nd international conference on neural information processing systems*, 1494–1504.
- Mishra, S.; Rizoıu, M.-A.; and Xie, L. 2016. Feature driven and point process approaches for popularity prediction. In *Proceedings of the 25th ACM international on conference on information and knowledge management*, 1069–1078.
- NASA FIRMS. 2023. MODIS Collection 6 Hotspot / Active Fire Detections MCD14ML.
- Nielsen, F.; and Garcia, V. 2009. Statistical exponential families: A digest with flash cards. *arXiv preprint arXiv:0911.4863*.
- Ogata, Y. 1988. Statistical models for earthquake occurrences and residual analysis for point processes. *Journal of the American Statistical association*, 83(401): 9–27.
- Rahimi, A.; and Recht, B. 2007. Random features for large-scale kernel machines. *Advances in neural information processing systems*, 20.
- Renner, I. W.; Elıth, J.; Baddeley, A.; Fithian, W.; Hastie, T.; Phillips, S. J.; Popovic, G.; and Warton, D. I. 2015. Point process models for presence-only analysis. *Methods in Ecology and Evolution*, 6(4): 366–379.
- Rudi, A.; and Ciliberto, C. 2021. PSD Representations for Effective Probability Models. In *Advances in Neural Information Processing Systems*, volume 34, 19411–19422.
- Sellier, J.; and Dellaportas, P. 2023. Sparse Spectral Bayesian Permanental Process with Generalized Kernel. In *International Conference on Artificial Intelligence and Statistics*, 2769–2791. PMLR.
- Shawe-Taylor, J.; and Cristianini, N. 2004. *Kernel methods for pattern analysis*. Cambridge university press.
- Sutherland, D. J.; and Schneider, J. 2015. On the error of random fourier features. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*, 862–871.
- Tsuchida, R.; Ong, C. S.; and Sejdinovic, D. 2023. Squared Neural Families: A New Class of Tractable Density Models. *Advances in neural information processing systems*, 37.
- Wainwright, M. J.; Jordan, M. I.; et al. 2008. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1–2): 1–305.
- Walder, C. J.; and Bishop, A. N. 2017. Fast Bayesian intensity estimation for the permanental process. In *International Conference on Machine Learning*, 3579–3588. PMLR.
- Williams, C.; and Seeger, M. 2000. Using the Nyström method to speed up kernel machines. *Advances in neural information processing systems*, 13.
- Williams, C. K.; and Rasmussen, C. E. 2006. *Gaussian processes for machine learning*. MIT press Cambridge, MA.
- Ziyin, L.; Hartwig, T.; and Ueda, M. 2020. Neural Networks Fail to Learn Periodic Functions and How to Fix It. In *Advances in Neural Information Processing Systems*, volume 33, 1583–1594.