Nonstationary Multivariate Gaussian Processes for Electronic Health Records

Rui Meng[+]   Braden Soper[*]   Herbert Lee[+]   Vincent X. Liu[**] John D. Greene[**]   Priyadip Ray[*]

+ University of California, Santa Cruz
* Lawrence Livermore National Laboratories
** Kaiser Permanente Division of Research

Abstract

We propose multivariate nonstationary Gaussian processes for jointly modeling multiple clinical variables, where the key parameters, length-scales, standard deviations and the correlations between the observed output, are all time dependent. We perform posterior inference via Hamiltonian Monte Carlo (HMC). We also provide methods for obtaining computationally efficient gradient-based maximum a posteriori (MAP) estimates. We validate our model on synthetic data as well as on electronic health records (EHR) data from Kaiser Permanente (KP). We show that the proposed model provides better predictive performance over a stationary model as well as uncovers interesting latent correlation processes across vitals which are potentially predictive of patient risk.

1 Introduction

The large-scale collection of electronic health records (EHRs) offers the promise of accelerating clinical research for understanding disease progression and improving predictive modeling of patient clinical outcomes [Cheng et al., 2017, Jung and Shah, 2015]. Typically, EHR data consists of rich patient information, including but not limited to, demographic information, vital signs, laboratory results, diagnosis codes, prescriptions and treatments. However, it is extremely challenging to develop models for EHR data. Contributing to these challenges are data

quality, data heterogeneity, complex dependencies across multiple time series, irregular sampling rates, systematically missing data, and statistical nonstationarity [Cheng et al., 2017, Ghassemi et al., 2015, Futoma et al., 2017a].

Despite these challenges, the promise of leveraging EHR data to improve patient outcomes has resulted in an explosive growth of research in the past decade. While the existing literature addresses many of the challenges in modeling EHR data, such as irregular sampling rates [Cheng et al., 2017, Li and Marlin, 2016, Futoma et al., 2017a], missing data [Ghassemi et al., 2015] and the modeling of complex dependencies across multiple streams of clinical data [Cheng et al., 2017, Ghassemi et al., 2015], violations of stationarity in EHR data [Jung and Shah, 2015] has received less attention. In this paper, we propose a novel statistical framework based on multivariate Gaussian processes (GPs) to model both nonstationarity and heteroscedasticity in EHR data. We explore both model predictive performance as well as inferred nonstationary correlation patterns across different clinical variables. Inference for the proposed model is performed in a fully Bayesian manner, providing full uncertainty quantification on predictions and all model parameters, such as time-varying correlations across clinical variables.

While biomedical processes can be both multivariate and nonstationary, models which handle both features have not been explored in the context of EHR data, to the best of our knowledge. Sepsis is a prime example of a disease in which correlated multivariate output and nonstationarity may be critical for early identification. Sepsis has been shown to exhibit highly nonstationary variations in the vitals of patients [Cao et al., 2004] while the cross-correlation of these vitals has been shown to be predictive of early onset [Fairchild et al., 2016]. While both multivariate and nonstationary models have been proposed for EHR

data, to the best of our knowledge ours is the first model for EHR data which is both nonstationary and multivariate.

We demonstrate our proposed framework by modeling a large EHR dataset composed of emergency department (ED) hospitalization episodes from Kaiser Permanente (KP). The patients were suspected to have an infection and, in a subset of cases, met the clinical criteria for sepsis [Fohner et al., 2019, Seymour et al., 2016]. Sepsis is a life-threatening organ dysfunction arising from a dysregulated host response to infection, affecting at least 30 million patients worldwide and resulting in 5 million deaths each year [Fohner et al., 2019, Klompas and Rhee, 2016].

We apply our proposed approach to jointly model systolic pressure, diastolic pressure, heart rate, respiratory rate, pulse pressure and oxygen saturation levels. We demonstrate improved model prediction performance and uncertainty quantification over the state-of-art. Since changes in cross-correlations across vital signs is often an indicator of onset of sepsis [Fairchild et al., 2016], we also explore the inferred cross-correlations across the vitals and their relationship with the hourly LAPS2 scores (LAPS2 is a KP specific measure of acute disease burden and is an indicator of the risk state of a patient) [Escobar et al., 2013].

2 Related Work

Gaussian processes have a long history in both spatio-temporal statistical modeling [Luttinen and Ilin, 2009] and in machine learning [Rasmussen and Williams, 2005a]. With the increasing use of EHR data to improve patient health outcomes, there has been an increased application of Gaussian processes to modeling EHR data. Our use of Gaussian processes is motivated by their flexibility in handling nonstationary and correlated multivariate data, which have been extensively applied to spatio-temporal statistical modeling [Cressie and Wikle, 2011]. In this section we briefly overview the recent literature on the use of Gaussian processes in EHR modeling and point out the main contributions of this paper.

EHR data consists of multiple correlated measurements taken over time. As such, multi-output or multi-task Gaussian processes (MTGPs) have been proposed as appropriate models for EHR data. A MTGP framework for modeling the correlation across multiple physiological time series was first proposed in [Drichen et al., 2014]. In [Ghassemi et al., 2015] the inferred hyper-parameters from a MTGP model were used as compact latent representations used to pre-

dict severity of illness in ICU patients. Online patient state prediction [Cheng et al., 2017] and online patient risk assessment [Alaa et al., 2018] were both proposed via a MTGP framework based on large-scale EHR data sets. Personalized treatment effects were predicted via MTGPs in [Alaa and van der Schaar, 2017]. Online MTGPs were combined with RNN classifiers for early sepsis prediction in hospital patients in [Futoma et al., 2017b, Futoma et al., 2017a]. While each of the above approaches are able to learn a correlation structure both within and between clinical time series, all models are both stationary and homoscedastic.

Because hospitalized patients can go through drastic physiological changes in short periods of time, nonstationary models are needed for EHR data [Cao et al., 2004, Jung and Shah, 2015, Hripcsak et al., 2015]. One effect of the biological nonstationarity is highly irregular sampling rates for EHR data. This is due to attending healthcare providers adjusting the sampling rates in response to observable changes in patient state. Nonstationary Gaussian processes have been proposed as means of correcting for these highly irregular sampling rates via time warping in [Lasko, 2015]. While this model does directly model nonstationarity in the clinical time series, it does not directly model heteroscedastic nor correlated multivariate data.

Other than directly modeling EHR data, Gaussian processes have been utilized in a variety of ways with EHR data. They have been used to smooth and regularize data for blackbox optimizers [Li and Marlin, 2016, Futoma et al., 2017b, Futoma et al., 2017a, Lasko et al., 2013], as priors for latent hazard functions in modulated point processes [Lasko, 2014], and as components of hierarchical generative models [Schulam et al., 2015].

The flexibility and expressiveness of Gaussian processes clearly offer a powerful framework for modeling complex EHR data. And while Gaussian process models have been proposed to handle either nonstationarity or correlated multivariate EHR data, to the best of our knowledge there has been no attempt to model both nonstationarity and correlated multivariate EHR data. Furthermore no heteroscedastic Gaussian process models have been proposed for EHR data previously. In the following section we present a novel multivariate, nonstationary, heteroscedastic Gaussian process model capable of handling complex EHR data.

3 Multivariate Nonstationary Gaussian Processes

Inpatient clinical time series data is composed of measurements of multiple correlated patient vital signs. Furthermore, the statistical properties of the data may not be constant across time due to physiological changes from acute disease onset. For these reasons we propose a multivariate nonstationary Gaussian process (MNGP) to model EHR data. Importantly, it is the first such multivariate Gaussian process to model both nonstationarity in the length-scale parameter, signal variance and the covariance matrix between dimensions of the output.

We briefly review some basic properties of both multivariate Gaussian processes in the following section. In section 3.2 we present our Generalized Nonstationary Multivariate Gaussian Process model in detail.

3.1 Background

A Multivariate Gaussian process (MGP) defines a distribution over multivariate functions $\boldsymbol{f}(t) = (f_1(t), \ldots f_M(t))^T$. For any collection of inputs $t_1, \ldots, t_N$, the function values $\boldsymbol{f}_n = \boldsymbol{f}(t_n)$ follow a multivariate normal distribution

$$\vec{\boldsymbol{f}} \equiv \text{vec}([\boldsymbol{f}_1, \ldots, \boldsymbol{f}_N]^T) \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{K}^f)$$

where vec is the vectorization operator. The covariance matrix $\boldsymbol{K}^f$ is generated from a covariance function K^f such that for any two inputs $t, t' \in \{t_1, ..., t_N\}$ and any two dimensions $m, m' \in \{1, ..., M\}$, the covariance between the values $\boldsymbol{f}(t)[m]$ and $\boldsymbol{f}(t')[m']$ is given by $K^f(t, t', m, m')$. A MGP is said to be separable when a decomposition exists such that $K^f(t, t', m, m') = B(m, m')K(t, t')$ for some functions B of the dimension indices only and K of the input dimensions only. In matrix notation this is equivalent to $\boldsymbol{K}^f = \boldsymbol{B} \otimes \boldsymbol{K}$, where the covariance matrix $\boldsymbol{B} \in \mathbb{R}^{M \times M}$ summarizes the relations across output dimensions, the covariance matrix $\boldsymbol{K}$ summarizes the relations across input, and $\otimes$ denotes the Kronecker product. A typical separable model would be to model the covariance matrix $\boldsymbol{B}$ directly, say by its Cholesky decomposition, $\boldsymbol{B} = \boldsymbol{L}\boldsymbol{L}^T$, while $\boldsymbol{K}$ is parameterized by a kernel function. A commonly used stationary kernel is the square exponential or radial basis function (RBF) kernel

$$k_{\text{RBF}}(t, t') = \sigma^2 \exp\left(-\frac{(t - t')^2}{2d^2}\right). \tag{1}$$

where the signal variance σ corresponds to the range scale of function and the length-scale d encodes how fast the function values can change with respect to the distance $|t - t'|$.

A continuous-time stochastic process is said to be wide-sense stationary if both its mean and autocovariance functions do not vary with time. For a zero-mean Gaussian process this property reduces to having a stationary covariance function, namely a positive definite kernel $k(\boldsymbol{x}, \boldsymbol{x}')$ which is only a function of $\boldsymbol{x} - \boldsymbol{x}'$. Thus for our purposes a nonstationary zero-mean Gaussian process has a covariance function which is not stationary.

3.2 A Generalized Nonstationary Multivariate Gaussian Process

We now present our Generalized Nonstationary Multivariate Gaussian Process (GNMGP) in detail. The hierarchical representation of the model is as follows:

$$\begin{aligned}
\boldsymbol{y}(t) &= \boldsymbol{L}(t)\boldsymbol{g}(t) + \boldsymbol{\epsilon}(t), \\
\boldsymbol{\epsilon}(t) &\sim \mathcal{N}(0, \sigma_{err}^2 I), \\
g_d(t) &\overset{iid}{\sim} \text{GP}(0, K) \qquad d = 1, 2, ..., M, \\
L_{ij}(t) &\sim \text{GP}(\mu_L, K_L) \qquad i \geq j \\
\sigma_{err}^2 &\sim \text{IG}(a, b).
\end{aligned} \tag{2}$$

We utilize a Gibbs kernel for the independent GPs where nonstationarity is achieved by placing a GP prior on the log length-scale process.

$$K(t, t') = \sqrt{\frac{2\ell(t)\ell'(t)}{\ell(t)^2 + \ell(t')^2}} \exp\left(-\frac{(t - t')^2}{\ell(t)^2 + \ell(t')^2}\right)$$

$$\log(\ell(t)) \sim \text{GP}(\mu_{\tilde{\ell}}, K_{\tilde{\ell}}(t, t'))$$

Finally $\mu_{\tilde{\ell}}$ and $K_{\tilde{\ell}}(t, t')$ are treated as hyperparameters of the model and can be chosen appropriately for the application.

Because $\boldsymbol{L}(t)$ is a lower triangular matrix and the components of $\boldsymbol{g}(t)$ are iid, the covariance function of the resulting multivariate GP $\boldsymbol{f}(t) = \boldsymbol{L}(t)\boldsymbol{g}(t)$ is given by

$$\begin{aligned}
K^f(t, t', m, m') &= [\boldsymbol{L}(t)\text{cov}(\boldsymbol{g}(t), \boldsymbol{g}(t'))\boldsymbol{L}^T(t')]_{mm'} \\
&= K(t, t')[\boldsymbol{L}(t)\boldsymbol{L}^T(t')]_{m,m'}. \tag{3}
\end{aligned}$$

The proposed GNMGP model can be understood as generalizations of existing GP models. For example, [Gelfand et al., 2004] proposed a nonstationary model by considering a spatially varying covaraince relation but did not emphasize the input-dependent length-scale, which is crucial in real applications. When using a matrix-variate inverse Wishart spatial process for the covariance matrix $\boldsymbol{B}(t) = \boldsymbol{L}^T(t)\boldsymbol{L}(t)$ and a stationary kernel, such as the RBF kernel, for the temporal kernel K, our model reduces to the spatially varying linear model.

In general the above model will be nonseparable, meaning the covariance function cannot be decomposed into components that are functions of either

time or dimension, but not both. A a special case of the above model that is separable can be derived as follows. Suppose $\boldsymbol{L}(t) = \sigma(t)\boldsymbol{L}$ for some positive function $\sigma(t)$ and constant lower-triangular matrix L. Letting $\boldsymbol{B} = \boldsymbol{L}\boldsymbol{L}^t$ we see from (3), that the GNMGP kernel function reduces to $K^f(t, t', m, m') = \sigma(t)\sigma(t')K(t, t')[\boldsymbol{B}]_{m,m'}$, which is separable. To finish the specification of this model we assume $\log(\sigma(t)) \sim \text{GP}(\mu_{\tilde{\sigma}}, K_{\tilde{\sigma}}(t, t'))$ and $L_{ij} \sim \mathcal{N}(\mu_L, \Sigma_L)$ for $i \geq j$. As before $\mu_{\tilde{\sigma}}, K_{\tilde{\sigma}}(t, t'), \mu_L, \Sigma_L$ are hyperparameters of the model to be chosen accordingly to the application.

We note that [Heinonen et al., 2016] proposed a fully nonstationary univariate kernel by extending the standard Gibbs kernel to include input-dependent signal variance and input-dependent signal noise. When considering input dependent length-scale and signal-variance, our proposed separable kernel can be seen as a multivariate extension of the nonstationary model proposed in [Heinonen et al., 2016].

Finally, we note that using different kernel functions is possible. For example, [Paciorek and Schervish, 2004] proposed a class of nonstationary kernels for univariate output with a Matérn kernels. Extensions of our model to the Matérn kernel is straightforward, providing a multivariate output alternative to the model in [Paciorek and Schervish, 2004].

4 Inference

We propose two inference approaches, maximum a posteriori (MAP) and Hamiltonian Monte Carlo inference. This section discusses the case of complete data, which means at each location or time stamp t, all observations $\boldsymbol{y}_t$ are available. Inference for incomplete data, where not all $\boldsymbol{y}_t$ are available at each t, follows from standard Gaussian process methods for marginalizing over missing data [Rasmussen and Williams, 2005b]. Note that for ease of exposition we introduce the following notation: $\tilde{\ell}(t) \equiv \log(\ell(t))$ and $\tilde{\sigma}(t) \equiv \log(\sigma(t))$.

4.1 Maximum A Posteriori (MAP)

This section considers maximum a posteriori inference for both models. In the separable model setting, model parameters $\sigma_{err}^2, \boldsymbol{L}, \tilde{\ell}, \tilde{\boldsymbol{\sigma}}$ are of interest. The marginal posteriors of these parameters are

$$
\begin{aligned}
p(\sigma_{err}^2, \boldsymbol{L}, \tilde{\ell}, \boldsymbol{\sigma}|\boldsymbol{y}, \boldsymbol{t}) &= \int p(\boldsymbol{f}, \sigma_{err}^2, \boldsymbol{L}, \boldsymbol{g}, \tilde{\ell}, \boldsymbol{\sigma}|\boldsymbol{y}, \boldsymbol{t}) d\boldsymbol{g} d\boldsymbol{f} \\
&\propto \mathcal{N}(\vec{\boldsymbol{y}}|\boldsymbol{0}, \boldsymbol{K}^f + \sigma_{err}^2 I) \prod_{i \geq j} \mathcal{N}(\boldsymbol{L}_{ij}|0, c) \mathcal{N}(\tilde{\boldsymbol{\sigma}}|\boldsymbol{\mu}_{\tilde{\boldsymbol{\sigma}}}, \boldsymbol{K}_{\tilde{\boldsymbol{\sigma}}}) \\
&\quad \mathcal{N}(\tilde{\ell}|\boldsymbol{\mu}_{\tilde{\ell}}, \boldsymbol{K}_{\tilde{\ell}}) \text{IG}(\sigma_{err}^2|a, b) \quad (4)
\end{aligned}
$$

The most expensive computation comes from $\mathcal{N}(\boldsymbol{y}|\boldsymbol{0}, \boldsymbol{K}^f + \sigma_{err}^2 I)$. Since this setting is separable $\boldsymbol{K}^f = \boldsymbol{B} \otimes \boldsymbol{K}$, methods exploiting Kronecker structure [Saatçi, 2012, Wilson, 2014] are discussed. Directly computing the likelihood costs $O(M^3 N^3)$, due to the computation of $(\boldsymbol{B} \otimes \boldsymbol{K} + \sigma_{err}^2 I)^{-1}$ and $\log\det(\boldsymbol{B} \otimes \boldsymbol{K} + \sigma_{err}^2 I)$.

Efficient computation approaches for the two terms are proposed through eigen-decomposition $\boldsymbol{B} = U_B D_B U_B^T$ and $\boldsymbol{K} = U_K D_K U_K^T$. Then we rewrite the two terms:

$$
\begin{aligned}
(\boldsymbol{B} \otimes \boldsymbol{K} + \sigma_{err}^2 I)^{-1} &= (U_B D_B U_B^T \otimes U_K D_K U_K^T + \sigma_{err}^2 I)^{-1} \\
&= (U D U^T + \sigma_{err}^2 I)^{-1} \\
&= U(D + \sigma_{err}^2 I)U^T
\end{aligned}
$$

where $U = U_B \otimes U_K$ is a unitary matrix and $D = D_B \otimes D_K$ is a diagonal matrix. And

$$
\begin{aligned}
\log\det(\boldsymbol{B} \otimes \boldsymbol{K} + \sigma_{err}^2 I) &= \log\det(U(D + \sigma_{err}^2 I)U^T) \\
&= \log\det(D + \sigma_{err}^2 I).
\end{aligned}
$$

Then applying Algorithm 14 in [Saatçi, 2012], the total computation cost is reduced to $O(\max(MN, M^3, N^3)) = O(\max(M^3, N^3))$.

In the general nonseparable setting, model parameters $\sigma_{err}^2, \boldsymbol{L}, \tilde{\ell}$ are of interest. Here $\boldsymbol{L}_{ij}(t)$ is a three dimensional tensor. The marginal posteriors of these parameters are

$$
\begin{aligned}
p(\sigma_{err}^2, \boldsymbol{L}, \tilde{\ell}, \boldsymbol{\sigma}|\boldsymbol{y}, \boldsymbol{t}) &= \int p(\boldsymbol{f}, \sigma_{err}^2, \boldsymbol{L}, \boldsymbol{g}, \tilde{\ell}, \boldsymbol{\sigma}|\boldsymbol{y}, \boldsymbol{t}) d\boldsymbol{g} d\boldsymbol{f} \\
&\propto \mathcal{N}(\vec{\boldsymbol{y}}|\boldsymbol{0}, \boldsymbol{K}^f + \sigma_{err}^2 I) \prod_{i \geq j} \mathcal{N}(\boldsymbol{L}_{ij}|\boldsymbol{\mu}_L, \boldsymbol{K}_L) \\
&\quad \mathcal{N}(\tilde{\ell}|\boldsymbol{\mu}_{\tilde{\ell}}, \boldsymbol{K}_{\tilde{\ell}}) \text{IG}(\sigma_{err}^2|a, b) \quad (5)
\end{aligned}
$$

4.2 Hamiltonian Monte Carlo

This section describes fully Bayesian inference via Hamiltonian Monte Carlo (HMC) [Brooks et al., 2011]. We implement HMC using automatic differentiation in pytorch [Baydin et al., 2018]. Implementing HMC on the marginal posterior (4), we obtain posterior samples for $\sigma_{err}^{2(s)}, \boldsymbol{L}^{(s)}, \tilde{\ell}^{(s)}, \tilde{\boldsymbol{\sigma}}^{(s)}$. The posterior samples of the correlation matrix $\boldsymbol{C}^{(s)}$, and the standard deviation $\overline{\boldsymbol{\sigma}}^{(s)}(t)$ at time t are derived as follows:

$$
\boldsymbol{C}^{(s)} = \boldsymbol{D}^{-1/2(s)} \boldsymbol{B}^{(s)} \boldsymbol{D}^{-1/2(s)}
$$

where $\boldsymbol{B}^{(s)} = \boldsymbol{L}^{(s)} \boldsymbol{L}^{T(s)}$ and $\boldsymbol{D}^{(s)} = \text{diag}(\boldsymbol{B}^{(s)})$. And

$$
\begin{aligned}
\overline{\boldsymbol{\sigma}}^{(s)}(t) &= \boldsymbol{\sigma}^{(s)}(t)\sqrt{\text{diag}\boldsymbol{B}^{(s)}} \\
&= \exp(\tilde{\boldsymbol{\sigma}}^{(s)}(t))\sqrt{\text{diag}\boldsymbol{B}^{(s)}}.
\end{aligned}
$$

Implementing HMC on the marginal posterior (5), we obtain posterior samples for $\sigma_{err}^{2(s)}, \boldsymbol{L}^{(s)}, \tilde{\boldsymbol{\ell}}^{(s)}$. The posterior samples of the correlation matrix $\boldsymbol{C}^{(s)}(t)$ and the standard deviation $\overline{\boldsymbol{\sigma}}^{(s)}(t)$ at time t are derived as follows:

$$\boldsymbol{C}^{(s)}(t) = \boldsymbol{D}^{-1/2(s)}(t)\boldsymbol{B}^{(s)}(t)\boldsymbol{D}^{-1/2(s)}(t)$$

where $\boldsymbol{B}^{(s)}(t) = \boldsymbol{L}^{(s)}(t)\boldsymbol{L}^{T(s)}(t)$ and $\boldsymbol{D}^{(s)}(t) = \text{diag}(\boldsymbol{B}^{(s)}(t))$. And

$$\overline{\boldsymbol{\sigma}}^{(s)}(t) = \sqrt{\text{diag}\boldsymbol{B}^{(s)}(t)}\,.$$

4.3 Model Prediction

For both separable and nonseparable models, given a new time stamp t^* with corresponding latent vector $\boldsymbol{f}^*$, the joint distribution of $(\boldsymbol{y}, \boldsymbol{f}^*)$ is

$$\begin{pmatrix} \vec{\boldsymbol{y}} \\ \boldsymbol{f}^* \end{pmatrix} \sim \mathcal{N}\left(\boldsymbol{0}, \begin{pmatrix} \boldsymbol{K}^f + \sigma_{err}^2 I & \boldsymbol{k}^f \\ \boldsymbol{k}^{fT} & \boldsymbol{K}^{f*} \end{pmatrix}\right)$$

where $\boldsymbol{K}^f = K^f(\boldsymbol{t}, \boldsymbol{t})$, $\boldsymbol{K}^* = K^f(t^*, t^*)$ and $\boldsymbol{k}^{fT} = \text{cov}(\boldsymbol{f}^*, \vec{\boldsymbol{f}} + \boldsymbol{\epsilon}) = \text{cov}(\boldsymbol{f}^*, \vec{\boldsymbol{f}})$. Since $\text{cov}(\text{vec}[\boldsymbol{f}_1, \ldots, \boldsymbol{f}_N], \boldsymbol{f}^*) = A^f L^f(t^*)^T$, where

$$A^f = \begin{pmatrix} k(t^*, t_1)L^f(t_1) \\ k(t^*, t_2)L^f(t_2) \\ \vdots \\ k(t^*, t_N)L^f(t_N) \end{pmatrix},$$

it follows that

$$\begin{aligned} \boldsymbol{k}^f &= \text{cov}(\vec{\boldsymbol{f}}, \boldsymbol{f}^*) \\ &= \text{cov}(P(\text{vec}[\boldsymbol{f}_1, \ldots, \boldsymbol{f}_N]), \boldsymbol{f}^*) \\ &= P(\text{cov}(\text{vec}(\boldsymbol{f}), \boldsymbol{f}^*)), \end{aligned}$$

where P is a permutation operator such that $P(\text{vec}([\boldsymbol{f}_1, \ldots, \boldsymbol{f}_N])) = \text{vec}([\boldsymbol{f}_1, \ldots, \boldsymbol{f}_N]^T)$. Therefore, the predictive distribution is a multivariate Normal distribution:

$$\boldsymbol{f}^*|\boldsymbol{y} \sim \mathcal{N}(\boldsymbol{\mu}^*, \Sigma^*) \tag{6}$$

where $\boldsymbol{\mu}^* = \boldsymbol{k}^{fT}(\boldsymbol{K}^f + \sigma_{err}^2 I)^{-1}\overline{\boldsymbol{y}}$ and $\Sigma^* = \boldsymbol{K}^{f*} - \boldsymbol{k}^{fT}(\boldsymbol{K}^f + \sigma_{err}^2 I)^{-1}\boldsymbol{k}^f$.

5 Experiments

We validate our proposed models on synthetic data and Kaiser Permanente's Electronic Health Records (EHR) data [Fohner et al., 2019]. We implemented all models/inference algorithms in Python using the open source Pytorch library.

5.1 Synthetic Data

We uniformly generate 200 timestamps on a unit interval $(0, 1)$. Next we generate a bi-variate Gaussian timeseries, in which the shared log length-scale process $l(t)$ is generated using a smooth Gaussian process prior $\text{GP}(0, \text{RBF}(\sigma = 4, d = 0.4))$ and the individual standard deviation process is generated using a log Gaussian process prior $\log\text{GP}(0, \text{RBF}(\sigma = 1, d = 0.1))$ and correlation process $r(t)$ between two covariates is generated via a deterministic function $r(t) = \cos(\pi t)$. Based on these three processes and assuming error variance $\sigma_{err}^2 = 1e-6$, we generate observed data using our proposed nonseparable kernel (3). We perform maximum a posteriori (MAP) inference using weak priors where, $\tilde{\ell}(t) \sim \text{GP}(0, \text{RBF}(\sigma = 5, d = 0.1))$, $L_{ij}(t) \sim \text{GP}(0, \text{RBF}(\sigma = 5, \ell = 0.1))$ and $\sigma_{err}^2 \sim \text{IG}(a = 1, b = 1)$. For MAP estimation, we initialize $\tilde{\ell}(t)$ by the empirical semivariogram and $L_i j(t)$ by the sample covariance matrix, using data in the window $[t - w, t + w]$, where w is the window size. We set the learning rate to 0.01. Simulation results show that the MAP estimates strongly depend on initialization, as well as on the window size w due to the nonconvexity of the objective function. The true latent processes and the inferred MAP processes and are shown in Figure 1. Because latent GPs involve unobserved functions generating the observed data, inference of the latent processes is especially difficult when data is limited. Nevertheless, Figure 1 shows that we are able to recover the correct trends for the latent processes reasonably well.

5.2 Kaiser Permanente Electronic Health Records Data

We demonstrate our proposed framework by modeling the time-varying vital signs such as, systolic blood pressure (BPSYS), diastolic blood pressure (BP-DIA), pulse pressure (PP), heart rate (HRTRT) and oxygen saturation (O2SAT) of patients admitted to the emergency department (ED) with confirmed or suspected infection. The Kaiser Permanente (KP) dataset is an anonymized EHR dataset where a patient's hospital stay is identified by an episode ID [Fohner et al., 2019, Seymour et al., 2016].

For our analysis we removed all measured vitals that had missing values to obtain a complete dataset (however as discussed in Section 4, missing data may be handled via marginalization). To better demonstrate the utility of our model we selected episodes that exhibited a high-degree of nonstationarity based on a unit root test [Maddala and Wu, 1999]. Episodes that had p-values greater than 0.1 were considered to have failed the stationarity test and were thus kept. Finally,
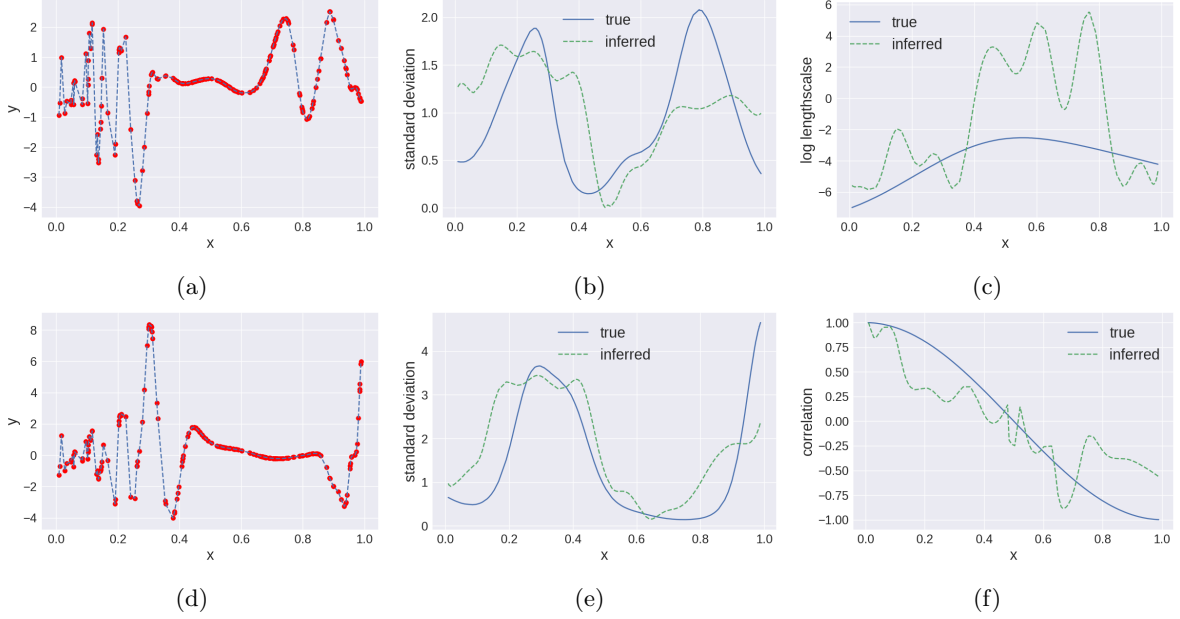
(a) (b) (c)

(d) (e) (f)

Figure 1: (1a): Generated data for dimension 1, (1d): Generated data for dimension 2, (1c): log length-scale processes, (1b) standard deviation process for dimension 1, (1e): standard deviation process for dimension 2, (1f): correlation process across dimensions 1 and 2.

for demonstration purposes we restricted analysis to episodes that had between 100 and 200 observations. Of these episodes we randomly selected a cohort of 205 episodes. For the KP dataset, we use priors identical to those used for inference on the synthetic dataset, except for L, where $L_{ij}(t) \sim \mathrm{GP}(0, \mathrm{RBF}(\sigma = 5, d = 0.2))$.

5.2.1 Prediction results

Prediction of future clinical observations for a patient is of significant interest for improved medical decision making, improved diagnosis, and clinical intervention. For model validation, we obtain the posterior distribution of our model parameters, based on all but the last five observations for a patient. We next predict the last five observations. In Table 1, we provide the mean and standard deviation of root mean square error (RMSE) as well as the log predictive density (LPD) for 205 episodes for stationary Gaussian processes model and non-stationary Gaussian processes model (separable and non-separable kernels). For visualization purposes, in Figure 2, we show the predictive performance for a single patient.

The simulation results clearly demonstrate that both the non-separable as well as the separable non-stationary Gaussian processes provide better predictive performance as well as uncertainty quantification over the stationary Gaussian process. However, as seen in Figure 2, some episodes require a non-stationary

model, and the non-separable non-stationary GP substantially outperforms the other models.

5.2.2 Inference of latent processes

Recent literature on sepsis [Fairchild et al., 2016, Cao et al., 2004] suggests that increased non-stationarity and/or increased correlation of vitals are often an early indicator of sepsis. In this section, we look at the inferred correlation processes across the vitals and plot them against the hourly LAPS2 scores. LAPS2 is a Kaiser Permanente-developed metric for acute disease burden that provides a measure of the acute severity of illness of a patient by evaluating a set of 15 common laboratory values, 5 vital signs, and neurologic status [Escobar et al., 2013]. Due to space constraints, we show results for three episodes. In Figure 3 the inferred time-varying correlation across heart rate (HRTRT) and oxygen saturation levels (O2SAT) seems to be in close agreement with the hourly LAPS2 score for episodes A and B. This might indicate that increased correlation across heart rate (HRTRT) and oxygen saturation levels (O2SAT) may potentially be an early indicator of increased risk for a patient. It is also interesting and intuitively pleasing to note that the uncertainty (as observed from the posterior samples) associated with the inferred correlation functions increases, for periods when a patient is monitored less frequently. In Figure 4 we also show the inferred correlation matrices across all vitals at three time-points during the inpatient stay for episode

	Stationary GP	Non-stationary GP (separable)	Non-stationary GP (non-separable)
RMSE	13.858 (6.375)	13.176 (6.159)	13.010 (6.193)
LPD	-3.916 (0.899)	-3.906 (0.989)	-3.863 (0.953)

Table 1: Predictive root mean square error (RMSE) and log predictive density (LPD) are summarized for stationary multivariate Gaussian processes (SMGP), nonstationary multivariate Gaussian processes (NMGP) and generalized nonstationary multivariate Gaussian processes (GNMGP) based on MAP inference. Mean and standard deviation (in brackets) for RMSE based on 205 episodes are provided.
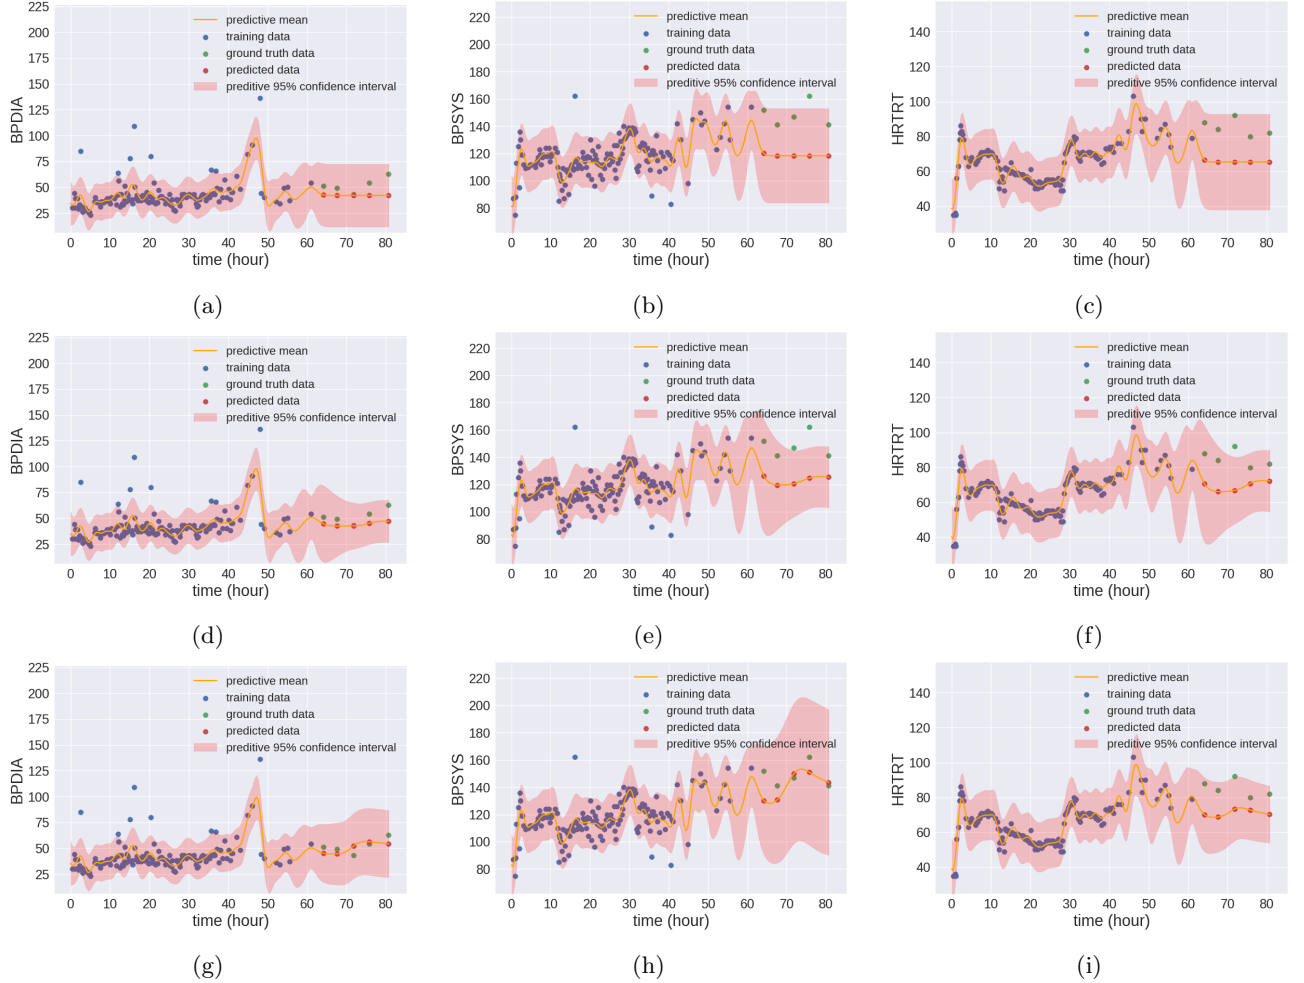

Figure 2: Prediction results for BPDIA, BPSYS and HTRT for (a,b,c) stationary model (RMSE: 20.056, LPD: -4.210), (d,e,f) separable non-stationary model (RMSE: 15.514, LPD: -3.957), and (g,h,i) non-separable non-stationary model (RMSE: 10.669, LPD: -3.804)

C. The correlation matrices show significant changes in correlation patterns across the patient's hospital stay, which is in direct contrast to an assumption almost invariably present in mathematical models for EHR in the existing literature, i.e., a stationary correlation matrix. We have performed similar analysis on the entire cohort of patients selected for this study. However, detailed analysis of such results are beyond the scope of this paper.

6 Conclusion

We have developed a novel nonstationary Gaussian process framework for modeling multiple correlated clinical variables. To the best of our knowledge, this is the first multivariate statistical model for EHR data which is both nonstationary and heteroscedastic. Both MAP and HMC inference procedures were developed for the proposed model. The model was then validated on both synthetic data and real EHR data from Kaiser
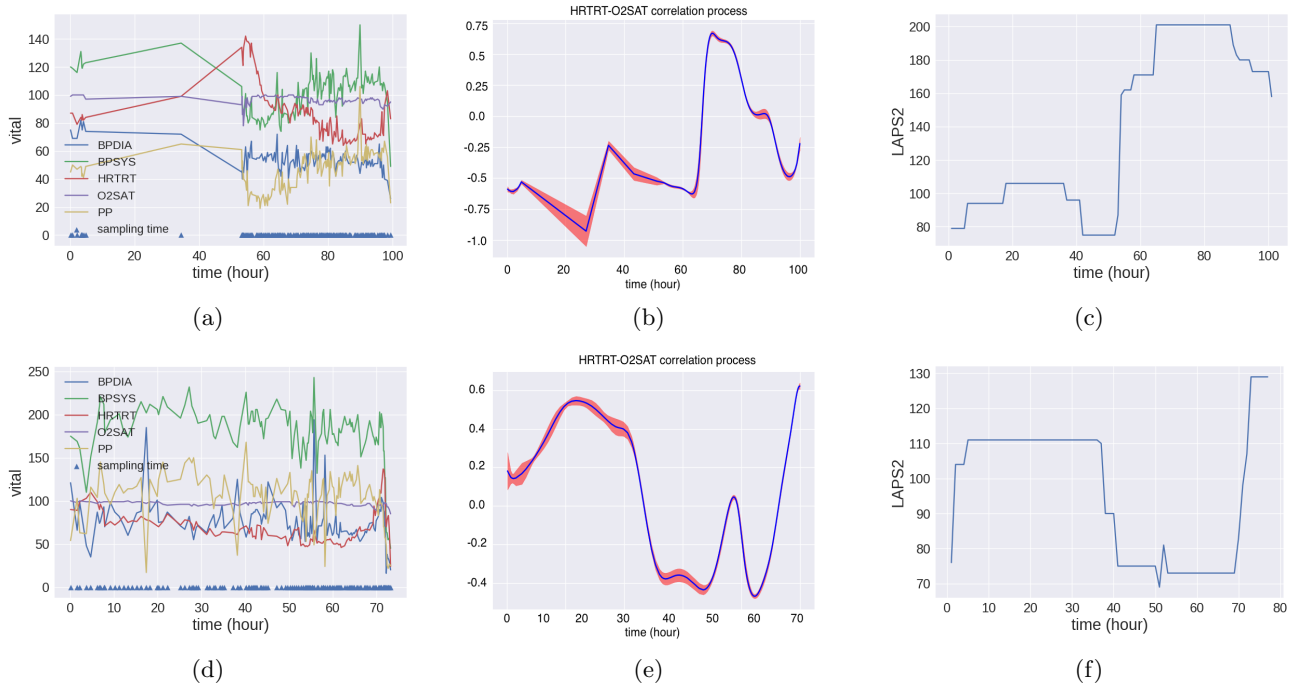
Figure 3: Observed vitals for episode A (a) and episode B (d). 100 HMC samples from the posterior of the HRTRT-O2SAT correlation process for episode A (b) and episode B (e). Hourly LAPS2 scores for episode A (c) and episode B (f).
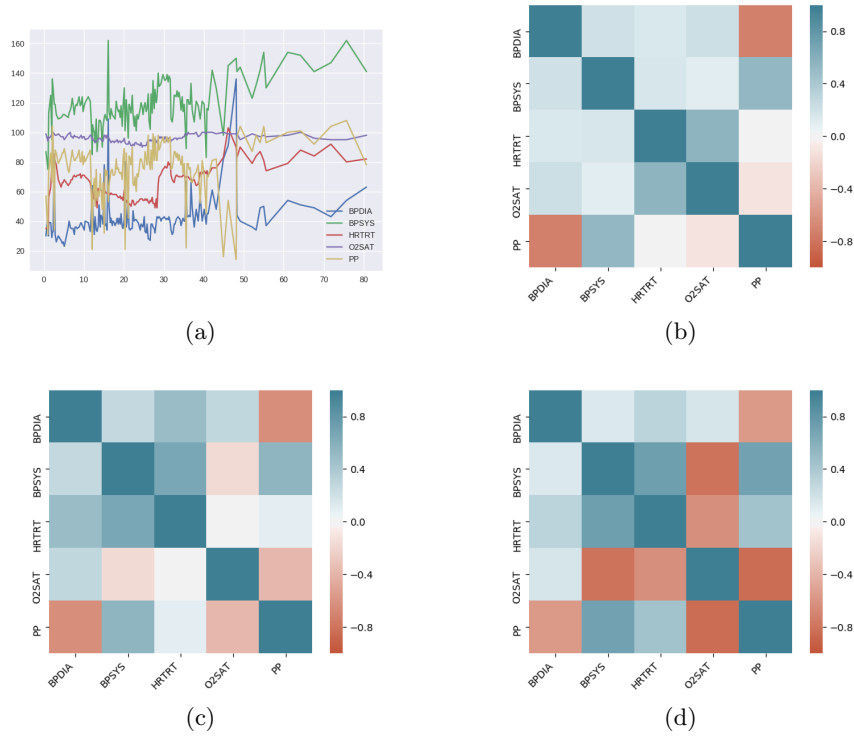

Figure 4: (a) Observed vitals for episode C and heat map of correlation matrix across all vitals at time (b) 20.8h, (c) 40.1h and (d) 71.8h.

Permanente. We demonstrated both improved prediction performance over stationary models as well as inferred latent time-varying correlation processes which are potentially indicative of patient risk. Future work includes a more detailed and systematic study of the inferred correlation processes and their relationship to a patient's risk profile.

References

[Alaa and van der Schaar, 2017] Alaa, A. M. and van der Schaar, M. (2017). Bayesian Inference of Individualized Treatment Effects using Multi-task Gaussian Processes. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R., editors, *Advances in Neural Information Processing Systems 30*, pages 3424–3432. Curran Associates, Inc.

[Alaa et al., 2018] Alaa, A. M., Yoon, J., Hu, S., and Schaar, M. v. d. (2018). Personalized Risk Scoring for Critical Care Prognosis Using Mixtures of Gaussian Processes. *IEEE Transactions on Biomedical Engineering*, 65(1):207–218.

[Baydin et al., 2018] Baydin, A. G., Pearlmutter, B. A., Radul, A. A., and Siskind, J. M. (2018). Automatic differentiation in machine learning: a survey. *Journal of machine learning research*, 18(153).

[Brooks et al., 2011] Brooks, S., Gelman, A., Jones, G., and Meng, X.-L. (2011). *Handbook of markov chain monte carlo*. CRC press.

[Cao et al., 2004] Cao, H., Lake, D. E., Griffin, M. P., and Moorman, J. R. (2004). Increased Nonstationarity of Neonatal Heart Rate Before the Clinical Diagnosis of Sepsis. *Annals of Biomedical Engineering*, 32(2):233–244.

[Cheng et al., 2017] Cheng, L.-F., Darnell, G., Dumitrascu, B., Chivers, C., Draugelis, M. E., Li, K., and Engelhardt, B. E. (2017). Sparse Multi-Output Gaussian Processes for Medical Time Series Prediction. *arXiv e-prints*, page arXiv:1703.09112.

[Cressie and Wikle, 2011] Cressie, N. and Wikle, C. K. (2011). *Statistics for Spatio-Temporal Data*. Wiley.

[Drichen et al., 2014] Drichen, R., Pimentel, M. A. F., Clifton, L., Schweikard, A., and Clifton, D. A. (2014). Multi-task Gaussian process models for biomedical applications. In *IEEE-EMBS International Conference on Biomedical and Health Informatics (BHI)*, pages 492–495.

[Escobar et al., 2013] Escobar, G. J., Gardner, M. N., Greene, J. D., Draper, D., and Kipnis, P. (2013). Risk-adjusting hospital mortality using a comprehensive electronic record in an integrated health care delivery system. *Medical care*, pages 446–453.

[Fairchild et al., 2016] Fairchild, K. D., Lake, D. E., Kattwinkel, J., Moorman, J. R., Bateman, D. A., Grieve, P. G., Isler, J. R., and Sahni, R. (2016). Vital signs and their cross-correlation in sepsis and NEC: a study of 1,065 very-low-birth-weight infants in two NICUs. *Pediatric Research*, 81:315.

[Fohner et al., 2019] Fohner, A. E., Greene, J. D., Lawson, B. L., Chen, J. H., Kipnis, P., Escobar, G. J., and Liu, V. X. (2019). Assessing clinical heterogeneity in sepsis through treatment patterns and machine learning. *Journal of the American Medical Informatics Association*.

[Futoma et al., 2017a] Futoma, J., Hariharan, S., and Heller, K. (2017a). Learning to Detect Sepsis with a Multitask Gaussian Process RNN Classifier. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML'17, pages 1174–1182, Sydney, Australia. JMLR.org. eventplace: Sydney, NSW, Australia.

[Futoma et al., 2017b] Futoma, J., Hariharan, S., Heller, K. A., Sendak, M., Brajer, N., Clement, M., Bedoya, A., and O'Brien, C. (2017b). An Improved Multi-Output Gaussian Process RNN with Real-Time Validation for Early Sepsis Detection. In *Proceedings of the Machine Learning for Health Care Conference, MLHC 2017, Boston, Massachusetts, USA, 18-19 August 2017*, pages 243–254.

[Gelfand et al., 2004] Gelfand, A. E., Schmidt, A. M., Banerjee, S., and Sirmans, C. (2004). Nonstationary multivariate process modeling through spatially varying coregionalization. *Test*, 13(2):263–312.

[Ghassemi et al., 2015] Ghassemi, M., Pimentel, M. A., Naumann, T., Brennan, T., Clifton, D. A., Szolovits, P., and Feng, M. (2015). A Multivariate Timeseries Modeling Approach to Severity of Illness Assessment and Forecasting in ICU with Sparse, Heterogeneous Clinical Data. *Proceedings of the ... AAAI Conference on Artificial Intelligence. AAAI Conference on Artificial Intelligence*, 2015:446–453.

[Heinonen et al., 2016] Heinonen, M., Mannerström, H., Rousu, J., Kaski, S., and Lähdesmäki, H. (2016). Non-stationary gaussian process regression with hamiltonian monte carlo. In *Artificial Intelligence and Statistics*, pages 732–740.

[Hripcsak et al., 2015] Hripcsak, G., Albers, D. J., and Perotte, A. (2015). Parameterizing time in electronic health record studies. *Journal of the Amer-*

ican Medical Informatics Association, 22(4):794–804.

[Jung and Shah, 2015] Jung, K. and Shah, N. H. (2015). Implications of non-stationarity on predictive modeling using EHRs. *Journal of Biomedical Informatics*, 58:168 – 174.

[Klompas and Rhee, 2016] Klompas, M. and Rhee, C. (2016). The cms sepsis mandate: right disease, wrong measure. *Annals of internal medicine*, 165(7):517–518.

[Lasko, 2014] Lasko, T. A. (2014). Efficient Inference of Gaussian-Process-Modulated Renewal Processes with Application to Medical Event Data. *Uncertainty in artificial intelligence : proceedings of the ... conference. Conference on Uncertainty in Artificial Intelligence*, 2014:469–476.

[Lasko, 2015] Lasko, T. A. (2015). Nonstationary Gaussian Process Regression for Evaluating Clinical Laboratory Test Sampling Strategies. *Proceedings of the ... AAAI Conference on Artificial Intelligence. AAAI Conference on Artificial Intelligence*, 2015:1777–1783.

[Lasko et al., 2013] Lasko, T. A., Denny, J. C., and Levy, M. A. (2013). Computational Phenotype Discovery Using Unsupervised Feature Learning over Noisy, Sparse, and Irregular Clinical Data. *PLOS ONE*, 8(6):1–13.

[Li and Marlin, 2016] Li, S. C.-X. and Marlin, B. (2016). A Scalable End-to-end Gaussian Process Adapter for Irregularly Sampled Time Series Classification. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, NIPS'16, pages 1812–1820, USA. Curran Associates Inc. event-place: Barcelona, Spain.

[Luttinen and Ilin, 2009] Luttinen, J. and Ilin, A. (2009). Variational gaussian-process factor analysis for modeling spatio-temporal data. In *Advances in neural information processing systems*, pages 1177–1185.

[Maddala and Wu, 1999] Maddala, G. S. and Wu, S. (1999). A comparative study of unit root tests with panel data and a new simple test. *Oxford Bulletin of Economics and statistics*, 61(S1):631–652.

[Paciorek and Schervish, 2004] Paciorek, C. J. and Schervish, M. J. (2004). Nonstationary covariance functions for gaussian process regression. In *Advances in neural information processing systems*, pages 273–280.

[Rasmussen and Williams, 2005a] Rasmussen, C. E. and Williams, C. K. I. (2005a). *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press.

[Rasmussen and Williams, 2005b] Rasmussen, C. E. and Williams, C. K. I. (2005b). *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning series)*. The MIT Press.

[Saatçi, 2012] Saatçi, Y. (2012). *Scalable inference for structured Gaussian process models*. PhD thesis, Citeseer.

[Schulam et al., 2015] Schulam, P., Wigley, F., and Saria, S. (2015). Clustering longitudinal clinical marker trajectories from electronic health data: Applications to phenotyping and endotype discovery.

[Seymour et al., 2016] Seymour, C. W., Liu, V. X., Iwashyna, T. J., Brunkhorst, F. M., Rea, T. D., Scherag, A., Rubenfeld, G., Kahn, J. M., Shankar-Hari, M., Singer, M., et al. (2016). Assessment of clinical criteria for sepsis: for the third international consensus definitions for sepsis and septic shock (sepsis-3). *Jama*, 315(8):762–774.

[Wilson, 2014] Wilson, A. G. (2014). *Covariance kernels for fast automatic pattern discovery and extrapolation with Gaussian processes*. PhD thesis, University of Cambridge.