
Patient-Specific Effects of Medication Using Latent Force Models with Gaussian Processes

Li-Fang Cheng*
Princeton University

Bianca Dumitrascu
Institute for Advanced Study

Michael Zhang
Princeton University

Corey Chivers
University of Pennsylvania Health System

Michael Draugelis
University of Pennsylvania Health System

Kai Li
Princeton University

Barbara E Engelhardt
Princeton University

Abstract

A multi-output Gaussian process (GP) is a flexible Bayesian nonparametric framework that has proven useful in jointly modeling the physiological states of patients in medical time series data. However, capturing the short-term effects of drugs and therapeutic interventions on patient physiological state remains challenging. We propose a novel approach that models the effect of interventions as a hybrid Gaussian process composed of a GP capturing patient baseline physiology convolved with a latent force model capturing effects of treatments on specific physiological features. The combination of a multi-output GP with a time-marked kernel GP leads to a well-characterized model of patients' physiological state across a hospital stay, including response to interventions. Our model leads to analytically tractable cross-covariance functions that allow for scalable inference. Our hierarchical model includes estimates of patient-specific effects but allows sharing of support across patients. Our approach achieves competitive predictive performance on challenging hospital data, where we recover patient-specific response to the administration of three common drugs: one antihypertensive drug and two anticoagulants.

1 Introduction

In the era of digital medicine, modern medical devices enable clinicians to accurately and frequently measure the physiological state of their patients. Heart rate, blood pressure, arterial oxygen saturation, and white blood cell counts are just a few of the physiological measurements used to monitor patient state and to predict outcomes such as hospital discharge rate, patient survival, or readmission (Ranganath et al., 2016; Rajkomar et al., 2018). To model the complex relationships governing patient state dynamics, a large body of work exploits the flexible properties of Gaussian processes (GPs). Both single output (Stegle et al., 2008; Lasko et al., 2013) and multi-output methods (Nemati et al., 2012; Ghassemi et al., 2015; Cheng et al., 2017; Futoma et al., 2017) use carefully designed kernels to improve prediction accuracy of patient states across sparse, noisy, and irregularly sampled time series horizons. These methods predict correlations between patient vital signs that give insight into imputing hard-to-sample state variables, detecting patient trends, and making decisions under uncertainty. Modeling patient state dynamics is made difficult by the administration of drugs and other therapeutic interventions, which directly impact state features for a limited time window and often return to a baseline. Accurate prediction of patient state following an intervention allows a clinician to consider multiple options for intervention and make informed treatment decisions.

Recent papers have attempted to model the complex effects of interventions on patient state (Xu et al., 2016; Schulam and Saria, 2017). These models do not incor-

porate a control systems understanding of the complex mechanisms governing the tightly controlled responses of physiological traits to interventions. This leads to inaccurate and non-generalizable predictions when insufficient training data is available, which is often the case when predictions include a patient-specific component to allow for personalized response estimation.

One motivation for modeling intervention effects on patient state with dynamical systems is to learn optimal treatment policies. For instance, prior work used reinforcement learning (RL) to derive a closed-loop anesthesia controller to regulate mean arterial pressure based on a dynamical model of a patient (Padmanabhan et al., 2015). Other applications, including finding multi-drug therapies for human immunodeficiency virus (HIV), have used dynamical systems models (Adams et al., 2004).

In this paper, we introduce a Bayesian nonparametric framework for estimating the dynamics of clinical traits in response to drug interventions through electronic health records (EHRs) from hospital patients. Specifically, we develop an approach to learn the patient-specific response of clinical traits to treatments from medical time series data. Our approach convolves a baseline multi-output Gaussian process (GP) with latent force models (LFMs) (Alvarez et al., 2009), which we model using GPs with kernels derived from differential equations representing dynamical systems.

To model the effects of interventions on clinical traits, we use latent force functions sampled from GPs with “causal” (time-marked) kernels (Cunningham et al., 2012). Here, the term “causal” refers to the physical constraint that an observation can only be affected by events that occurred earlier in time, not simultaneously or later in time; we show that this constraint is important for accurately capturing the effects of interventions. By estimating shared and patient-specific parameters that capture the effects of interventions on specific clinical traits, our model offers novel mechanistic insights and patient specificity, and also achieves competitive predictive accuracy when compared to state-of-the-art methods (Soleimani et al., 2017).

The contributions of this work are two-fold: First, we use time-marked kernels to model medical time-series data in order to capture the latent dynamics of intervention effects using a GP. Second, we incorporate these treatment effects in a model of patient state by convolving this GP latent force model with a multi-output GP that captures patient state absent interventions. Our approach is a necessary step towards achieving optimal control of patient health through personalized treatment policies (Särkkä et al., 2017). We show the predictive value of our approach on a large hospital

patient data set.

2 Background

2.1 Gaussian Processes for Time Series

In the medical time series setting, data are collected from n patients, indexed by i , across T_i time points indexed by t . Single output data typically correspond to time-varying covariates encoding physiological states such as blood pressure or heart rate. The measured pairs $\mathcal{D} = \{x_i, y_i\}_{i=1}^n$, where x_i corresponds to the time at which the covariate $y_i \in \mathbb{R}$ was recorded, are then modeled through an underlying latent function $f(\cdot)$ such that $y_i = f(x_i) + \epsilon_i$, where $\epsilon_i \sim \mathcal{N}(0, \sigma_i^2)$ is independent Gaussian noise. The goal is to estimate and to evaluate $f(\cdot)$ at future time points to enable prediction and early detection of physiologically abnormal states.

Gaussian processes (GPs) represent a versatile generative framework for modeling the distribution of an arbitrary real-valued function $f(\cdot)$. GPs are nonparametric stochastic processes specified by mean and covariance functions:

$$f(\mathbf{x}) \sim \mathcal{GP}(\mu(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')), \quad (1)$$

where $\mu(\mathbf{x})$ is the *mean function* $\mu(\mathbf{x}) = \mathbb{E}[f(\mathbf{x})]$ and $k(\mathbf{x}, \mathbf{x}')$ is the *covariance function* or *kernel*: $k(\mathbf{x}, \mathbf{x}') = \mathbb{E}[(f(\mathbf{x}) - \mu(\mathbf{x}))(f(\mathbf{x}') - \mu(\mathbf{x}'))]$. The mean function $\mu(\mathbf{x})$ is often assumed to be zero (Rasmussen and Williams, 2006). An immediate result is the multivariate Gaussian form of the joint $[f(x_1), f(x_2), \dots, f(x_n)] \sim \mathcal{N}(\mathbf{0}, K)$, where K is the $n \times n$ kernel matrix with entries $K_{i,j} = k(x_i, x_j)$.

Properties of the function $f(\mathbf{x})$ such as smoothness or periodicity are determined by the kernel function $k(\mathbf{x}, \mathbf{x}')$. One commonly used kernel is the squared exponential (SE) kernel

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

which is parameterized by a length scale ℓ and a scale factor σ . The functions generated by a GP with an SE kernel are smooth because the kernel function is infinitely differentiable. SE kernels capture *stationary* processes, as the covariance between two vectors depends on the difference in time (or other covariate) $\|\mathbf{x} - \mathbf{x}'\|$ but not on absolute time.

GPs have been studied extensively in the context of time series (Roberts et al., 2012), and are especially useful when the data are sparsely or irregularly sampled. In medical time series, the clinical measurements are recorded sporadically and sometimes sparsely across time. To better model multiple correlated measurements, multi-output GPs (MOGPs), which capture the

covariance structure between multiple measurements, have been adapted for use on medical data (Ghassemi et al., 2015; Futoma et al., 2017; Cheng et al., 2017). However, kernels used to capture temporal dependencies between sequential clinical measurements are mostly stationary. This is a limitation since, during a patient’s stay in the hospital, clinical events and interventions occur that affect physiological state for a period of time.

2.2 Latent Force Models for Patient Data

Physiological dynamics have been long studied by physiologists using systems of differential equations. For heart rate and blood pressure, for example, the cardiac conduction system is assumed to be a network of self-excitatory pacemakers leading to systems of nonlinear oscillators (Glass, 2001). Medical time series-based prediction of these covariates has relied on linear and nonlinear regression models, including but not limited to GPs; yet an explicit connection with control theory has been lacking. Recently, several methods were proposed that bridge the gap between stochastic control methods and nonparametric time series models (Gao et al., 2008; Alvarez et al., 2009). In particular, multi-output GPs may be used to represent latent force models (LFMs), where the covariance functions (kernels) are derived using ordinary differential equations (ODEs). In the LFM setup, we would like to model the system dynamics between a set of observed processes, $\{g_q(t)\}_{q=1}^Q$, and a set of unobserved *latent forces*, $f_m(t)$, assuming that they interact according to differential equations capturing those dynamics. For example, in a first-order LFM, the following equation holds:

$$\frac{dg_q(t)}{dt} + D_q g_q(t) = B_q + \sum_{m=1}^M S_{qm} f_m(t), \quad (3)$$

with B_q and D_q representing the base level value and underlying decay parameter of each output q , and S_{qm} the influence constants from each latent force f_m to each output $g_q(t)$. This formulation recovers the latent force functions $f_m(t)$ and input functions $g_q(t)$ from discrete observations. Solving the ODE yields

$$g_q(t) = \frac{B_q}{D_q} + \sum_{m=1}^M S_{qm} \exp(-D_q t) \int_0^t f_m(\tau) \exp(D_q \tau) d\tau. \quad (4)$$

In LFMs, the latent forces f_m are modeled independently as samples from their respective GPs. For certain classes of covariance functions, such as SE kernels, one can show that the outputs g_q are also GPs with analytically closed-form covariance functions, as well as cross-covariances between latent forces and the outputs (Alvarez et al., 2009).

While sampling the latent forces f_m from independent GPs is computationally convenient, important information can be lost. For example, when provided with historical patient data, we might want to know the physiological dynamics shared across patients from related covariate groups (i.e., same age, sex, BMI) or across patients receiving similar treatments. Limited numbers of observations also motivate a hierarchical approach to this problem, in order to share strength across patients.

2.3 Treatment Effect Estimation

In addition to modeling physiological state, accurate modeling of patient data requires incorporating treatment effects, to allow the control or stabilization of the physiological states of patients. Treatments are often drug interventions that can be characterized by drug name, administration type (e.g., oral, intravenous), and dosage. Estimating the effect of a treatment on a patient’s physiological state is paramount for comparing and choosing from potential treatments.

While GPs are commonly used to model medical time series data, several extensions have been proposed to estimate the effects of medical treatments. Counterfactual Gaussian processes (CGPs) (Schulam and Saria, 2017) use marked point processes (MPP) as an event model to account for dependencies between actions and observed physiological trajectories. In Xu et al. (2016), a class of parametric function is introduced to model the effects of dialysis for patients with acute kidney injury. The functions were designed to model different types of effects, including delay and decay. To explain heterogeneity across patients, a Dirichlet process was used for clustering patients. In Futoma et al. (2017), the treatment effects are modeled in the prior mean function of multi-output GPs, formulated as the sum of multiple exponential decay functions. These approaches require the response dynamics to conform to a specific functional form, whereas in practice these dynamics are often heterogeneous.

More recently, Soleimani et al. (2017) introduced treatment effects as the output of a linear time-invariant (LTI) system. The inputs are the observed drug administrations (e.g., type and dosage), and the effects are estimated based on a chosen form of a second-order filter. Patient-specific filter parameters were estimated and regularized using a global prior across patients.

To allow the treatment response to take on arbitrary functional forms with scalable effects sizes and directions, we use a hierarchical GP to model the latent forces. In particular, we replace the independent GPs with a causal-kernel GP whose hyperparameters are shared across patients to allow arbitrary functional

response and to share strength by capturing shared responses across patient groups.

3 Causal Convolutional GPs

Here, we propose a flexible framework to model medical time series data. Our method brings together ideas from GP latent force models (Alvarez et al., 2009; Särkkä et al., 2017) and causal GPs (Cunningham et al., 2012) to address a challenge in modeling medical time series data—the *systematic inclusion of multi-treatment effects on the dynamics of multiple physiological covariates*. We first introduce the notation in the context of time-marked medical data, and then introduce the model. We also discuss the details of implementation and inference methods.

3.1 Medical Time Series with Treatments

We denote the observed medical time series data $y_{i,j}^t$ from $i = \{1, 2, \dots, n\}$ patients characterized by $j = \{1, 2, \dots, J\}$ physiological dynamic covariates across irregularly sampled time points $t = \{1, 2, \dots, T_{i,j}\}$, as noisy samples from a Gaussian process with a patient- and covariate-specific mean function $\mu_{i,j}(t)$ and kernel $k_{i,j}^b(t, t')$:

$$\begin{aligned} y_{i,j}(t) &= f_{i,j}(t) + \epsilon_j, \quad \epsilon_j \sim \mathcal{N}(0, \sigma_{i,j}^2) \\ f_{i,j}(t) &\sim \mathcal{GP}(\mu_{i,j}(t), k_{i,j}^b(t, t')). \end{aligned} \quad (5)$$

Here, the kernel $k_{i,j}^b$ accounts for stationary temporal fluctuations of physiological signals, such as circadian rhythms. We choose this kernel as a sum of one SE kernel and one periodic kernel:

$$\begin{aligned} k_{i,j}^b &= k_{\text{SE}} + k_{\text{PER}} = \sigma_{1,i,j}^2 \exp\left[-\frac{(t-t')^2}{2\ell_{1,i,j}^2}\right] + \\ &\quad \sigma_{2,i,j}^2 \exp\left[-\frac{\sin^2(\pi||t-t'||/p_{i,j})}{2\ell_{2,i,j}^2}\right]. \end{aligned} \quad (6)$$

For each patient i , we index treatments as $m = 1, 2, \dots, M_i$, and denote the time of treatment as $t_1^i, t_2^i, \dots, t_{M_i}^i$. For each treatment, we use function $\tau_i : 1, 2, \dots, M_i \rightarrow \mathcal{T}$ to map the treatment index to a treatment set $\mathcal{T}$ representing the treatment type. Since the dosage-response curve of the same drug usually has a nonlinear curve that varies across dosages (Myers and Thiessen, 1980; Ghassemi et al., 2014), and the characteristics of absorption vary across different routes, we treat the same drug with different dosages or taken via different routes (e.g., oral or injection) as different treatments. Whenever clear from the context, we drop the patient index from these variables.

We assume there are different latent forces induced by each type of treatment. For a treatment m given

at time t_m , we model the latent force as a function of time $f_m(t; t_m)$ drawn from a Gaussian process. We also assume that each patient has a patient-specific latent force, and use a hierarchical model to share support for latent force models across patients.

3.2 Causal Treatment Dynamics

The dynamic behavior of a treatment’s response to the clinical covariates are represented in our setup as a latent force model. In particular, we model the mean function $\mu_{i,j}(t)$ of the clinical traits through the the first-order dynamical system LFM:

$$\frac{d\mu_{i,j}(t)}{dt} + D_{i,j}\mu_{i,j}(t) = B_{i,j} + \sum_{m=1}^{M_i} S_{i,j,m}f_m(t; t_m), \quad (7)$$

where the decay $D_{i,j}$, the baseline covariate output $B_{i,j}$, and treatment effect size $S_{i,j,m}$ are patient-specific parameters that control the dynamics of the treatment response. We assume these patient-specific parameters come from a population-wise empirical prior based on demographic data, such as age and weight. We assume the latent force function $f_m(t; t_m)$ of the same treatment is shared across patients, and is sampled from a Gaussian process with a causal covariance kernel:

$$f_m(t; t_m) \sim \mathcal{GP}(0, k_{f,f'}(t, t'; t_m)). \quad (8)$$

We use the term “causal” to refer to the constraint that medication effects may only act forward in time. To do this, we define the kernel function as,

$$k_{f,f'}(t, t'; t_m) = \exp\left\{-\frac{[h(t-t_m) - h(t'-t_m)]^2}{\ell_m^2}\right\}, \quad (9)$$

where $h(t) = t\mathcal{I}(t > 0)$ is the clipping function warping the input space and enforcing forward-time causality, while preserving the GP structure (Cunningham et al., 2012). Through $h(t)$, function $f_m(t; t_m)$ is constant before the current time t_m . In our model, we introduce an additional condition $f_m(t; t_m) = 0$ for $t < t_m$.

3.3 Kernel Convolution

The structure of the latent force model leads to a natural composition with the causal Gaussian process prior, leading to an analytic computation of output covariances and cross-covariances—in other words, a *convolution of kernels*. This fact allows for simple gradient descent-based inference. Closed-form kernels were derived following the integral in Eqn. 4. For instance, the cross-covariance between $\mu_{i,j}(t)$ and one

latent force $f_m(t)$ when $t, t' > t_m$ is computed as

$$\begin{aligned} & k_{\mu_{i,j},f_m}(t, t'; t_m) \\ &= S_{i,j,m} \exp(-D_{i,j}t) \exp\left[-\left(\frac{t' - t_m}{\ell_m}\right)^2\right] \\ &\times \frac{1}{D_{i,j}} [\exp(D_{i,j}t_m) - 1] \\ &+ \frac{\sqrt{\pi}\ell}{2} S_{i,j,m} \exp[-D_{i,j}(t - t')] \exp(\nu_{i,j,m}^2) \\ &\times \left[\operatorname{erf}\left(\frac{t - t'}{\ell_m} - \nu_{i,j,m}\right) + \operatorname{erf}\left(\frac{t' - t_m}{\ell_m} + \nu_{i,j,m}\right) \right], \end{aligned} \quad (10)$$

where $\nu_{i,j,m} = \frac{\ell_m D_{i,j}}{2}$. Details of the computation of the closed-form kernels corresponding to the cases $t > t_m > t'$, $t_m > t, t'$, and $t' > t_m > t$ are in the Supplementary Material.

3.4 Hyperparameter Learning for Medication Effects Model

To learn the hyperparameters for each patient, we optimize the marginal likelihood of the GP model with respect to the vector of observations across covariates $\mathbf{x}_i$ and $\mathbf{y}_i$.

$$\begin{aligned} \log p(\mathbf{y}_i | \mathbf{x}_i, \boldsymbol{\theta}) &= -\frac{1}{2}(\mathbf{y}_i - \boldsymbol{\mu}_i)^\top (K_{i|\boldsymbol{\theta}} + \epsilon I)^{-1} (\mathbf{y}_i - \boldsymbol{\mu}_i) \\ &\quad - \frac{1}{2} \log |K_{i|\boldsymbol{\theta}} + \epsilon I| - \left(\frac{\sum_{j=1}^J T_{i,j}}{2} \right) \log(2\pi), \end{aligned} \quad (11)$$

where $\boldsymbol{\mu}_i = \mu_{i,j}(\mathbf{x}_i)$. For each patient, we learn a set of hyperparameter $\boldsymbol{\theta}$, which consists of covariate-specific hyperparameters for baseline kernel, $\{\sigma_{1,i,j}, \ell_{1,i,j}, \sigma_{2,i,j}, \ell_{2,i,j}, p_{i,j}\}$, and hyperparameters shared across covariates through causal LFM kernel, $\{B_{i,j}, D_{i,j}, S_{i,j,m}, \ell_{i,j,m}\}$. We assume that the patient- and treatment-specific hyperparameters are shared across treatments of the same type, while the treatment effect size parameter differs for different dosages or modes of administration. Our implementation is based on GPy (GPy, since 2012), and we optimize the hyperparameters using scale conjugate gradient methods. We derived the gradients using the SymPy package (Meurer et al., 2017).

4 Experiments

In this section, we show the effectiveness of our method by modeling multiple treatment effects on electronic healthcare record (EHR) data from the Hospital of University of Pennsylvania (HUP). We briefly describe the data and preprocessing steps, and then we discuss results from our method fitted to patient subsets motivated by two clinical applications. We show empirical results of our method using the metrics of prediction

accuracy. We compare results with baseline methods with GPs using basic kernels and mean functions from related work.

4.1 Inpatient Hospital Data

We evaluated our method using clinical data collected at HUP. The data set consists of 139k patients with access to demographic details (e.g., age, weight, gender), as well as 139 clinical measurements consisting of vital signs and lab tests, and administrated medications during the patients' stay in the hospital. We normalize each clinical trait by subtracting the empirical mean for each patient from each measurement. We tested our method on two challenging applications—modeling the effects of antihypertensive agents and anticoagulants. We chose to focus on the patients with a primary diagnosis of myocardial infarction (MI; i.e., heart attack) in our data set, resulting in total 1,716 adult patients, as they usually received both types of treatments.

We first modeled the effects of the most frequently administered antihypertensive drug in our data set, metoprolol, on heart rate (HR) and systolic blood pressure (SBP). Metoprolols are beta-blockers that are used to treat high blood pressure or angina due to heart disease. We filtered MI patients to include patients with at least five observations in both heart rate or blood pressure, reducing the number of patients to 594. We removed patients that were administrated metoprolol jointly with other antihypertensive agents, resulting in 233 patients. Finally, we included treatments of metoprolol that were administered at least 20 times across all patients, resulting in 181 patients with six type of treatments, including four different dosages of metoprolol tartrate, and one dosage for metoprolol succinate ER and metoprolol injection, for a total of 310 treatment administrations.

Second, we modeled the effects of two different types of anticoagulants: heparin and warfarin. We filtered the MI patients to include those with at least five observations on two lab test results that reflect a patient's ability to form blood clots: partial thromboplastin time (PTT) and international normalized ratio (INR), resulting in 581 patients. Among all treatments, we considered the top three most frequently administered, and we include patients that received at least one of them. The filtered data includes 404 patients with a total of 592 treatment administrations (Table 1).

4.2 Evaluation Metrics

We evaluated our model by comparing predictive performance with three state-of-the-art GP models: (i) univariate GPs with a constant mean function and the baseline (squared exponential and periodic) ker-

Antihypertensive ($N = 181$)	Count
Heart Rate (HR)	9,798
Systolic Blood Pressure (SBP)	7,804
Metoprolol Tartrate (6.25 mg)	42
Metoprolol Tartrate (12.5 mg)	92
Metoprolol Tartrate (25 mg)	93
Metoprolol Tartrate (50 mg)	40
Metoprolol Succinate ER (25 mg)	22
Metoprolol Injection (5 mg)	21
Anticoagulants ($N = 404$)	Count
Partial Thromboplastin Time (PTT)	4,911
International Normalized Ratio (INR)	4,348
Heparin Injection (5000 units)	246
Heparin Infusion (25000 units)	319
Warfarin (5mg)	27

Table 1: **Data statistics of the two drug types used in the experiments.** Total number of observations for the targeted vital signs and lab results, and the count of targeted treatments.

nel $k_{i,j}^b(t, t')$ (Eqn. 6; SE+PER), (ii) univariate GPs with an exponential decay mean function and the Ornstein-Uhlenbeck (OU) kernel (Futoma et al., 2017) (OU+EXP), and (iii) Matérn- $3/2$ kernel with a second-order LTI filter for effect modeling (Soleimani et al., 2017) (MAT32+LTI). For the comparative method (ii) and (iii), we added a constant component in the mean function in the proposed setup to account for patient-specific baseline values of each covariate. For all methods, we used the first 70% of observations for each patient for training, and the remaining 30% for testing. We computed the mean absolute error (MAE) of predictions on test data to evaluate model performance.

4.3 Clinical Impact

When comparing the MAE on test data for the two experiments (antihypertensives and anticoagulants; Table 2), our method performed competitively across experiments on the predictive tasks for the covariates responding to antihypertensives. Our model performs better than related methods on the task of predicting traits responding to anticoagulants, in particular the blood clot formation trait PTT.

While predictive performance is similar across related methods, our method shows important advantages in model flexibility. The proposed model uses GP priors for the medication effect function, which allow a flexible functional form unlike the parametric functions used in previous work (Futoma et al., 2017; Soleimani

et al., 2017). We demonstrate this through two case studies: First, we study predictive trajectories and the inferred effects on blood pressure and heart rate of a treatment on one patient (Fig. 1). The patient received 50 mg of metoprolol tartrate. We compare the predictive trajectory with uncertainty from a related method (Mat32+LTI) (Fig. 1 a–b) with results from our method (Fig. 1 c–d). Our method estimates an explicit treatment-induced latent force with strong effects on both heart rate and systolic blood pressure (Fig. 1 e) that matches the time frame prescribed by clinical medicine (Broegden et al., 1977; Kjekshus, 1986; Kezerashvili et al., 2012). Indeed, the direction of estimated effects from the related method on SBP (Fig. 1, a and b) is the opposite of the clinical usage; this error may be due to delayed effect of the drug on SBP. While for this patient the MAE of the related method is slightly lower than our method (4.68 vs. 5.37 in SBP; 7.90 vs. 8.09 in HR), results from our method are closer to clinical ground truth with substantially smaller uncertainty. Furthermore, with the GP model of latent force, our estimated effect is more robust than a standard latent force model when presented with additional uncertainty and noise in the data, such as delayed effects from the time of treatment administration, which is important in the clinical setting.

Next, we find similar robust and clinically interpretable behavior are achieved for a patient receiving multiple heparin infusions (25,000 units) and heparin injections (5,000 units; Fig. 2). For both our method and the related method (MAT32+LTI), the longer-term effects on PTT is estimated (Fig. 2, a and c). Both methods estimated negative effects for the heparin injections (Fig. 2f), while in general heparin is assumed not to affect INR (Katzung and Trevor, 2015).

5 Discussion

We developed a framework using latent force models (LFMs) to capture treatment effects on patients’ physiological state estimated using medical time series data. By modeling treatment effects as latent forces convolved with a multi-output GP, we create a flexible framework that bridges the gap between the smooth, stationary dynamics of patient state and a mechanistic understanding of the forces that impact clinical traits. A GP model of latent forces provides a flexible probabilistic framework with convenient inference properties; we enforce appropriate effect dynamics using a causal kernel. Two key contributions—the latent force and multi-output GP model convolution and the causal kernel in the latent force GP—lead to a computationally tractable solution with low variance and clinically-relevant interpretation of personalized treatment effects. Further improvements in speed may

Covariate	Antihypertensive Agents			
	SE+PER	OU+EXP	MAT32+LTI	Proposed
SBP	10.688 ± 0.242	11.547 ± 0.292	10.502 ± 0.240	10.882 ± 0.242
HR	7.694 ± 0.173	7.780 ± 0.186	7.791 ± 0.173	7.851 ± 0.179
Covariate	Anticoagulants			
	SE+PER	OU+EXP	MAT32+LTI	Proposed
PTT	12.646 ± 0.379	12.431 ± 0.512	12.436 ± 0.376	12.248 ± 0.368
INR	0.204 ± 0.009	0.339 ± 0.013	0.596 ± 0.011	0.180 ± 0.009

Table 2: **Prediction results on test data.** MAE ($\pm$ standard error) was computed for results from our method and three comparison methods using univariate GPs with different mean functions and kernels. Our method performs competitively across covariates for the antihypertensives and better than the comparison methods for the anticoagulants.




Figure 1: **Prediction of systolic blood pressure (SBP) and heart rate (HR) for one patient under metoprolol tartrate treatments.** (a) and (b): results from a related method MAT32+LTI. (c) and (d): results from the proposed method. (e) The learned latent force for the first metoprolol tartrate (50 mg) treatment and the effects on SBP and HR; the recovered effects on SBP and HR are in their original units. Our method achieves higher confidence (lower predictive variance) and greater consistency with known clinical effects than the related model.

be found by adapting a recent kernel approximation (Guarnizo and Lopez, 2018).

There are several directions to improve our method for clinical treatment effect estimation. Our framework assumes that the effect of each treatment is independent

of any other, and interactions between treatments are not modeled. These interactions could be modeled by modifying the kernel for the latent force component. In addition, the method assumes the decay parameters ($D_{i,j}$) of the treatment effect for a single patient are treatment-independent and constant throughout the




Figure 2: **Prediction of partial thromboplastin time (PTT) and international normalized ratio (INR) for one patient under heparin treatments.** (a) and (b): results from the related method MAT32+LTI. (c) and (d): results from the proposed method. (e) and (f): the estimated latent force for one heparin infusion (25000 units) and one heparin injection (5000 units), and their effects on PTT and INR; the recovered effects on PTT and INR are in their *normalized* units.

patient’s hospital stay. As the decay parameters reflect the physiology of drug absorption, which may change as a function of patient state, we might model this parameter as a stochastic process itself. For future research, the latent force model encourages an optimal control perspective: estimating treatment effect sizes of each patient for each clinical covariate, coupled with accurate patient state prediction, can provide a path toward treatment prioritization and decision-making in clinical interventions.

Acknowledgements

BD completed part of this work while visiting The Statistical and Applied Mathematical Sciences Institute in Durham, NC, under the kind support of the NSF grant DMS-1638521.

References

B. M. Adams, H. T. Banks, H.-D. Kwon, and H. T. Tran. Dynamic multidrug therapies for HIV: Op-

timal and STI control approaches. *Mathematical Biosciences and Engineering*, 1(2):223–241, 2004.

M. Alvarez, D. Luengo, and N. Lawrence. Latent force models. In *Artificial Intelligence and Statistics*, pages 9–16, 2009.

R. Brogden, R. Heel, T. Speight, and G. Avery. Metoprolol: a review of its pharmacological properties and therapeutic efficacy in hypertension and angina pectoris. *Drugs*, 14(5):321–348, 1977.

L.-F. Cheng, G. Darnell, B. Dumitrascu, C. Chivers, M. E. Draugelis, K. Li, and B. E. Engelhardt. Sparse multi-output Gaussian processes for medical time series prediction. *ArXiv e-prints*, Mar. 2017.

J. Cunningham, Z. Ghahramani, and C. Rasmussen. Gaussian processes for time-marked time-series data. In N. D. Lawrence and M. Girolami, editors, *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics*, volume 22 of *Proceedings of Machine Learning Research*, pages

- 255–263, La Palma, Canary Islands, 21–23 Apr 2012. PMLR.
- J. Futoma, S. Hariharan, M. Sendak, N. Brajer, M. Clement, A. Bedoya, C. O’Brien, and K. Heller. An improved multi-output Gaussian process RNN with real-time validation for early sepsis detection. *arXiv preprint arXiv:1708.05894*, 2017.
- P. Gao, A. Honkela, M. Rattray, and N. D. Lawrence. Gaussian process modelling of latent chemical species: applications to inferring transcription factor activities. *Bioinformatics*, 24(16):i70–i75, 2008.
- M. Ghassemi, M. A. F. Pimentel, T. Naumann, T. Brennan, D. A. Clifton, P. Szolovits, and M. Feng. A multivariate timeseries modeling approach to severity of illness assessment and forecasting in ICU with sparse, heterogeneous clinical data. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 446–453, 2015.
- M. M. Ghassemi, S. E. Richter, I. M. Eche, T. W. Chen, J. Danziger, and L. A. Celi. A data-driven approach to optimized medication dosing: a focus on heparin. *Intensive care medicine*, 40(9):1332–1339, 2014.
- L. Glass. Synchronization and rhythmic processes in physiology. *Nature*, 410(6825):277, 2001.
- GPY. GPY: A Gaussian process framework in Python. <http://github.com/SheffieldML/GPY>, since 2012.
- C. Guarnizo and A. Lopez. Fast kernel approximations for latent force models and convolved multiple-output Gaussian processes. In *Proceedings of the Thirty-Fourth Conference in Uncertainty and Artificial Intelligence (2018)*. Sheffield, 2018.
- B. G. Katzung and A. J. Trevor. *Basic & Clinical Pharmacology*. McGraw-Hill New York, NY, 2015.
- A. Kezerashvili, K. P. Marzo, and J. D. de Leon. Beta blocker use after acute myocardial infarction in the patient with normal systolic function: When is it “ok” to discontinue? In *Current cardiology reviews*, 2012.
- J. K. Kjekshus. Importance of heart rate in determining beta-blocker efficacy in acute and long-term acute myocardial infarction intervention trials. *The American journal of cardiology*, 57(12):F43–F49, 1986.
- T. A. Lasko, J. C. Denny, and M. A. Levy. Computational phenotype discovery using unsupervised feature learning over noisy, sparse, and irregular clinical data. *PloS One*, 8(6):1–13, 2013.
- A. Meurer, C. P. Smith, M. Paprocki, O. Čertík, S. B. Kirpichev, M. Rocklin, A. Kumar, S. Ivanov, J. K. Moore, S. Singh, T. Rathnayake, S. Vig, B. E. Granger, R. P. Muller, F. Bonazzi, H. Gupta, S. Vats, F. Johansson, F. Pedregosa, M. J. Curry, A. R. Terrel, v. Roučka, A. Saboo, I. Fernando, S. Kulal, R. Cimrman, and A. Scopatz. Sympy: symbolic computing in Python. *PeerJ Computer Science*, 3:e103, Jan. 2017. ISSN 2376-5992. doi: 10.7717/peerj-cs.103. URL <https://doi.org/10.7717/peerj-cs.103>.
- M. G. Myers and J. J. Thiessen. Metoprolol kinetics and dose response in hypertensive patients. *Clinical Pharmacology & Therapeutics*, 27(6):756–762, 1980.
- S. Nemati, L.-W. H. Lehman, R. P. Adams, and A. Malhotra. Discovering shared cardiovascular dynamics within a patient cohort. In *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 6526–6529, 2012.
- R. Padmanabhan, N. Meskin, and W. M. Haddad. Closed-loop control of anesthesia and mean arterial pressure using reinforcement learning. *Biomedical Signal Processing and Control*, 22:54–64, 2015.
- A. Rajkomar, E. Oren, K. Chen, A. M. Dai, N. Hajaj, M. Hardt, P. J. Liu, X. Liu, J. Marcus, M. Sun, et al. Scalable and accurate deep learning with electronic health records. *npj Digital Medicine*, 1(1):18, 2018.
- R. Ranganath, A. Perotte, N. Elhadad, and D. Blei. Deep survival analysis. *arXiv preprint arXiv:1608.02158*, 2016.
- C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- S. Roberts, M. Osborne, M. Ebdon, S. Reece, N. Gibson, and S. Aigrain. Gaussian processes for time-series modelling. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 371(1984), 2012.
- S. Särkkä, M. A. Álvarez, and N. D. Lawrence. Gaussian process latent force models for learning and stochastic control of physical systems. *arXiv preprint arXiv:1709.05409*, 2017.
- P. Schulam and S. Saria. Reliable decision support using counterfactual models. In *Advances in Neural Information Processing Systems*, pages 1696–1706, 2017.
- H. Soleimani, A. Subbaswamy, and S. Saria. Treatment-response models for counterfactual reasoning with continuous-time, continuous-valued interventions. In *33rd Conference on Uncertainty in Artificial Intelligence, UAI 2017*. AUAI Press Corvallis, 2017.
- O. Stegle, S. V. Fallert, D. J. C. MacKay, and S. Brage. Gaussian process robust regression for noisy heart rate data. *IEEE Transactions on Biomedical Engineering*, 55(9):2143–2151, 2008.
- Y. Xu, Y. Xu, and S. Saria. A Bayesian nonparametric approach for estimating individualized treatment-

response curves. In *Machine Learning for Healthcare Conference*, pages 282–300, 2016.