Predictive Models for Point Processes

by

Wenzhao Lian

Department of Statistical Science
Duke University

Date: _____

Approved:

_____
David Dunson, Supervisor

_____
Jerry Reiter

_____
Lawrence Carin

Thesis submitted in partial fulfillment of the requirements for the degree of
Master of Science in the Department of Statistical Science
in the Graduate School of Duke University
2015

Abstract

Predictive Models for Point Processes

by

Wenzhao Lian

Department of Statistical Science
Duke University

Date: _____

Approved:

_____
David Dunson, Supervisor

_____
Jerry Reiter

_____
Lawrence Carin

An abstract of a dissertation submitted in partial fulfillment of the requirements for the degree of Master of Science in the Department of Statistical Science in the Graduate School of Duke University
2015

Abstract

Point process data are commonly observed in fields like healthcare and social science. Designing predictive models for such event streams is an under-explored problem, due to often scarce training data. In this thesis, a multitask point process model via a hierarchical Gaussian Process (GP) is proposed, to leverage statistical strength across multiple point processes. Nonparametric learning functions implemented by a GP, which map from past events to future rates, allow analysis of flexible arrival patterns. To facilitate efficient inference, a sparse construction for this hierarchical model is proposed, and a variational Bayes method is derived for learning and inference. Experimental results are shown on both synthetic data and as well as real electronic health-records data.

Contents

List of Tables

List of Figures

Acknowledgements

I would like to thank my advisor David Dunson for his insightful suggestions. David's broad knowledge in Bayesian statistics, deep understanding in modeling, ability to spot key issues in problems, are really helpful in my research and study. I would also like to thank my committee member Jerry Reiter for his helpful suggestions in the thesis writing. Jerry is the first one who introduced me Bayesian statistics, his passion, knowledge, and the way of communnication make me interested in this area and form a solid foundation for research. I would also like to thank my PhD advisor Larry Carin who provides a lot of insightful opinions in my research. Larry also provides me financial support and all kinds of resources including the access to datasets and the opportunities to collaborate with others.

I would like to thank my coauthor Vinayak Rao at Purdue for helpful discussions. Vinayak's critical thinking, solid knowledge in statistics, and insightful ideas affect me a lot in research. I would also like to thank Ricardo Henao and Joe Lucas for their helpful suggestions in this work. They have expertise in both statistics and healthcare, and help me to formulate a meaningful problem to solve in both fields.

I would also like to thank Chris Meek and Asela Gunawardana at Microsoft research. The discussions with Chris and Asela during my internship there were very helpful for me to explore the field of point processes.

1

Introduction

Point process data have seen increased attention in fields like neuroscience (Rad and Paninski, 2011), electronic commerce (Xu et al., 2014), and healthcare analysis (Lasko, 2014). One thread of work has focused on learning arrival rates by imposing smoothness on a latent rate function (Adams et al., 2009; Rao and Teh, 2011; Lloyd et al., 2014). Another interesting problem consists of predicting future arrivals as a direct function of past observations (Pillow et al., 2008; Gunawardana et al., 2011). Taking healthcare analysis as a motivating example, we focus on the latter: given a patient's hospital visit history up to time t, (i) when will the next visit happen? and (ii) how many visits will the patient have in $[t, t+L]$? Answering such questions provides a quantitative evaluation of the patient's risk, which helps in making treatment plans and allocating hospital resources more efficiently (Amarasingham et al., 2010). Similar problems also arise in other fields, such as predicting purchasing behavior for individual customers, or predicting failures in distributed computer systems.

A few works have explored the prediction problem in point processes by learning a functional mapping from history features to the current intensity rate (Pillow et al., 2008; Rajaram et al., 2005; Gunawardana et al., 2011). In Gunawardana et al. (2011),

the intensity function is constrained to be piecewise-constant, learned using decision trees and used for prediction. In our proposed multitask point process model, we build upon this piecewise-constant intensity model. To allow for flexibility of the function mapping from history features to future rate, and to capture the uncertainty of estimation, we use a nonparametric method, by imposing a Gaussian process (GP) prior on the intensity rate. However, when building such predictive models for event arrival processes, one difficulty is that the available training data are scarce for each subject/task. Therefore, we treat each individual arrival process as a task and follow a *multitask learning* approach to share information from all tasks in a *hierarchical* manner.

Methods for learning GPs from multiple tasks have been proposed (Yu et al., 2005), but involve a shared global mean function, inferred at all observed inputs (history features) across all the tasks. The posterior of this function cannot be directly updated due to non-conjugacy of the point process likelihood to GP priors. One approximation method, variational Bayes, is often applied, leading to the number of unknown parameters scaling as $\mathcal{O}(N^2)$, where N is the number of unique features from all tasks. Borrowing from the framework of *pseudo inputs* in the GP literature (Snelson and Ghahramani, 2006; Titsias, 2009), we constrain the rate functions to an M-dimensional latent space, where $M \ll N$, reducing the parameter space to $\mathcal{O}(M^2 + MP)$, where P is the GP input dimension. By adjusting the locations of the pseudo inputs, we can effectively share data across tasks and efficiently represent these functions.

Learning the history-to-rate mapping functions allows one to predict future events by analytical integration or forward sampling. We consider both, evaluating our model and inference methodology on both synthetic and a real Electronic Health Records (EHR) dataset. The latter involves different categories of health problems, and we demonstrate that future hospital visits for some types of diseases are pre-

dictable even with simple history features. Our work has two main contributions: (*i*) providing an efficient approach to share data/parameters in hierarchical/multitask GP models; and (*ii*) building a predictive model for arrival data from multiple event streams, using point processes in a multitask scheme.

The rest of the thesis is outlined as follows. In Chapter 2, we review a few existing works on point process intensity models and predictive models. In Chapter 3, we highlight the piecewise-constant conditional intensity model and detail the building blocks of our proposed methods in both Chapter 4 and 5. A single task predictive model is proposed and evaluated on synthetic data in Chapter 4, which provides insights into how to design predictive models leveraging the conditional intensity model assumption. In Chapter 5, a multitask point process model is proposed, which is the key contribution of this thesis. Model properties, inference method, and experimental results on both synthetic and real-world healthcare data are presented in details. Last in Chapter 6, we summarize the thesis and list a few interesting future directions. The work on multitask point process predictive modeling has been published by the author in Lian et al. (2015).

2

Background

GP-modulated point processes are a popular approach for modeling event streams (Adams et al., 2009; Rao and Teh, 2011; Lloyd et al., 2014; Lasko, 2014). Assuming a smoothly varying intensity function, the intensity rate, as well as its uncertainty, can be estimated from observed streams. Extrapolation can be used for short-term prediction. There are two main limitations of these models. First, the smoothness assumption does not hold in many scenarios. Sudden rate changes often happen upon event arrivals, *e.g.*, the risk of a patient's hospital visit might change significantly after a single visit. Second, these models involve a common modulating GP, so that multiple streams have to be aligned, something not always appropriate or possible. For example, similar arrival patterns might appear at different periods of time for different streams, which cannot be captured by such models.

Another relevant line of work on point process predictive modeling comes from the neural decoding literature (Kulkarni and Paninski, 2007; Rad and Paninski, 2011; Pillow et al., 2008), where multiple subjects/neurons are affected by common stimuli (features), generating event streams. A generalized linear model can be trained on the stimuli or spiking history to learn the rate function, and further predict future

spiking events. Meanwhile, the network structure across neurons can be inferred, which in turn, helps with prediction. However, temporal alignment is also assumed in these models, which is valid in the setting of neural decoding, but not in the cases we consider.

Also related is the work of Weiss and Page (2013), which integrates the multiplicative forest Continuous Time Bayesian Network (mfCTBN) with the piecewise-constant intensity model of Gunawardana et al. (2011). More precisely, they learn forests mapping from demographic and event history features to intensity rates, but under the assumption that all subjects have the same function, *i.e.*, each event stream is a realization trajectory of some underlying model. In our work, we consider a different setting, where the variability across subjects cannot be ignored. This difference is crucial in scenarios such as healthcare analysis and electronic commerce, where variability among members of a population affects arrival pattern discovery and predictive performance.

3

Piecewise-constant Conditional Intensity Model

Multi-task point process observations can be summarized as sequences of arrival time stamps $\{y_n^u\}$, with $n = 1, \cdots, D^u$ and $u = 1, \cdots, U$, where y_n^u represents the time stamp of the n-th arrival of subject/task u. Our goal is to model the sequences and make predictions for future event arrivals.

The event streams can be naturally modeled using an intensity model, with a hazard rate function $\gamma(t)$. In an infinitesimal time interval Δ, the probability of an event occuring in this interval is given by $\Delta\gamma(t)$. To build the temporal dependency between past observations and future events, we choose the hazard rate as a function of past observations $\boldsymbol{h}(t)$, denoted as $\gamma(t|\boldsymbol{h}(t))$, in which $\boldsymbol{h}(t)$ summarizes the history of past observations (Gunawardana et al., 2011); details on the potential form of $\boldsymbol{h}(t)$ are discussed subsequently. Denoting $\mathcal{Y}^u = \{y_1^u, \cdots, y_{D^u}^u\}$ as the set of arrival time stamps in task u, the likelihood is

$$p(\mathcal{Y}^u) = \prod_{n=1}^{D^u} \gamma^u(y_n^u|\boldsymbol{h}^u(y_n^u)) \exp(-\int_{y_n^u}^{y_{n+1}^u} \gamma^u(\tau|\boldsymbol{h}^u(\tau))d\tau)\,.$$

Assuming $\gamma^u(t|\boldsymbol{h}^u(t))$ as piecewise constant with N^u change points (pieces) at $\{t_i^u\}_{i=1}^{N^u}$,

and piece length $\Delta_i^u = t_{i+1}^u - t_i^u$, we have

$$p(\mathcal{Y}^u) = \prod_{i=1}^{N^u} \gamma^u(t_i^u|\boldsymbol{h}^u(t_i^u))^{\mathbb{I}(t_i^u \in \mathcal{Y}^u)} \times \exp\{-\Delta_i^u \gamma^u(t_i^u|\boldsymbol{h}^u(t_i^u))\}\,. \tag{3.1}$$

Many approaches exist to extract features $\boldsymbol{h}(t)$ from past event arrivals. One possible feature-construction approach uses empirical rates at recent time points (Rajaram et al., 2005), *i.e.*, $\boldsymbol{h}^u(t_i^u) \in \mathbb{R}_+^P$ with $h_p^u(t_i^u) = \bar{\gamma}(t_i^u - L_p)$ for $L_1, L_2, \cdots, L_P$ predefined lengths of memory. Here $\bar{\gamma}(t_i^u - L_p)$ refers to the empirically computed rate at time $t_i^u - L_p$. Keeping the algorithm simple, we adopt a construction similar in spirit to Gunawardana et al. (2011): $\boldsymbol{h}^u(t_i^u)$ is P dimensional, where its p-th element denotes the count of arrivals in $[t_i^u - L_p, t_i^u]$. Therefore, $\boldsymbol{h}^u(t)$ is a piecewise constant function. Revisiting the likelihood in (3.1), we observe that N^u is the number of change points (pieces) of the feature function $\boldsymbol{h}^u(t)$ going through the whole event stream $\{y_n^u\}_{n=1}^{D^u}$, and that N^u scales linearly with D^u (the number of events). Side information in the form of covariates can also be considered as a natural extension, by augmenting the feature space. Constructing richer features from history observations is left as an interesting direction for future work.

To complete the intensity model, the functional mapping from feature space to intensity rate must be specified. Define the space of the history features $\boldsymbol{h}^u(t_i^u)$ as $\mathcal{H}$, where each point $\boldsymbol{h}^u(t_i^u) \triangleq \boldsymbol{h}_i^u \in \mathcal{H}$ is a possible feature vector. Here we impose a GP prior on a set of functions $f^u(\cdot) : \mathcal{H} \to \mathbb{R}$, followed by a transformation to ensure the non-negativity of intensity rates.

Multiple choices exist for this transformation, including the square transformation ($\gamma^u(\cdot) = \{f^u(\cdot)\}^2$) and the exponential transformation ($\gamma^u(\cdot) = \exp\{f^u(\cdot)\}$). Other transformations such as the normal cumulative density function and logistic functions also worth mentioning but they require a pre-defined maximum intensity

value, which is not plausible. Therefore in this thesis, we focus on exploring the square and exponential transformations. We remark that no "right" choice exists for a given model; instead, the transformation is chosen to fascilitate the inference. In Chapter 4, the commonly used exponential transformation is adopted because of the simplicity of log-likelihood evaluation during the Markov Chain Monte Carlo (MCMC) inference. In Chapter 5, we use a square transformation rather than the exponential transformation ($\gamma^u(\cdot) = \exp\{f^u(\cdot)\}$), because using the variational inference algorithm proposed for that model, the uncertainty of $f^u(\cdot)$ cannot be properly estimated in the latter construction, as discussed by Lloyd et al. (2014); this issue is addressed thoroughly in Section 5.2. To resolve the ambiguity in the square transformation caused by $\{f^u(\cdot)\}^2 = \{-f^u(\cdot)\}^2$, a prior specification favoring the positive half-space is imposed and works well in practice.

4

A Single Task Point Process Predictive Model

4.1 Model

In this chapter, we propose a predictive model for single task point process observations, *i.e.*, the number of event streams $U = 1$. For brevity, we drop the index u in this chapter. Inspired by the GP modulated renewal process (Lasko, 2014), we use a kernel to impose contraints on the rate function $\gamma(t)$, more precisely, on the transformed rate function $f(t)$ at different time points. Use $\boldsymbol{K}$ as an infinite dimensional square matrix, defined as follows,

$$\begin{aligned} \boldsymbol{k}(\boldsymbol{h}(t_1), \boldsymbol{h}(t_2)) &= \tau \exp\{-(\boldsymbol{h}(t_1) - \boldsymbol{h}(t_2))^T \boldsymbol{D} (\boldsymbol{h}(t_1) - \boldsymbol{h}(t_2))\} \\ &\quad (1 + \tau_t \exp\{-\frac{(t_1 - t_2)^2}{l}\}), \qquad (4.1) \\ \text{where } \boldsymbol{D} &= \text{Diag}(\lambda_1^{-1}, \cdots, \lambda_P^{-1}). \end{aligned}$$

Inspecting (4.1), $\boldsymbol{k}(\boldsymbol{s}(t_1), \boldsymbol{s}(t_2))$, a "spatial-temporal" kernel, encourages instances with similar history features ("spatial locations") and time to have similar intensity values. τ and λ are the hyper-parameters for the spatial kernel, controlling the amplitude and correlation length in the history feature space. The spatial kernel

discovers similar history features across the whole event stream, while the temporal kernel imposes smoothness in nearby time points on the rate function via varying the correlation length in time l. τ_t controls the strength of the smoothness influencing the rate function.

Since our goal is prediction on future events, we need to note that the predictive power is mainly endowed by the spatial kernel, while the temporal kernel smoothes the prediction, preventing abrupt changes in the intensity function. In the model proposed in Chapter 5, we only keep the spatial kernel while discard the temporal kernel component. In that case, how the intensity function varies is governed by the changes in the feature function $\boldsymbol{h}(t)$.

The generative process is described as follows in a discrete formulation (discretization with right continuity) as $t_i - t_{i-1} = \Delta$, $\Delta \to 0$, and it naturally generalizes to the continuous scenario. Because of being a conditional intensity model, the generative model proceeds in a forward manner: at time t_i, with $\boldsymbol{h}(t_j)$ and $\gamma(t_j), t_j < t_i$ known, $\gamma(t_i)$ and the events falling in $[t_{i-1}, t_i)$ are generated as follows,

$$
\begin{aligned}
N(t_{i+1}) - N(t_i) &\sim Pois(\gamma(t_i)\Delta), && (4.2)\\
\log \gamma(t_i) &\sim \mathcal{N}(\mu_i, \sigma_i^2), && (4.3)\\
\mu_i &= \boldsymbol{k}(\cdot, \boldsymbol{h}(t_i))^T \boldsymbol{K}_{i-1}^{-1}(\boldsymbol{f}_{i-1} - \boldsymbol{f}_0), && (4.4)\\
\sigma_i^2 &= \boldsymbol{k}(\boldsymbol{h}(t_i), \boldsymbol{h}(t_i)) - \boldsymbol{k}(\cdot, \boldsymbol{h}(t_i))^T \boldsymbol{K}_{i-1}^{-1} \boldsymbol{k}(\cdot, \boldsymbol{h}(t_i)). && (4.5)
\end{aligned}
$$

In above, $\boldsymbol{k}(\cdot, \boldsymbol{h}(t_i)) = [\boldsymbol{k}(\boldsymbol{h}(t_1), \boldsymbol{h}(t_i)), \cdots, \boldsymbol{k}(\boldsymbol{h}(t_{i-1}), \boldsymbol{h}(t_i))]^T$, $\boldsymbol{K}_{i-1} \in \mathbb{R}_+^{(i-1)\times(i-1)}$, with $\{j,k\}$-th element as $\boldsymbol{k}(\boldsymbol{h}(t_j), \boldsymbol{h}(t_k))$, $\boldsymbol{f}_{i-1} = [\log \gamma(t_1), \cdots, \log \gamma(t_{i-1})]^T$, and $\boldsymbol{f}_0$ is the mean log intensity (fixed at a prior value like 0 or based on the empirical mean intensity). As can be seen, the forward generative process for transformed rates $f(t_i)$ are consecutive conditional Gaussian distributins. Therefore, the marginal prior distribution of the rate vector $[f(t_1), \cdots, f(t_N)]^T$ is also Gaussian, which makes the inference feasible by borrowing techniques for GP latent variable models.

4.2 Experiments

In this section, we evaluate the proposed model on synthetic data. Given an event stream, the variables to infer include kernel parameters $\{\tau, \{\lambda_p\}_{p=1}^P, \tau_t, l\}$ and the rate function $\gamma(t)$. Slice sampling is used to sample the kernel parameters and $\lambda(t)$ can be updated with an auxilary sampler (Murray and Adams, 2010; Lasko, 2014).

The event stream is generated according to an intensity function $\tilde{\gamma}(t) = \boldsymbol{\beta}^T \boldsymbol{h}(t)$, where $h_p(t)$ is the number of events in the window $[t - t_p, t]$. $p = 1, 2$ and $t_1 = 1, t_2 = 5$. Note that this is not exactly the same generative process as defined by the model proposed in Section 4.1; instead, it is a piecewise constant conditional intensity point process where the conditional intensity is constructed via linear dependence on event history features. To bound the intensity function (otherwise, too frequenty event arrivals cause computational issues and zero event stream trivially leads to a constant-zero intensity function), we constrain the intensity $\tilde{\lambda}(t)$ by upper and lower bounds $\tilde{\lambda}_{max}$ and $\tilde{\lambda}_{min}$. In the experiments, $\tilde{\lambda}_{max} = 1.0$ and $\tilde{\lambda}_{min} = 0.2$. Following $\tilde{\lambda}(t)$ defined above, we can, in a forward manner, generate the event streams over $[0, T]$. Synthetic event streams are generated 10 times to obtain the mean and standard deviation of the quantitative metrics of the model performance. For each observed stream, the segment of $[0, T/2]$ is used as training and the other segment (with the same length) over $[T/2, T]$ is used as testing. In the experiments, $T = 200$.

Three methods are compared in this set of experiments: (i) the proposed method with $\boldsymbol{h}(t)$ constructed from empirical rates, *i.e.*, $h_p(t) = \bar{\gamma}(t - L_p)$; the empirical rate is estimated using the kernel estimation method for point process data in Diggle (1985), shorted as kernel-ST; (ii) the proposed method using the raw counts in previous windows as $\boldsymbol{h}(t)$, shorted as count-ST; and (iii) the GP modulated renewal process in Lasko (2014), shorted as GP-renewal.

We first evaluate the intensity function inferred on the training segment. To

compare with the exising GP modulated renewal process model, we compute the log-likelihood of the training segments using intensity functions learned. As demonstrated in Table 4.1, all three methods achieve similar data log-likelihood on the training segments. The difference is insignificant considering the variation of the performance. As the GP modulated renewal process model is a state-of-the-art intensity estimation method for point processes, it is safe to conclude that the proposed model is competitive/useful in intensity estimation (when data satisfy the conditional intensity model assumptions) even they are designed for forward prediction tasks. To keep the algorithms simple, in Chapter 5, the raw counts in previous windows are used as history features. We leave history feature construction as an open problem for future work since it is found crucial in many real-world applications.

Table 4.1: Comparison of data log-likelihood obtained using the proposed method with spatial-temporal kernel using two different history features: empirical rate using kernel estimation methods and raw counts, and the GP modulated renewal process.

	log-likelihood, training	log-likelihood, testing
kernel-ST	-73.284±13.416	-74.021±21.175
count-ST	-68.735±11.167	-71.856± 21.337
GP-renewal	-72.501±11.119	-210.068±354.220

We then evaluate the predictive power of the methods proposed, where the data log-likelihood of the testing segments is computed. Specifically, using the models trained, we extrapolate the intensity functions over the testing segments, and evaluating the likelihood on those segments. As can be seen, the proposed methods, including kernel-ST and count-ST, achieve similar data log-likelihood on the testing set as on the training set. However, using the GP modulated renewal process, the model does not have predictive power; especially for testing time points farther away from the observation intervals, the predictive intensity degrades to the prior mean since the correlation between testing points $t \in [T/2, T]$ and training points $t \in [0, T/2]$ decreases to 0 in the GP covariance matrix. In contrast, in the proposed

methods, for a testing point t, the model finds among training points t' with features $\boldsymbol{h}(t')$ similar with $\boldsymbol{h}(t)$, and outputs a weighted sum of the corresponding intensity rates $\gamma(t')$ as the predicted rate $\gamma(t)$.

To summarize, the proposed method achieves reasonable goodness-of-fit and predictive power with sufficient training data. However, in practice, point process data is often scarce from a single event stream, for instance, the purchasing behaviour of a customer is sparse and the hospital visit history of a patient is short. In the next chapter, we will discuss an approach to utilize multiple correlated event streams for predictive model learning.

5

A Multitask Point Process Predictive Model

5.1 Model

As mentioned in Chapter 1 and 4, sharing information across tasks is necessary when learning the rate functions from sparse data. Accordingly, in this chapter, we consider a hierarchical GP construction (not built upon the model designed in Chapter 4).

Denote the rate at feature vector $\boldsymbol{h}_i^u$ (time t_i and event stream u) as γ_i^u, with the corresponding transformed rate denoted as $f_{N,i}^u$. The generative process from history features to rates may be described as follows

$$\boldsymbol{\mu}_N^0 \sim \mathcal{GP}\left(\boldsymbol{g}, \frac{1}{\xi}\boldsymbol{K}_{NN}\right), \tag{5.1}$$

$$\boldsymbol{f}_N^u \sim \mathcal{N}\left(\boldsymbol{\mu}_N^0, \boldsymbol{K}_{NN}\right), \tag{5.2}$$

$$\gamma_i^u = \{f_{N,i}^u\}^2. \tag{5.3}$$

In (5.1), $\boldsymbol{g}$ is a prior mean, which may be set according to prior knowledge or the empirical average intensity rate over all tasks. Parameter ξ controls the complexity of the hyper-prior. The covariance matrix $\boldsymbol{K}_{NN}$ can be obtained from a squared-exponential or Matérn function with automatic relevance determination (ARD) co-

variance kernel, for example,

$$\boldsymbol{K}_{NN,ij} = \tau^2 \exp\left(-\sum_{p=1}^{P} \frac{(h_{pi} - h_{pj})^2}{2\lambda_p^2}\right) + \sigma^2 \mathbb{I}(i = j)\,, \tag{5.4}$$

where $\boldsymbol{h}_i \in \mathcal{H}$ may be any possible history feature. In (5.1), $\boldsymbol{\mu}_N^0$ determines the global mean function of transformed rate functions, $\boldsymbol{f}_N^u$, from all U tasks. For each individual task, the task-specific transformed rate function, $\boldsymbol{f}_N^u$, is a "noisy version" of the global mean function $\boldsymbol{\mu}_N^0$. Our construction is inspired by Yu et al. (2005), which can be shown analogous to a hierarchical construction of the multitask linear regression model

$$\boldsymbol{w}^0 \sim \mathcal{N}\left(\boldsymbol{v}, \frac{1}{\xi}\boldsymbol{I}\right)\,, \tag{5.5}$$

$$\boldsymbol{w}^u \sim \mathcal{N}\left(\boldsymbol{w}^0, \boldsymbol{I}\right)\,, \tag{5.6}$$

$$f_{N,i}^u \sim \mathcal{N}\left(\boldsymbol{w}^{u\top}\boldsymbol{h}_i, \sigma^2\right)\,, \tag{5.7}$$

where $\boldsymbol{w}^0$ and $\boldsymbol{w}^u$ are the global and task-specific regression parameters, respectively, and $\boldsymbol{v}$ is the hyper-mean for $\boldsymbol{w}^0$. Choosing a linear kernel for (5.1) and (5.2) ($\boldsymbol{K}_{NN,ij} = \boldsymbol{h}_i^T\boldsymbol{h}_j$), and letting $\boldsymbol{v} = \boldsymbol{g} = 0$, the construction through (5.1)-(5.2) is equivalent to the multitask linear regression model in (5.5)-(5.7) on a finite observed dataset.

If there exists enough training data, *i.e.*, $\{\boldsymbol{h}_i^u\}_{i=1}^{N^u}$ densely covers the feature space $\mathcal{H}$, $\boldsymbol{f}_N^u$ is mainly determined by the data in task u. Otherwise, $\boldsymbol{f}_N^u$ is significantly affected by the prior $\boldsymbol{\mu}_N^0$, which pools together information from rate functions $\{\boldsymbol{f}_N^u\}$ from all tasks.

However, a problem with the hierarchal GP construction using (5.1) and (5.2) is that the dimension of $\boldsymbol{\mu}_N^0$ and $\boldsymbol{f}_N^u$ is typically massive, in fact it is not smaller than the number of unique history features from all tasks, *i.e.*, $|\cup_u \{\boldsymbol{h}_i^u\}_{i=1}^{N^u}|$, where $|\mathcal{S}|$ refers to the cardinality of a set $\mathcal{S}$. Thus, the number of unknown parameters in the

model scales with $\mathcal{O}(U\sum_{u=1}^{U}N^u)$, or $\mathcal{O}(U(\sum_{u=1}^{U}N^u)^2)$ if variance estimation is also considered. In point process models in particular, this is impractical because the influence of the likelihood is weak and noise levels are high. Therefore, we propose an alternative approach inspired by the pseudo input framework for GPs (Snelson and Ghahramani, 2006; Titsias, 2009). Specifically, we propose a two-step process by introducing an M-dimension vector $\boldsymbol{f}_M^u$ for each task, where $M \ll |\cup_u \{\boldsymbol{h}_i^u\}_{i=1}^{N^u}|$. Then, using a so-called *conditional Gaussian Processes*, we define the generative process as

$$\begin{aligned}
\boldsymbol{\mu}_M^0 &\sim \mathcal{N}\left(\boldsymbol{g}, \frac{1}{\xi}\boldsymbol{K}_{MM}\right), && (5.8)\\
\boldsymbol{f}_M^u &\sim \mathcal{N}\left(\boldsymbol{\mu}_M^0, \boldsymbol{K}_{MM}\right), && (5.9)\\
\boldsymbol{f}_N^u|\boldsymbol{f}_M^u &\sim \mathcal{GP}\Big(\boldsymbol{g} + \boldsymbol{K}_{NM}^u\boldsymbol{K}_{MM}^{-1}(\boldsymbol{f}_M^u - \boldsymbol{g}), && (5.10)\\
&\qquad (1+\frac{1}{\xi})(\boldsymbol{K}_{NN}^u - \boldsymbol{K}_{NM}^u\boldsymbol{K}_{MM}^{-1}{\boldsymbol{K}_{NM}^u}^\top)\Big),
\end{aligned}$$

where $\boldsymbol{f}_N^u$ contains the transformed rates at all possible feature vectors $\{\boldsymbol{h}_i^u\}_{i=1}^{N^u}$ in task u, while $\boldsymbol{f}_M^u$ is a M-dimensional vector, consisting of the transformed rates at M feature-vector locations $\{\boldsymbol{s}_i\}_{i=1}^{M}$.

In (5.10), for each task u, only the function values at feature vectors appearing in task u are required, while in (5.2), for each task, the function values at all feature vectors appearing in all tasks need to be specified.

Following the GP literature Titsias (2009), we refer to $\{\boldsymbol{s}_i\}$ as pseudo inputs. These need not appear in any of the tasks (or even $\mathcal{H}$); we only require that a distance can be properly defined between $\boldsymbol{s}_i$ and $\boldsymbol{h}_j^u$. And similar to (5.4), $\boldsymbol{K}_{MM}$ and $\boldsymbol{K}_{NM}^u$ can be obtained through the squared-exponential or the Matérn function with

ARD kernel, with $\boldsymbol{h}_j^u$ and $\boldsymbol{s}_i$ as covariates. Specifically,

$$\boldsymbol{K}_{MM,ij} = \tau^2 \exp\left(-\sum_{p=1}^{P} \frac{(s_{pi} - s_{pj})^2}{2\lambda_p^2}\right) + \sigma^2 \mathbb{I}(i=j)\,,$$

$$\boldsymbol{K}_{NM,ij}^u = \tau^2 \exp\left(-\sum_{p=1}^{P} \frac{(h_{pi}^u - s_{pj})^2}{2\lambda_p^2}\right).$$

It can be shown that the construction through (5.8)-(5.10) results in the same marginal prior distribution for $\boldsymbol{f}_N^u$ as (5.1)-(5.2) (see Appendix A) while having a significantly reduced computational cost. Overfitting problems can also be alleviated because model complexity is reduced. This is especially beneficial when sophisticated features are constructed, resulting in very large history feature spaces, however, with the rate function likely living in a low-dimensional manifold. As a result, assuming $\{\boldsymbol{s}_i, f_{M,i}^u\}_{i=1}^M$ captures the characteristics of the function $\{\boldsymbol{h}_i^u, f_{N,i}^u\}_{i=1}^{N^u}$, while achieving computational savings, and improved estimation accuracy. The proposed full generative model is summarized by equations (5.8)-(5.10), (5.3) and (3.1).

5.2 Inference

Model parameters $\boldsymbol{\Theta}$ include the pseudo-input locations, the global mean of transformed rates, and the GP hyperparamters: $\boldsymbol{\Theta} = \{\{\boldsymbol{s}_m\}_{m=1}^M, \boldsymbol{\mu}_M^0, \sigma^2, \tau^2, \{\lambda_p\}_{p=1}^P\}$. Obtaining a full posterior distribution for $\boldsymbol{\mu}_M^0$ is straightforward due to local conjugacy. However we observed no significant difference in performance against just a point estimate, and for simplicity we focus here on the latter.

Maintaining $p(\boldsymbol{f}_M^u, \boldsymbol{f}_N^u | \mathcal{Y}, \boldsymbol{\Theta})$, the full posterior distribution over the intensity functions of each task is important for transfer learning across tasks with different numbers of observations, where $\mathcal{Y}$ refer to the complete collection of event streams $\mathcal{Y}^1, \cdots, \mathcal{Y}^U$. However, this is not straightforward due to the non-conjugate point process likelihood, and borrowing ideas from variational learning for sparse GPs (Lloyd et al., 2014; Titsias, 2009), we propose a variational form to approximate the posterior $p(\boldsymbol{f}_M^u, \boldsymbol{f}_N^u | \mathcal{Y}, \boldsymbol{\Theta})$. Specifically, we use

$$\begin{aligned} q(\boldsymbol{f}_N^u, \boldsymbol{f}_M^u) &= p(\boldsymbol{f}_N^u | \boldsymbol{f}_M^u) q(\boldsymbol{f}_M^u) \,. \\ &= p(\boldsymbol{f}_N^u | \boldsymbol{f}_M^u) \mathcal{N}(\boldsymbol{f}_M^u; \boldsymbol{\mu}^u, \boldsymbol{\Sigma}^u) \,. \end{aligned} \tag{5.11}$$

Note that we allow a free-form Gaussian distribution for $\boldsymbol{f}_M^u$, the task-specific transformed rates at the pseudo inputs. However the transformed rates evaluated at all features in each task, $\boldsymbol{f}_N^u$, are constrained by the low-dimensional function $\boldsymbol{f}_M^u$ via (5.10). In the following, we use a simplified notation q^u to denote $q(\boldsymbol{f}_M^u) = \mathcal{N}(\boldsymbol{f}_M^u; \boldsymbol{\mu}^u, \boldsymbol{\Sigma}^u)$. Because of the special factorized form in (5.11), $\{\boldsymbol{\mu}^u, \boldsymbol{\Sigma}^u\}$ are the only variational parameters in the algorithm.

The inference objective is to maximize the variational lower bound (Beal, 2003)

(called ELBO for evidence lower bound optimization):

$$\log p(\mathcal{Y}, \mathbf{\Theta}) \geqslant \mathbb{E}_{\prod_{u=1}^{U} q(\boldsymbol{f}_N^u, \boldsymbol{f}_M^u)}[\log p(\mathcal{Y}, \mathbf{\Theta}) - \log(\prod_{u=1}^{U} q(\boldsymbol{f}_N^u, \boldsymbol{f}_M^u))]$$

$$= \log p(\mathbf{\Theta}) + \sum_{u=1}^{U} \left\{ \mathbb{E}_{q(\boldsymbol{f}_N^u, \boldsymbol{f}_M^u)}[\log \frac{p(\mathcal{Y}^u, \boldsymbol{f}_N^u, \boldsymbol{f}_M^u)}{q(\boldsymbol{f}_N^u, \boldsymbol{f}_M^u)}] \right\} . \tag{5.12}$$

Note that in (5.12), $q(\boldsymbol{f}_N^u, \boldsymbol{f}_M^u)$ and $p(\mathcal{Y}^u, \boldsymbol{f}_N^u, \boldsymbol{f}_M^u)$ implicitly depend on $\mathbf{\Theta}$. Since we only impose a Gaussian prior on $\boldsymbol{\mu}_M^0$ in (5.8), $\log p(\mathbf{\Theta})$ is simplified to $\log p(\boldsymbol{\mu}_M^0)$. Priors on other model parameters may also be imposed, *e.g.*, truncated prior on GP hyper-parameters, but here we learn their Maximum Likelihood Estimator (MLE) instead. To show explicit dependence, we denote the ELBO as $\mathcal{F}(q^1, \cdots, q^U, \mathbf{\Theta})$. We maximize $\mathcal{F}(q^1, \cdots, q^U, \mathbf{\Theta})$ in (5.12), giving a variational Expectation Maximization (EM) algorithm guaranteed to converge to a local optimum (Beal, 2003). In practice, we alternate between a variational E-step, where $\mathbf{\Theta}$ is fixed and $\mathcal{F}(q^1, \cdots, q^U, \mathbf{\Theta})$ is maximized *w.r.t.* $\{q^u\}_{u=1}^U$, and a variational M-step, where $\{q^u\}_{u=1}^U$ is fixed and $\mathcal{F}(q^1, \cdots, q^U, \mathbf{\Theta})$ is maximized *w.r.t.* $\mathbf{\Theta}$. Derivation details are standard; we list the key steps below:

$$\mathcal{F}(q^1, \cdots, q^U, \mathbf{\Theta}) = \sum_{u=1}^{U} \mathbb{E}_{q^u} \left[\mathbb{E}_{p(\boldsymbol{f}_N^u | \boldsymbol{f}_M^u)}[\log p(\mathcal{Y}^u | \boldsymbol{f}_N^u)] \right]$$

$$+ \sum_{u=1}^{U} \mathbb{E}_{q^u} \left[\log \frac{p(\boldsymbol{f}_M^u | \boldsymbol{\mu}_M^0)}{q(\boldsymbol{f}_M^u | \boldsymbol{\mu}^u, \mathbf{\Sigma}^u)} \right] + \log p(\boldsymbol{\mu}_M^0)$$

$$\triangleq \mathcal{F}_1 + \mathcal{F}_2 + \mathcal{F}_3 . \tag{5.13}$$

The first term $\mathcal{F}_1$ (5.14) measures how the functions specified by q^u (determining the distribution of function values at pseudo inputs) fit the observations. Because of the sparsity assumption of the conditional Gaussian process, we can integrate out $\boldsymbol{f}_N^u$ and leave $\mathcal{F}_1$ as a function of only $\{\boldsymbol{\mu}^u, \mathbf{\Sigma}^u\}$ and $\mathbf{\Theta}$:

$$\mathcal{F}_1 = \sum_{u=1}^{U} \int q(\boldsymbol{f}_N^u|\boldsymbol{\mu}^u, \boldsymbol{\Sigma}^u) \log p(\mathcal{Y}^u|\boldsymbol{f}_N^u) d\boldsymbol{f}_N^u \tag{5.14}$$
$$= \sum_{u=1}^{U} \sum_{i=1}^{N^u} \left\{ \mathbb{I}(t_i^u \in \mathcal{Y}^u) \mathbb{E}[\log(f_{N,i}^u)^2] - \Delta_i^u \mathbb{E}[(f_{N,i}^u)^2] \right\} ,$$

The expectation is *w.r.t.* $q(\boldsymbol{f}_N^u|\boldsymbol{\mu}^u, \boldsymbol{\Sigma}^u) = \mathcal{N}(\boldsymbol{b}^u, \boldsymbol{B}^u)$, with parameters in (5.15) and (5.16):

$$\boldsymbol{b}^u = \boldsymbol{g} + \boldsymbol{K}_{NM}^u \boldsymbol{K}_{MM}^{-1} (\boldsymbol{\mu}^u - \boldsymbol{g}) , \tag{5.15}$$

$$\boldsymbol{B}^u = \left(1 + \frac{1}{\xi}\right) \left(\boldsymbol{K}_{NN}^u - \boldsymbol{K}_{NM}^u \boldsymbol{K}_{MM}^{-1} \boldsymbol{K}_{NM}^{u\top}\right)$$
$$+ \boldsymbol{K}_{NM}^u \boldsymbol{K}_{MM}^{-1} \boldsymbol{\Sigma}^u (\boldsymbol{K}_{NM}^u \boldsymbol{K}_{MM}^{-1})^\top . \tag{5.16}$$

$\mathbb{E}[\log(f_{N,i}^u)^2]$ in (5.14) can be calculated using confluent hypergeometric functions (Lloyd et al., 2014). Further details can be found in Appendix B, where we also provide a robust approximation to tackle the the well known numerical instability issue in confluent hypergeometric function evaluations (Ancarani and Gasaneo, 2008).

As a side note on why we prefer a squared transformation over the exponential, when using the latter, $\mathcal{F}_1$ is modified as (5.17), denoted as $\tilde{\mathcal{F}}_1$:

$$\tilde{\mathcal{F}}_1 = \sum_{u=1}^{U} \sum_{i=1}^{N^u} \{ \mathbb{I}(t_i^u \in \mathcal{Y}^u) \mathbb{E}[f_{N,i}^u] - \Delta_i^u \mathbb{E}[\exp(f_{N,i}^u)] \} , \tag{5.17}$$

Now the variance of $f_{N,i}^u$ does not affect $\mathbb{E}[f_{N,i}^u]$ at event times $t_i^u \in \mathcal{Y}^u$ (Lloyd et al., 2014), but only contributes to the second term through $\mathbb{E}[\exp(f_{N,i}^u)]$. This leads to instability issues during inference. We also implemented the algorithm using this transformation, but the estimated function f^u diverged even with strong prior constraints.

$\mathcal{F}_2$, as specified in (5.18), penalizes the task-specific function's deviation from the global mean function $\boldsymbol{\mu}_M^0$. This is especially important for tasks with few training

data, *e.g.*, short event streams or few event arrivals.

$$\mathcal{F}_2 = -\tfrac{U}{2}\log|\boldsymbol{K}_{MM}| + \tfrac{1}{2}\sum_{u=1}^{U}\log|\boldsymbol{\Sigma}^u| \tag{5.18}$$

$$-\textstyle\sum_{u=1}^{U}\tfrac{1}{2}tr\left[\boldsymbol{K}_{MM}^{-1}\left(\boldsymbol{\mu}^u\boldsymbol{\mu}^{u\top}+\boldsymbol{\Sigma}^u+\boldsymbol{\mu}_M^0\boldsymbol{\mu}_M^{0\top}-2\boldsymbol{\mu}^u\boldsymbol{\mu}_M^{0\top}\right)\right].$$

$\mathcal{F}_3$ involves the hyper-prior on the global function $\boldsymbol{\mu}_M^0$, with ξ controlling the belief strength:

$$\mathcal{F}_3 = \frac{1}{2}\{\log|\xi\boldsymbol{K}_{MM}^{-1}| - tr(\xi\boldsymbol{K}_{MM}^{-1}(\boldsymbol{\mu}_M^0-\boldsymbol{g})(\boldsymbol{\mu}_M^0-\boldsymbol{g})^\top)\}.$$

Having defined the variational objective, we can maximize the objective alternately *w.r.t.* to $\{\boldsymbol{\mu}^u, \boldsymbol{\Sigma}^u\}$ and $\boldsymbol{\Theta}$. In the variational E-step, we update the variational parameters $\boldsymbol{\mu}^u$ and $\boldsymbol{\Sigma}^u$ using gradient descent methods. As a practical consideration, to preserve the positive definiteness of $\boldsymbol{\Sigma}^u$, we optimize it instead over its lower Cholesky decomposition, $\boldsymbol{L}^u$, where $\boldsymbol{\Sigma}^u = \boldsymbol{L}^u\boldsymbol{L}^{u\top}$, with positiveness constraints on the diagonal elements. In the variational M-step, updates for $\boldsymbol{\mu}_M^0$ are obtained in closed form due to local conjugacy. Gradient methods are needed for updating other parameters, including locations of pseudo inputs and GP hyperparameters (see Appendix C for inference details).

When the number of tasks is large, the algorithm can be readily distributed over a cluster, with the E-step optimization procedures for tasks distributed over individual cores/machines.

5.3 Experiments

We evaluate our model on both synthetic data as well as an EHR dataset. At any time t, the basic challenge is to predict the event arrival patterns in the interval $[t, t+L]$. One baseline is a simplified model where each task is learned independently (fixing $\boldsymbol{\mu}_M^0$ in (5.9) to $\boldsymbol{g}$) to demonstrate the benefit of sharing across tasks. We refer to our proposed model as MTPP (for multitask point process model), and the simplified model as IPP (for independent point process model). Another baseline is Poisson regression (PoiR), specialized for the prediction problem considered. Here, at each time stamp t, using observations in a window of length L_P (the maximum memory length), we calculate the history features $\boldsymbol{h}(t)$, and treat them as predictors for the number of arrivals in the succeeding window $[t, t+L]$. While this is not a generative processs, the Poisson regression model can be trained using observations and tested on prediction tasks.

5.3.1 Synthetic Experiments

For the first experiment, we synthesize a rate function for each of $U = 10$ tasks using a two step process. First the rates at $M = 10$ feature locations (pseudo inputs) are generated from a common Gaussian distribution. The rate function is then drawn conditioned on these, as specified by (5.10). The features constructed at time t are the event counts in intervals $[t-5, t]$ and $[t-1, t]$. The GP hyperparameters are set as, $\lambda_1 = \lambda_2 = 4, \tau = 1.2$, and $\sigma = 0.01$, to produce moderate rate functions, *e.g.*, the empirical average rate over each event stream is between 0.05 and 1.5. Using the conditional intensity model introduced in Chapter 3, we generate event streams for each of the U tasks. We train the models using observations up to time T, and perform forward prediction.

For inference, the pseudo points were initialized using K-means clustering over

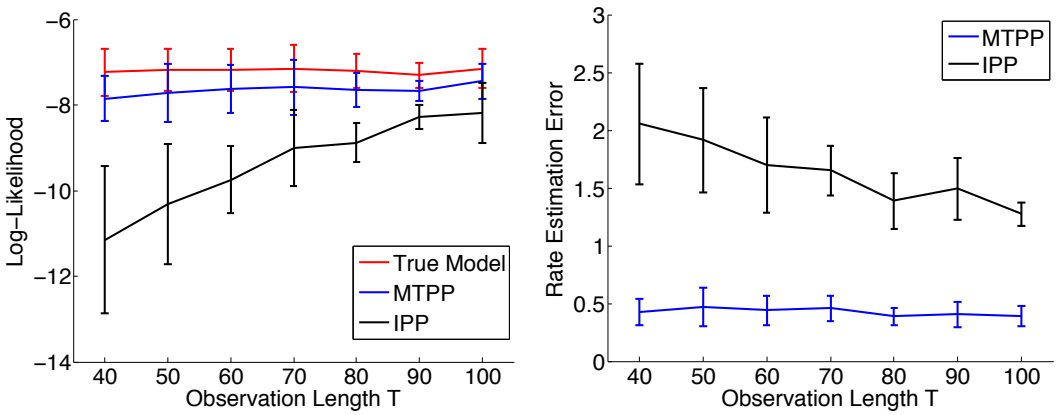

FIGURE 5.1: Left: Predictive log-likelihood per unit time; Right: mean normalized estimation error of intensity rates.

the history features appearing in the training set. The GP hyperparameters are initialized setting the length-scale parameters, $\{\lambda_p\}$, as the standard deviation of observed feature vectors, the magnitude, τ, as the standard deviation of square roots of empirical rates (computed via binning methods), and the noise term, σ^2, as a small value, $0.01\tau^2$. The algorithm converges within 20 iterations, after which the relative successive increase of the variational objective is negligible.

Model fit: We first evaluate the model and inference by comparing MTPP and IPP with the ground truth model. We vary the observation length of the training set T; the experiments for each setting are repeated for 15 runs, with mean and standard deviation reported.

The left panel of Figure 5.1 compares the predictive log-likelihood per unit time (data log-likelihood discounted by event stream length) on unseen streams. We see that with an increasing length of training observations, the predictive log-likelihood of both MTPP and IPP gradually approaches the log-likelihood obtained with the true model. When T is small, which is true in many practical scenarios, IPP cannot fit the data well while MTPP learns the model accurately. This is again illustrated in the right panel of Figure 5.1, where for both MTPP and IPP, we compute the

rate estimation error (ℓ_2-norm), normalized *w.r.t.* the true rates, and averaged over all history features occurred in the data. We observe that MTPP outperforms IPP, especially in scarce-data scenarios, demonstrating the benefit of sharing across tasks.

Forward prediction: We evaluate the prediction performance using two metrics (mirroring the two questions posed in Chapter 1). The first is a binary classification problem on whether at least one event occurs in $[t, t+L]$, and the second, estimating the distribution of event arrival counts in $[t, t+L]$. For the former, at time t, the probability of no event occurring until $t+L$ can be analytically computed as $\exp(-\int_t^{t+L} \gamma(\tau|\boldsymbol{h}(\tau))d\tau)$. This is easily solvable because $\boldsymbol{h}(\tau)$ is piecewise-constant. For the latter, because the intensity model is trained given observations, *i.e.*, the history-to-rate functions are learned, at time t, we can generate sample paths in a forward manner. Using Monte Carlo sampling, we estimate the quantities of interest, *e.g.*, the distribution of event arrival counts in $[t, t+L]$.

To better illustrate how the model performs forward prediction, the left panel of Figure 5.2, shows for one testing instance, the empirical distribution of event arrival counts obtained from 100 Monte Carlo sample paths. Qualitatively, the distribution predicted by MTPP matches the one predicted by the true model best. The actual observed count can be considered as a draw from the distribution produced by the true model. To quantitively measure the performance, we compute the Kullback-Leibler (KL) divergence between predicted distributions using the learned models and the true model, varying with observation length T. As indicated in the right panel of Figure 5.2, the KL divergence between predicted distributions using IPP and the true model is large for short streams, decreasing only with longer streams. For MTPP, there is a much smaller divergence (which decreases to a value around 0.75).

We next compare prediction results from MTPP, IPP, and PoiR, where for PoiR all tasks are trained independently. The testing instances are constructed by taking

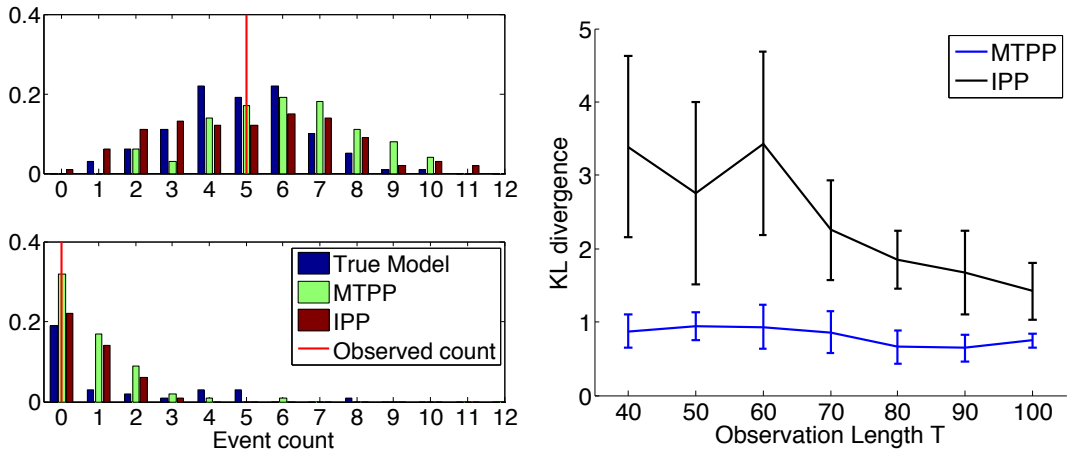

FIGURE 5.2: Left: Predicted empirical distributions of event counts in $[t, t+5]$; Right: KL divergence between predicted distributions of event counts in $[t, t+5]$ using the learned and the true model.

a sequence of snapshots on the unseen streams. In particular, we start at T with a length L sliding window, extract history features, record the event count, and move forward with stepsize $\frac{L}{2}$ until the end of the stream. Figure 5.3 demonstrates the prediction performance as a function of prediction window L. Shown in the left panel is the Area Under the Curve (AUC) for the binary problem of whether at least one event occurs in $[t, t+L]$. MTPP performs best, and is closest to the AUC achieved using the true model. The independent models perform poorly due to insufficient training data. One interesting phenomenon is that the prediction accuracy increases first and starts decreasing after a change point. This phenomenon results from a trade-off between two effects. First, for small L, the probability of no event occurring decreases from a moderate value (around 0.5) to 0, leading to an easier prediction problem. Second, as L increases further, the accumulating Poisson noise makes it harder. When L is small, the first effect dominates leading to a higher AUC. However, after a change point, the accumulated Poisson noise dominates, driving prediction accuracy down. The right panel of Figure 5.3 shows the Mean Absolute Error (MAE) of event arrival counts in $[t, t+L]$. Unlike the AUC for binary prediction, the MAE for

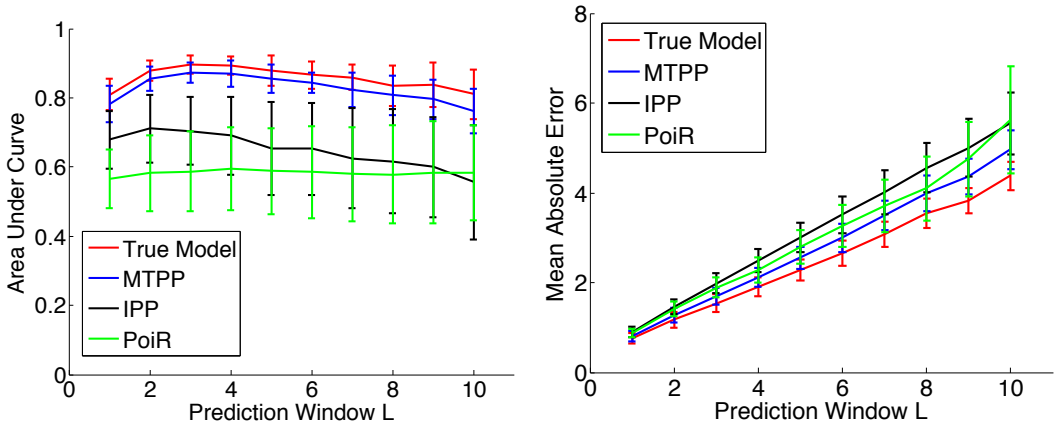

FIGURE 5.3: Left: AUC of binary prediction on whether at least one event occurs in $[t, t + L]$; Right: MAE of event arrival counts in $[t, t + L]$.

all methods increases monotonically as prediction window length increases. MTPP outperforms other baselines, achieving the MAE closest to the one obtained by the true model. The increment is linear because of the linearly accumulating Poisson noise.

5.3.2 Applications on Electrical Health Records

We now consider a real-world application involving an EHR dataset. We use the New Zealand national minimum dataset [1], covering the years 2007 through 2011 (inclusive). The data contains approximately 3.3 million inpatient visits from 1.5 million unique individuals with ages from 18 to 65. Available variables include ICD-10-AM (Australian Modification) diagnosis and procedure codes which are grouped into 22 broad categories (World Health Organization, 2010).

We focus on hospital visits for each disease category, associated with a block of ICD-10-AM billing codes (*e.g.* billing codes in the range C00 to D48 all refer to neoplasms). Treating clinical data as point process observations is a relatively novel approach to the best of our knowledge, and has only been adopted in Lasko (2014). In our experiments, for each disease category, we record all patients' visits associated with billing codes belonging to it, filter out patients with infrequent visits (fewer than 50), and split visit sequences for each patient into training and testing (split at the time stamp when half of the number of visits are observed).

After preprocessing, we have visit streams of multiple patients for each disease category, with the number of patients varying from 36 to 118, and the number of visits per patient varying from 51 to 98. In total, we analyze 6 disease types, listed in Table 5.1, mostly related with chronic diseases. For example, neoplasms include malignant and benign ones, and metabolic problems include type-I and type-II diabetes. Note we are not exploring the correlation across visits for different disease types, which is also interesting, and left as future work. The aim of this application to EHR data is to show that using the multitask point process model proposed, with a simple feature construction approach, the visit patterns for some disease types are reasonably predictable. For each patient's visit stream of a disease category, the

[1] http://www.health.govt.nz/nz-health-statistics

features constructed include the number of his/her hospital visits for this disease in the previous week, month, and six months. For inference, we set the number of pseudo inputs $M = 15$, and the initialization procedure is the same as in the synthetic experiments. The algorithm generally converges in tens of iterations.

Similar with the above synthetic experiments, we evaluate the algorithm from two perspectives: model fit and predictive ability. For model fit, we evaluate the data log-likelihood, comparing with two state-of-the-art approaches using GP modulated renewal processes: a direct inference method in Lasko (2014) and a thinning approach in Rao and Teh (2011). For the prediction tasks, unlike the synthetic experiments, we do not have access to the ground truth model, *i.e.*, knowledge about heterogeneity across population/tasks. Thus, we compare MTPP with PoiR using the same model for all patients. IPP and PoiR with individual models for each patient were also tested, but because the patients' visits are sparse, the individual model learned with this scarce training data performs poorly, hence only results of pooled MTPP and PoiR are reported.

Model Fit: GP modulated renewal processes are trained using both direct inference as in Lasko (2014) and thinning approach as in Rao and Teh (2011), for each patient's visit sequence in the neoplasm category. We then compare these two methods with MTPP on the data log-likelihood per time unit (day) over all sequences. As shown in the top-left panel of Figure 5.4, MTPP is comparable with the other two methods, with a smaller variation due to the sharing across patients. Shown in the top-right and bottom panels are the mean intensity functions inferred of anonymous patients (error bars are omitted for clarity), each associated with the raster plot of the hospital visits in the bottom. Relative dates, instead of calendar dates are used for confidentiality. As can be seen, the two temporal GP methods infer a smooth function modeling the rate, which fails to capture sudden changes. For example, see the small-to-large change rate around the 33rd month in the top-right panel. Also

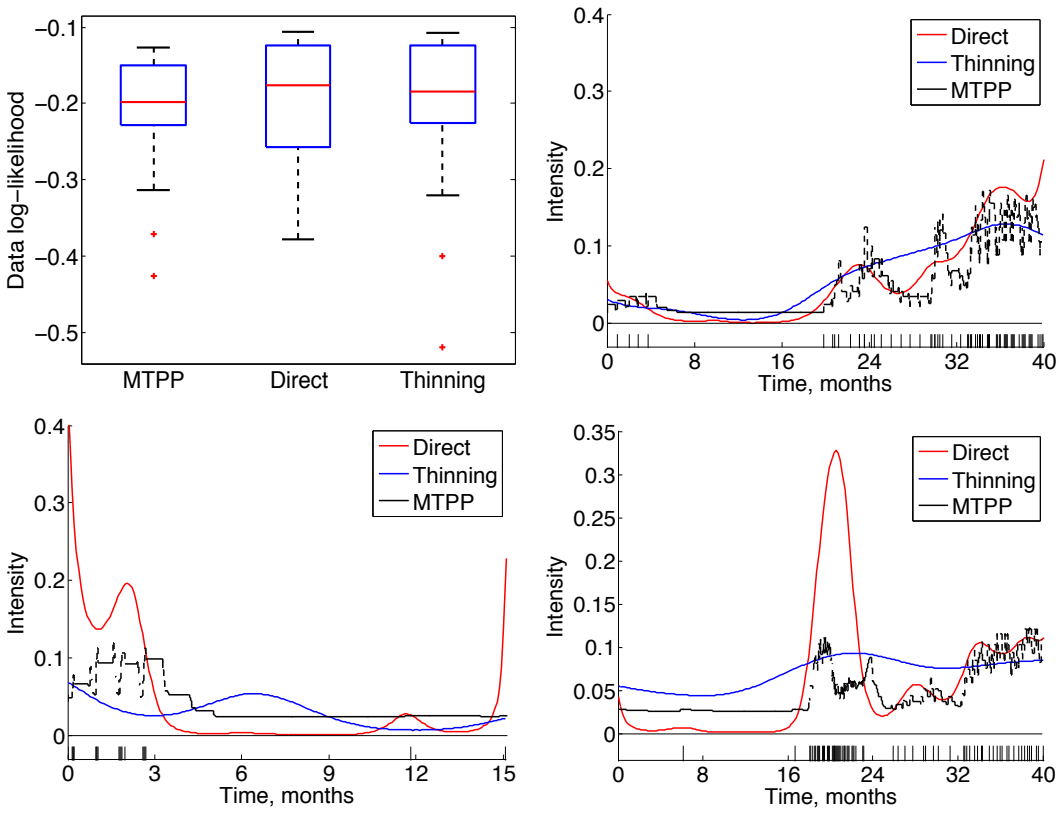

FIGURE 5.4: Model fit results. Top-left: Comparison of data log-likelihood per day of MTPP, direct inference, and thinning approach; Top-right and bottom row: Intensity functions inferred of anonymous patients' arrival sequences.

note in the bottom-left panel, MTPP performs well at the beginning and ending of the sequence, with no boundary effects of GPs. Such differences are consistent over the dataset. Note that the intensity function inferred via our method is piecewise-constant, with change points located where recent history changes. The goodness of fit suggests this current rate's dependence on history enables forward prediction.

Forward Prediction: Figure 5.5 (left panel) shows the learned history-to-rate mapping function globally shared by all patients in the neoplasm category. In particular, we plot the value of the function in terms of two features, namely, number of arrivals in windows $[t-7, t]$ and $[t-30, t]$ (the previous week and month). As indicated, a patient with many visits in a month and few visits in a week has a larger

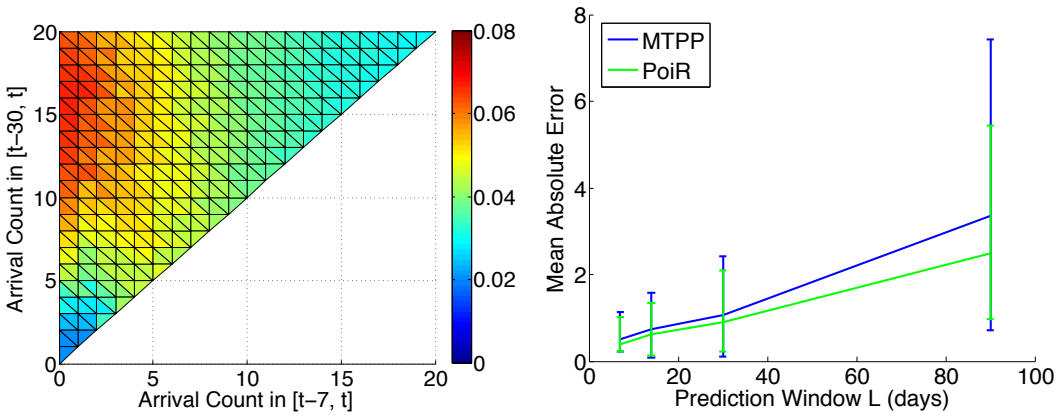

FIGURE 5.5: Neoplasms results. Left: Global rate function inferred by MTPP as a function of the number of arrivals in windows $[t-7, t]$ and $[t-30, t]$; Right: MAE as a function of the prediction window length, L, in days.

rate of a revisit.

The right panel in Figure 5.5 shows that for neoplasms, the MAE increases monotonically as the prediction window length increases up to 90 days. This is consistent with the results from synthetic data shown in Figure 5.3. Interestingly, because PoiR requires a separate model for each prediction interval, no noise variance accumulates in the model, and the increase in MAE results from the data itself. On the other hand, we just train MTPP once and generate sample paths for different window lengths, resulting in a Poisson noise that increases linearly with prediction interval. As a result, for MTPP the MAE diverges faster than for PoiR as L increases.

Results in Table 5.1 show AUC values for the 6 disease categories considered. Two settings for the length of the prediction windows, namely, 1 week and 1 month are evaluated. We see that consistent with the results for artificial data in Figure 5.3, AUCs tend to increase for moderate sizes of L. We also verified (results not shown) that further increasing the prediction window size has an accordingly negative impact on the AUCs. In fact, predictions become no better than random as L approaches 6 months for some disease types. When comparing MTPP to PoiR, we observe that the former outperforms the latter in 4 of 6 disease types. This supports the hypothesis

Table 5.1: AUC of binary predictions of events occurring in weekly and monthly windows for 6 disease types.

	1 week		1 month	
Disease Type	MTPP	PoiR	MTPP	PoiR
Neoplasms	**0.7379**	0.7249	**0.8136**	0.8058
Metabolic	**0.6807**	0.6170	**0.6778**	0.6195
Nervous	0.6926	**0.7241**	**0.7978**	0.7167
Circulatory	**0.6807**	0.6778	0.6170	**0.6195**
Respiratory	0.5733	**0.6302**	0.6308	**0.6322**
Digestive	**0.6050**	0.5562	**0.6555**	0.6170

that the changes of hospital visit patterns among the population are disease specific. However, correlations across diseases and clustering of patient subpopulations could be exploited, and are left as future work.

6

Conclusions

In this thesis, we formulate and tackle the problem of point process prediction. First in the single event stream scenario, we explore the feasibility of predictive models built upon conditional intensity models via a GP construction. Then we mainly focus on the problem of analyzing multiple streaming point processes in a multi-task setting, and propose a simple predictive strategy. The proposed model, using hierarchical GPs, leverages information across the tasks in a nonparametric manner, which is demonstrated in experimental results on synthetic and real-world datasets.

Exploring multi-task marked point processes, designing richer feature construction approaches for predictive models, and clustering tasks/event streams based on their arrival patterns, are left as interesting open problems.

Appendix

A. Hierarchical Gaussian process construction

Using the construction in (5.1) and (5.2), we can integrate over $\boldsymbol{\mu}_N^0$ to obtain the marginal prior distribution for $\boldsymbol{f}_N^u$ as

$$\boldsymbol{f}_N^u \sim \mathcal{N}\left(\boldsymbol{g}, \left(1 + \frac{1}{\xi}\right)\boldsymbol{K}_{NN}\right) . \tag{1}$$

Using the construction in (5.8)-(5.10), we can integrate out $\boldsymbol{f}_M^u$ using (5.9) and (5.10), obtaining,

$$\boldsymbol{f}_N^u|\boldsymbol{\mu}_M^u \sim \mathcal{N}\Bigg(\boldsymbol{g} + \boldsymbol{K}_{NM}^u\boldsymbol{K}_{MM}^{-1}(\boldsymbol{\mu}_M^0 - \boldsymbol{g}),$$

$$\boldsymbol{K}_{NN}^u + \frac{1}{\xi}\left(\boldsymbol{K}_{NN}^u - \boldsymbol{K}_{NM}^u\boldsymbol{K}_{MM}^{-1}{\boldsymbol{K}_{NM}^u}^\top\right)\Bigg) .$$

Then, integrating out $\boldsymbol{\mu}_M^0$ using (5.8), we can get the same marginal prior as in (1),

$$\boldsymbol{f}_N^u \sim \mathcal{N}\left(\boldsymbol{g}, \left(1 + \frac{1}{\xi}\right)\boldsymbol{K}_{NN}^u\right) . \tag{2}$$

The only difference between (1) and (2) is that we need a realization of the function at possible features from all tasks in (1), while only the features from task u need to be specified in (2).

B. Confluent hypergeometric functions

When $|x|$ is small, for example $|x| \leqslant 30$, we use the power series to compute the confluent hypergeometric function:

$$ {}_1F_1(a,b,x) = \sum_{k=0}^{\infty} \frac{(a)_k x^k}{(b)_k k!} \,, $$

where $(a)_0 = 1$, $(a)_k = \prod_{j=0}^{k-1}(a+j)$. In practice, the summation can be terminated at a sufficiently large number to achieve a given error tolerance level.

When $|x|$ is large, for example $|x| > 30$, we use the following computation (Thompson, 1997):

$$ {}_1F_1(a,b,x) = \frac{\Gamma(b)(-x)^{-a}}{\Gamma(b-1)} \sum_{k=0}^{\infty} \frac{(a)_k (a+a-b)_k}{k!(-x)^k} \,. $$

Having ${}_1F_1(a,b,x)$, the gradient $\tilde{G}(x) = G(0, \frac{1}{2}, z)$ can be numerically computed, where $G(a,b,x) = \frac{\partial_1 F_1(a,b,x)}{\partial a}$ (Ancarani and Gasaneo, 2008).

Finally, the expectation needed in (5.14) during variational E-step can be estimated as (Lloyd et al., 2014):

$$ \mathbb{E}[\log(f^u_{N,i})^2] = -\tilde{G}(-\frac{b_i^{u2}}{2B^u_{ii}}) + \log(\frac{B^u_{ii}}{2}) - const \,. $$

C. Key quantities in variational inference

Variational E-step:

$$\frac{d\mathcal{F}_1}{d\boldsymbol{\mu}^u} = \sum_{u=1}^{U}\sum_{i=1}^{N^u}\Big\{-\Delta_i^u\big[2b_i^u\boldsymbol{K}_{MM}^{-1}\boldsymbol{K}_{NM,i\cdot}^{u\top}\big] + \mathbb{I}(t_i^u \in \mathcal{T}^u)\left[\tilde{G}'\left(-\frac{(b_i^u)^2}{2B_{ii}^u}\right)\frac{b_i^u}{B_{ii}^u}\boldsymbol{K}_{MM}^{-1}\boldsymbol{K}_{NM,i\cdot}^{u\top}\right]\Big\},$$

$$\frac{d\mathcal{F}_1}{d\boldsymbol{\Sigma}^u} = \sum_{u=1}^{U}\sum_{i=1}^{N^u}\Big\{-\Delta_i^u[\boldsymbol{K}_{MM}^{-1}\boldsymbol{K}_{NM,i\cdot}^{u\top}\boldsymbol{K}_{NM,i\cdot}^{u}\boldsymbol{K}_{MM}^{-1}] + \mathbb{I}(t_i^u \in \mathcal{T}^u)\left[\left(-\tilde{G}'\left(-\frac{(b_i^u)^2}{2B_{ii}^u}\right)\frac{(b_i^u)^2}{2(B_{ii}^u)^2} + \frac{1}{B_{ii}^u}\right)\boldsymbol{K}_{MM}^{-1}\boldsymbol{K}_{NM,i\cdot}^{u\top}\boldsymbol{K}_{NM,i\cdot}^{u}\boldsymbol{K}_{MM}^{-1}\right]\Big\},$$

$$\frac{d\mathcal{F}_2}{d\boldsymbol{\mu}^u} = -\boldsymbol{K}_{MM}^{-1}(\boldsymbol{\mu}^u - \boldsymbol{\mu}_M^0),$$

$$\frac{d\mathcal{F}_2}{d\boldsymbol{\Sigma}^u} = -\frac{1}{2}\boldsymbol{K}_{MM}^{-1} + \frac{1}{2}\boldsymbol{\Sigma}^{u-1},$$

$$\frac{d\mathcal{F}_3}{d\boldsymbol{\mu}^u} = 0,$$

$$\frac{d\mathcal{F}_3}{d\boldsymbol{\Sigma}^u} = 0.$$

Variational M-step:

Updates for $\boldsymbol{\mu}_M^0$ can be derived in closed form as follows,

$$\hat{\boldsymbol{\mu}}_M^0 = \frac{1}{\xi + U}\left(\xi\boldsymbol{g} + \sum_{u=1}^{U}\boldsymbol{\mu}^u\right).$$

Gradient methods are needed for updating other parameters (denoted as θ_k), including pseudo input positions and GP hyper-parameters, with key quantities summarized below:

$$\frac{d\mathcal{F}_1}{d\theta_k} = \sum_{u=1}^{U}\sum_{i=1}^{N^u}\left\{\mathbb{I}(t_i^u \in \mathcal{T}^u)\left[-\tilde{G}'\left(-\frac{(b_i^u)^2}{2B_{ii}^u}\right)\left(-\frac{1}{2(B_{ii}^u)^2}\right)\right.\right.$$

$$\times\left\{2B_{ii}^u b_i^u(\boldsymbol{\mu}^u - \boldsymbol{g})^T\frac{d\boldsymbol{K}_{MM}^{-1}\boldsymbol{K}_{NM,i\cdot}^{u}{}^{\top}}{d\theta_k} - (b_i^u)^2\frac{dB_{ii}^u}{d\theta_k}\right\}$$

$$\left.\left.+\frac{1}{B_{ii}^u}\frac{dB_{ii}^u}{d\theta_k}\right]-\Delta_i^u\left[2b_i^u(\boldsymbol{\mu}^u - \boldsymbol{g})^T\frac{d\boldsymbol{K}_{MM}^{-1}\boldsymbol{K}_{NM,i\cdot}^{u}{}^{\top}}{d\theta_k} + \frac{dB_{ii}^u}{d\theta_k}\right]\right\},$$

$$\frac{dB_{ii}^u}{d\theta_k} = \frac{d\boldsymbol{K}_{NN,ii}}{d\theta_k} - \frac{d\boldsymbol{K}_{NM,i\cdot}^{u}\boldsymbol{K}_{MM}^{-1}\boldsymbol{K}_{NM,i\cdot}^{u}{}^{\top}}{d\theta_k}$$

$$+\ 2(\boldsymbol{\Sigma}^u\boldsymbol{K}_{MM}^{-1}\boldsymbol{K}_{NM,i\cdot}^{u}{}^{\top})^{\top}\frac{d\boldsymbol{K}_{MM}^{-1}\boldsymbol{K}_{NM,i\cdot}^{u}{}^{\top}}{d\theta_k},$$

$$\frac{d\mathcal{F}_2}{d\theta^k} = -\frac{U}{2}tr\left(\boldsymbol{K}_{MM}^{-1}\frac{d\boldsymbol{K}_{MM}}{d\theta_k}\right)+\frac{1}{2}tr\left(\boldsymbol{K}_{MM}^{-1}\right.$$

$$\left.\times\sum_{u=1}^{U}\left(\boldsymbol{\mu}^u\boldsymbol{\mu}^{u\top} + \boldsymbol{\Sigma}^u + \boldsymbol{\mu}_M^0\boldsymbol{\mu}_M^{0}{}^{\top} - 2\boldsymbol{\mu}^u\boldsymbol{\mu}_M^{0}{}^{\top}\right)\boldsymbol{K}_{MM}^{-1}\frac{d\boldsymbol{K}_{MM}}{d\theta_k}\right),$$

$$\frac{d\mathcal{F}_3}{d\theta^k} = -\frac{1}{2}tr\left(\boldsymbol{K}_{MM}^{-1}\frac{d\boldsymbol{K}_{MM}}{d\theta_k}\right)$$

$$+\frac{1}{2}tr\left(\xi\boldsymbol{K}_{MM}^{-1}(\boldsymbol{\mu}_M^0 - \boldsymbol{g})(\boldsymbol{\mu}_M^0 - \boldsymbol{g})^{\top}\boldsymbol{K}_{MM}^{-1}\frac{d\boldsymbol{K}_{MM}}{d\theta_k}\right).$$

Bibliography

Adams, R. P., Murray, I., and MacKay, D. J. C. (2009), "Tractable Nonparametric Bayesian Inference in Poisson Processes with Gaussian Process Intensities," in *Proceedings of the 26th International Conference on Machine Learning (ICML)*.

Amarasingham, R., Moore, B. J., Tabak, Y. P., Drazner, M. H., Clark, C. A., Zhang, S., Reed, W. G., Swanson, T. S., Ma, Y., and Halm, E. A. (2010), "An automated model to identify heart failure patients at risk for 30-day readmission or death using electronic medical record data," *Medical care*, 48, 981–988.

Ancarani, L. U. and Gasaneo, G. (2008), "Derivatives of any order of the confluent hypergeometric function," *Journal of Mathematical Physics*, 49.

Beal, M. J. (2003), "Variational algorithms for approximate Bayesian inference," Ph.D. thesis, U. London.

Diggle, P. (1985), "A kernel method for smoothing point process data," *Applied statistics*, pp. 138–147.

Gunawardana, A., Meek, C., and Xu, P. (2011), "A model for temporal dependencies in event streams," in *Advances in Neural Information Processing Systems*, pp. 1962–1970.

Kulkarni, J. E. and Paninski, L. (2007), "Common-input models for multiple neural spike-train data," *Network: Computation in Neural Systems*, 18, 375–407.

Lasko, T. A. (2014), "Efficient Inference of Gaussian Process Modulated Renewal Processes with Application to Medical Event Data," in *Uncertainty in Artificial Intelligence*.

Lian, W., Rao, V., Eriksson, B., and Carin, L. (2015), "A Multitask Point Process Predictive Model," in *Proceedings of the 32st International Conference on Machine Learning (ICML-15)*, JMLR Workshop and Conference Proceedings.

Lloyd, C., Gunter, T., Osborne, M. A., and Roberts, S. J. (2014), "Variational Inference for Gaussian Process Modulated Poisson Processes," *arXiv:1411.0254*.

Murray, I. and Adams, R. P. (2010), "Slice sampling covariance hyperparameters of latent Gaussian models," in *Advances in Neural Information Processing Systems*, pp. 1732–1740.

Pillow, J. W., Shlens, J., Paninski, L., Sher, A., Litke, A. M., Chichilnisky, E., and Simoncelli, E. P. (2008), "Spatio-temporal correlations and visual signalling in a complete neuronal population," *Nature*, 454, 995–999.

Rad, K. R. and Paninski, L. (2011), "Information rates and optimal decoding in large neural populations," in *Advances in neural information processing systems*, pp. 846–854.

Rajaram, S., Graepel, T., and Herbrich, R. (2005), "Poisson-networks: A model for structured point processes," in *Proceedings of the 10th International Workshop on Artificial Intelligence and Statistics*, pp. 277–284, Citeseer.

Rao, V. and Teh, Y. W. (2011), "Gaussian process modulated renewal processes," in *Advances in Neural Information Processing Systems*, pp. 2474–2482.

Snelson, E. and Ghahramani, Z. (2006), "Sparse Gaussian processes using pseudo-inputs," in *Advances in neural information processing systems*, pp. 1257–1264.

Thompson, W. J. (1997), *Atlas for Computing Mathematical Functions: An Illustrated Guide for Practitioners with Programs in FORTRAN and Mathematica with Cdrom*, John Wiley & Sons, Inc.

Titsias, M. (2009), "Variational learning of inducing variables in sparse Gaussian processes," in *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics*.

Weiss, J. C. and Page, D. (2013), "Forest-based point process for event prediction from electronic health records," in *Machine learning and knowledge discovery in databases*, pp. 547–562, Springer.

World Health Organization, G. (2010), "International Classification of Diseases, 10th Revision (ICD-10)," .

Xu, L., Duan, J. A., and Whinston, A. (2014), "Path to purchase: A mutually exciting point process model for online advertising and conversion," *Management Science*.

Yu, K., Tresp, V., and Schwaighofer, A. (2005), "Learning Gaussian processes from multiple tasks," in *Proceedings of the 22nd international conference on Machine learning*, pp. 1012–1019.