

Personalized Dynamic Treatment Regimes in Continuous Time: A Bayesian Approach for Optimizing Clinical Decisions with Timing*

William Hua[†], Hongyuan Mei[‡], Sarah Zohar[§], Magali Giral[¶], and Yanxun Xu^{||}

Abstract. Accurate models of clinical actions and their impacts on disease progression are critical for estimating personalized optimal dynamic treatment regimes (DTRs) in medical/health research, especially in managing chronic conditions. Traditional statistical methods for DTRs usually focus on estimating the optimal treatment or dosage at each given medical intervention, but overlook the important question of “when this intervention should happen.” We fill this gap by developing a two-step Bayesian approach to optimize clinical decisions with timing. In the first step, we build a generative model for a sequence of medical interventions—which are discrete events in continuous time—with a marked temporal point process (MTPP) where the mark is the assigned treatment or dosage. Then this clinical action model is embedded into a Bayesian joint framework where the other components model clinical observations including longitudinal medical measurements and time-to-event data conditional on treatment histories. In the second step, we propose a policy gradient method to learn the personalized optimal clinical decision that maximizes the patient survival by interacting the MTPP with the model on clinical observations while accounting for uncertainties in clinical observations learned from the posterior inference of the Bayesian joint model in the first step. A signature application of the proposed approach is to schedule follow-up visitations and assign a dosage at each visitation for patients after kidney transplantation. We evaluate our approach with comparison to alternative methods on both simulated and real-world datasets. In our experiments, the personalized decisions made by the proposed method are clinically useful: they are interpretable and successfully help improve patient survival.

Keywords: Bayesian joint model, dynamic treatment regimes, electronic health records, marked temporal point process, policy gradient.

*This work was supported by National Science Foundation DMS1918854 and 1940107.

[†]Department of Applied Mathematics and Statistics, Johns Hopkins University

[‡]Department of Computer Science, Johns Hopkins University

[§]INSERM, U1138, team 22, Centre de Recherche des Cordeliers, Université Paris 5, Université Paris 6

[¶]Centre de Recherche en Transplantation et Immunologie (CRTI), UMR1064, INSERM, Université de Nantes

^{||}Department of Applied Mathematics and Statistics, Johns Hopkins University, yanxun.xu@jhu.edu

1 Introduction

In biomedical applications involving long-term personalized care of patients with chronic health conditions (e.g., diabetes, human immunodeficiency virus infections, and chronic kidney diseases), treatments often include a sequence of decision making and must be adaptive to the uniquely evolving disease progression of each patient. Such scenarios are called dynamic treatment regimes (DTRs). Patients with chronic diseases are usually required to follow up with their physicians from time to time and their clinical data are recorded longitudinally. Based on these clinical observations, physicians make clinical decisions such as scheduling follow-up visitations and prescribing the right dosages to optimize patient outcomes given a patient's individual characteristics and treatment history at each clinic visitation. Estimating treatment effects and optimizing sequential treatment assignments from observational data have been extensively studied in both statistics and machine learning, such as the G-computation formula (Robins, 1986), inverse probability of treatment weighting (Orellana et al., 2010), doubly robust methods (Zhao et al., 2015), and reinforcement learning (Zhao et al., 2011; Clifton and Laber, 2020). This paper develops a two-step Bayesian approach to optimize clinical decisions with timing. In the first step, we develop a Bayesian joint model consisting of a generative probabilistic submodel for clinical decisions with timing and a submodel for clinical observations (e.g., longitudinal clinical measurements and time-to-event data): these two submodels share certain structures and parameters in order to capture the mutual influence between the clinical observations and decisions. The posterior inference of the proposed Bayesian joint model can learn the parameters in the clinical decision submodel as the estimates of how physicians treat patients in the observed data and uncertainties in clinical observations. In the second step, we propose an optimization method that allows the decision model, by interacting with the other parts of the joint model, to learn to make the personalized optimal clinical decision at the right time while accounting for uncertainties in clinical observations. Such a joint model and the proposed optimization method will be useful in many biomedical applications. We elaborate on one signature application in Section 1.1, explain why existing methods won't work well on it in Section 1.2, and then give an overview of our method and its technical novelty in Section 1.3.

1.1 A signature application

A signature medical application of the proposed method would be the kidney transplantation, the most common type of organ transplantation and the primary therapy for patients with end-stage kidney diseases (Arshad et al., 2019). Compared to dialysis, kidney transplantation improves patients' long-term survival and quality of life but with a lower healthcare cost (Jarl et al., 2018). Despite significant advances, a number of complications after surgery still represent important causes of morbidity and mortality for kidney transplant recipients, such as infection, stroke, and graft failure (Lamb and Lodhi, 2011; Bicalho et al., 2019). To prevent graft rejection, patients are usually hospitalized for a few days initially to monitor signs of complications, then required to have frequent checkups at an outpatient center after being released. At each visitation, they are administered immunosuppressive drugs, such as tacrolimus, to keep their immune systems from attacking and rejecting the new kidney (Kasiske et al., 2010). One

crucial medical decision is to schedule patients' post-transplantation follow-up visitations. While follow-up visitation frequency varies from 0-12 months (Israni et al., 2014), patients with stable kidney function usually have less frequent follow-ups compared to non-stable patients. Another medical decision is to determine the right dosage of tacrolimus at each follow-up visitation since a dosage that is either too high or too low may cause serious adverse events. Higher tacrolimus levels have been reported to associate with adverse effects such as neurotoxicity, nephrotoxicity, and cancers (Naesens et al., 2009); while lower tacrolimus levels are associated with an increased likelihood of graft rejection (Staatz et al., 2001). Therefore, optimizing personalized follow-up schedules and prescribing the right dosage of tacrolimus tailored to each patient at each visitation (i.e., precision medicine) are critical and can have a significant impact on patients' survival.

Large-scale kidney transplantation databases, such as French computerized and validated data in kidney transplantation (DIVAT), provide us both opportunities and challenges to determine personalized optimal follow-up schedules and tacrolimus dosages. DIVAT is a database storing medical records for kidney transplantation in several French hospitals (e.g., Nantes, Paris Necker). Data are collected from the date of transplantation until the graft failure, defined as either returning to dialysis or death with a functioning graft. At each scheduled follow-up visitation, patients' creatinine levels, an important biomarker for measuring kidney function, are collected longitudinally to determine the next follow-up time and assign dosages by physicians. For example, Figure 1 presents one randomly selected patient's longitudinal creatinine levels and tacrolimus dosages versus his/her follow-up visitations from DIVAT. In the first several visitations after kidney transplantation, this patient's creatinine levels were high, indicating the kidney was not functioning well; therefore, the physician scheduled a high frequency of follow ups and prescribed high dosages of tacrolimus. As time went by, this patient's kidney function became stable indicated by slowly decreasing creatinine levels; then the prescribed tacrolimus dosages also slowly decreased accompanied with a decreasing frequency of visitations. For patients with kidney transplantation, a major clinical outcome of interest is the graft survival time, defined as the time between the transplantation and the first graft failure. Follow-up schedules and tacrolimus dosages should be made for the sake of maximizing patients' graft survival time.

1.2 Why not use existing methods?

Although many statistical and machine learning DTR methods have been developed to optimize sequential clinical decisions (Murphy et al., 2003; Chakraborty, 2013; Laber et al., 2014; Xu et al., 2016; Luckett et al., 2019), they don't model, and thus can't optimize, the timing of clinical decisions. Most DTR methods regard treatment schedules as known *a priori* and only learn to assign optimal treatments at pre-defined schedules. For example, Xu et al. (2016) developed a Bayesian nonparametric approach building upon a dependent Dirichlet process and a Gaussian process to determine the optimal treatment regimen containing a front-line chemotherapy and a salvage treatment for acute myelogenous leukemia patients. However, the timing of the salvage treatment was pre-defined as the time when patients became resistant to the front-line chemotherapy or achieved complete remission first then relapsed. Shardell and Ferrucci (2018)




Figure 1: Example data for one patient’s creatinine and tacrolimus levels on a log scale over time. The points represent actual visitations.

proposed a joint mixed effects model using G-computation in which the longitudinal outcome and treatment assignment models share common random effects to estimate causal effects with pre-defined treatment timing. Clifton and Laber (2020) reviewed the use of Q-learning, a general class of reinforcement learning methods, in estimating optimal treatment regimens taking the timing of treatments as given. Zhao et al. (2011) attempted to optimize the timing to initiate second-line therapy in the context of clinical trials with two-stage treatments using Q-learning, but only considered two options (i.e., immediately or delayed after induction therapy). Other work on optimizing intervention timing with a fixed number of treatment options include initiation of antiretroviral therapy in HIV (Robins et al., 2008), just-in-time adaptive interventions in mobile health (Nahum-Shani et al., 2018; Carpenter et al., 2020), and advantage doubly robust policy learning that optimizes when to treat (Nie et al., 2021). Guan et al. (2019) developed a Bayesian nonparametric method that learns to recommend a regular recall time for patients with periodontal diseases. However, their method only picks the recall time out of a few pre-defined choices (e.g., 3 months, 6 months, and 9 months) and thus is not applicable to complicated scenarios like the one introduced in Section 1.1: at each visitation after kidney transplantation, the next visitation time has to be carefully scheduled given the current clinical measurement in order to maximize the patient’s health outcome. For instance, when patients’ kidney function is relatively stable, they should be instructed to wait longer until the next visitation compared to those who are less stable.

1.3 Why use our method?

To the best of our knowledge, the proposed approach is the first general methodology for estimating personalized optimal clinical decisions with timing. The method is cutting-edge because (1) we build a generative probabilistic model that properly handles clinical decisions with timing; (2) we embed this decision process into a Bayesian joint model that also models clinical observations; (3) we propose a Bayesian optimization procedure to optimize personalized treatment schedules alongside other clinical decisions while accounting for uncertainties in clinical observations based on the posterior inference of the proposed Bayesian joint model.

Our decision model is a marked temporal point process (MTPP) (Aalen et al., 2008), which is a natural tool to model discrete events in continuous time. It has been widely applied and become increasingly popular in various domains, including social science (Butts and Marcum, 2017), medical analytics (Liu et al., 2018), finance (Hawkes, 2018), and stochastic optimal control (Tabibian et al., 2019). In our example application of Section 1.1, each follow-up visitation is an event: the visitation time is assumed to be stochastically scheduled according to the probability distribution characterized by the proposed MTPP; and the assigned tacrolimus dosage, when the visitation happens, is treated as the corresponding “mark.”

The proposed MTPP for clinical decisions is then embedded into a Bayesian joint model where it shares certain structures and parameters with the other components modeling clinical observations, including longitudinal creatinine measurements and patient survival in the example application of Section 1.1. Such design allows our model to capture the complicated mutual influence between clinical observations (e.g., creatinine levels) and decisions (e.g., treatment schedules and tacrolimus dosages). We fit the proposed Bayesian joint model to the data and obtain the posterior inference, which not only learns the dynamic patterns of clinical observations with uncertainty quantification, but also estimates the parameters in modeling clinical decisions as the estimates of how physicians treat patients in practice. The posterior means of these parameters can be used as the initial values in the optimization procedure so that clinical decisions can be refined to optimize patients’ survival.

Next, we let the decision model interact with the observation model in an optimization procedure. This technique is known as “reinforcement learning” (Sutton and Barto, 2018): the decision model (also called the “policy”) is reinforced, by the feedback from the observation model (also called the “environment”), to give personalized optimal treatment schedules and dosages that would improve the expected health outcome for each patient. The Bayesian nature of our approach allows the learning to account for parameter uncertainties in the observation model. Figure 2 illustrates the proposed two-step Bayesian approach. The left box displays the observational data including the clinical measurements and dosages at a sequence of visitations, and the right-headed arrow next to it stands for the Bayesian joint model (i.e., the first step), pointing to the right box, which represents the learned inference from the Bayesian joint model, including inference on longitudinal measurements, the hazard function of survival, dosages, and the intensity function of visitation timing. The vertical arrows inside the right box indicates the policy optimization method (i.e., the second step), in which the decision

model learns to achieve higher reward by interacting with the learned measurement model. The R package *doct* (short for “Decisions Optimized in Continuous Time”) implementing the proposed model and algorithm is available at <https://github.com/YanxunXu/doct>.




Figure 2: Illustration of the proposed method.

The rest of the paper is organized as follows. In Section 2, we outline the overall framework of the proposed two-step Bayesian approach. In Section 3, we elaborate on the first step of developing a Bayesian joint model consisting of the decision submodel (for visitation schedules and dosages) and the observation submodel (for clinical longitudinal measurements and patient survival). In Section 4, we elaborate on the second step of the optimization procedure for the decision model. We evaluate the proposed approach through simulation studies in Section 5 and applying it to the DIVAT kidney transplantation dataset in Section 6. Lastly, we conclude the paper with a discussion in Section 7.

2 Problem formulation

Our goal is to optimize personalized clinical decision including scheduling a patient’s follow-up visitations and prescribing dosages to maximize the patient’s health outcome, e.g., the graft survival time in the kidney transplantation application. In this section, we first introduce the notations, and then outline the framework of the proposed two-step Bayesian approach.

For each patient i , let T_i and C_i denote the graft survival and administration censoring times, respectively, $i = 1, \dots, I$. We observe only $\tilde{T}_i = \min(T_i, C_i)$ and the censoring indicator $\delta_i = \mathbb{1}(T_i \leq C_i)$. The graft survival would typically be affected by the patient's baseline risk factors, denoted by $\mathbf{x}_i$, including the patient's age when receiving the transplantation and the donor type. At each visitation time $t_{i,j}$, the clinical measurement $y_{i,j}$ of interest would be taken, $j = 0, \dots, J_i$, and $t_{i,J_i} \leq \tilde{T}_i$. Note here $t_{i,0} = 0$ denotes the transplantation date of patient i , and $y_{i,0}$ denotes the initial creatinine level. Then the physician would prescribe the dosage $d_{i,j}$, and schedule the next visitation time $t_{i,j+1}$. In the kidney transplantation application, $y_{i,j}$ is the logarithm of the creatinine level ($\mu\text{mol/l}$), and $d_{i,j}$ is the logarithm of the tacrolimus level (ng/ml). We represent the i -th patient's sequence of visitations and assigned dosages by $\mathbf{e}_{i,\tilde{T}_i} = \{(t_{i,0}, d_{i,0}), \dots, (t_{i,J_i}, d_{i,J_i})\}$. Denote $\mathbf{y}_i = (y_{i,0}, \dots, y_{i,J_i})$. Our data can be summarized as $\mathcal{D} = \{\mathbf{x}_i, \tilde{T}_i, \delta_i, \mathbf{y}_i, \mathbf{e}_{i,\tilde{T}_i}\}_{i=1}^I$.

Assume that the physician assigns the dosage $d_{i,j}$ and schedule the next visitation time $t_{i,j+1}$ based on the patient's baseline covariates $\mathbf{x}_i$ and the longitudinal creatinine measurements $\overline{y_{i,j}} = \{y_{i,j'} : j' \leq j\}$, we write the joint model of the observed clinical observations and decisions as

$$\begin{aligned} \prod_{i=1}^I p(\tilde{T}_i, \delta_i, \mathbf{y}_i, \mathbf{e}_{i,\tilde{T}_i} \mid \mathbf{x}_i) &= \prod_{i=1}^I \left[p(\tilde{T}_i, \delta_i \mid \mathbf{y}_i, \mathbf{e}_{i,\tilde{T}_i}, \mathbf{x}_i) \right. \\ &\quad \times \left. \prod_j p(y_{i,j} \mid t_{i,j}, d_{i,j-1}, \mathbf{x}_i) p(d_{i,j} \mid \mathbf{x}_i, \overline{y_{i,j}}) p(t_{i,j+1} \mid \mathbf{x}_i, \overline{y_{i,j}}) \right]. \end{aligned} \quad (2.1)$$

We leave the details of the chosen parametrization for this joint model to Section 3. Throughout this paper, we use $p(\cdot)$ to denote the probability density. In order to identify causal effects, we make the standard assumptions of sequential ignorability and consistency throughout this paper (Murphy et al., 2003). The positivity assumption is satisfied in our model due to the use of the marked temporal point process (MTPP) (Aalen et al., 2008) for modelling stochastic clinical decisions in continuous time, which will be introduced in Section 3.1.

Assume that the joint model (2.1) is parameterized with $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$, where $\boldsymbol{\theta}$ denotes the set of parameters related to clinical decisions that control patients' follow-up schedules and dosages at follow-up visitations, and $\boldsymbol{\phi}$ denotes the remaining parameters. We aim to learn the parameters $\boldsymbol{\theta}$ of the clinical decision model such that, for any (future/hypothetical) patient i , a certain reward R_i can be maximized by prescribing the right amount of dosages and scheduling the visitations at the right times. Formally, we aim to maximize the expected reward for any patient i with baseline covariates $\mathbf{x}_i$:

$$G_i(\boldsymbol{\theta}) = \int E_{(\mathbf{y}_i, T_i, \mathbf{e}_{i,T_i}) \sim p(\mathbf{y}_i, T_i, \mathbf{e}_{i,T_i} \mid \boldsymbol{\theta}, \boldsymbol{\phi}, \mathbf{x}_i)} [R_i] p(\boldsymbol{\phi} \mid \mathcal{D}) d\boldsymbol{\phi}. \quad (2.2)$$

Here the expectation is taken over all possible stochastic realizations of the clinical measurements $\mathbf{y}_i$, the survival time T_i , and the clinical decisions $\mathbf{e}_{i,T_i}$ from their joint distribution $p(\mathbf{y}_i, T_i, \mathbf{e}_{i,T_i} \mid \boldsymbol{\theta}, \boldsymbol{\phi}, \mathbf{x}_i)$; $\mathcal{D}$ denotes the observed data; and $p(\boldsymbol{\phi} \mid \mathcal{D})$ is the posterior distribution of $\boldsymbol{\phi}$ obtained from the posterior inference of the Bayesian

joint model (2.1). Note that the reward is stochastic, because the clinical measurements and decisions (i.e., dosages and visitation times), and their influence on the reward are all stochastic. In our kidney transplantation application, we can consider the patient's median survival time as the reward. More details on the reward will be introduced in Section 4.

To find the optimal clinical decision parameter $\tilde{\theta}_i$ that maximizes the expected reward $G_i(\theta)$ for patient i after integrating out the uncertainty in the longitudinal process and the survival distribution, we propose a two-step Bayesian approach. In the first step that will be introduced in Section 3, we will fit the joint model (2.1) to the observed data and assign priors for parameters θ and ϕ , then obtain posterior inference through Markov chain Monte Carlo (MCMC) simulations. In the second step that will be introduced in Section 4, we will propose a policy gradient method using stochastic gradient descent (SGD) (Ruder, 2016) to optimize personalized clinical decision. The uncertainties in the clinical observations will be incorporated by integrating over the posterior distribution of $p(\phi \mid \mathcal{D})$ when calculating the reward (2.2). The estimated posterior mean of θ from the first step can be used as an initial value in the optimization procedure for efficient learning.

3 First step: a Bayesian joint model

In this section, we describe the proposed Bayesian joint model for both clinical decisions and observations. In Section 3.1, we introduce the clinical decision model for follow-up visitation schedules and dosages; in Section 3.2, we introduce the clinical observation model for longitudinal measurements and time-to-event data, which are linked to the decision model through parameter sharing. To facilitate our presentation and readers' understanding, we will use the kidney transplantation example and the DIVAT data to illustrate the model. However, the proposed method is applicable to general medical settings since the patterns that the method can capture are not tied to this particular application.

3.1 Modeling clinical decisions

Modeling event data with marker information is important to learn the latent mechanisms that govern the observed stochastic event patterns over time in many domains, such as social science (Butts and Marcum, 2017) and medical analytics (Liu et al., 2018). Marked temporal point processes (Aalen et al., 2008) are a general framework for modeling such event data. Formally, a marked temporal point process is a random process, the realization of which consists of a sequence of events localized in time, i.e., $\mathcal{H} = \{(t_0, d_0), (t_1, d_1), \dots, (t_J, d_J)\}$ with the occurrence time of event j being $t_j \in \mathbb{R}^+$ and d_j is the associated mark. In our application, t_j represents the time when a patient visits an outpatient center and d_j represents the tacrolimus dosage assigned by the physician. The first event is defined as the day of transplantation at $t_0 = 0$ with an initial dose d_0 .

Denote the event history up to time t to be $\mathcal{H}_t = \{(t_j, d_j) \in \mathcal{H} \mid t_j < t\}$. Under MTPP, the instantaneous rate of the event is characterized by a conditional inten-

sity function $\lambda(t)$, namely $\lambda(t) = \lim_{dt \rightarrow 0} \frac{\Pr\{\text{event happens in } [t, t+dt) | \mathcal{H}_t\}}{dt}$. Common forms of the conditional intensity function $\lambda(t)$ include Poisson process (Zhu and Li, 2018), Gamma process (Shibue and Komaki, 2020), Hawkes process (Hawkes, 1971). However, these common models cannot capture complicated patterns in many medical applications. For instance, as shown in Figure 3(a) that plots the empirical intensity of the amount of time between visitations for different ranges of creatinine levels in the DIVAT data, the elapsed time between follow-up visitations depends on the creatinine level. Also, the empirical intensities of visitations are observed to quickly rise to a peak and then fall down accompanied by moderate oscillations. Such complication is beyond the capacity of the Poisson process that assumes a constant intensity and the Gamma process whose intensity function is monotonic. The Hawkes process assumes that the past events always elevate the intensities of future events and this “self-exciting” effect is additive—it is also apparently not the dynamics that the visitations in the DIVAT data actually follow.




Figure 3: Panel (a) shows the empirical intensity plot for the amount of time (in days) between follow-up visitations. Panel (b) plots an example of how creatinine levels and model parameters affect the visitation intensity, where $\xi = 2$, $\beta_\alpha = (10, -1.8)^T$, $\mu = -4.4$, $\nu_1 = 1.5$, and $\nu_2 = 1$.

We propose a flexible conditional intensity function that also incorporates human intuition: it takes longitudinal clinical measurements into account and captures patients’ heterogeneity. Recall that $y_{i,j}$ denotes the logarithm of the creatinine level for patient i at the j -th follow-up visitation occurring at time $t_{i,j}$ (days). Our conditional intensity function makes use of a Gamma density function as follows:

$$\lambda_i(t) = \underbrace{\exp(\mu)}_{\text{Baseline Intensity}} + \underbrace{\alpha_{i,j} (t - t_{i,j})^{\kappa-1} e^{-\gamma(t-t_{i,j})} \frac{\gamma^\kappa}{\Gamma(\kappa)}}_{\text{Gamma density}} \text{ for } t \in (t_{i,j}, t_{i,j+1}], \quad (3.1)$$

where $\alpha_{i,j} > 0$, $\gamma > 0$, $\kappa \geq 1$. The parameter $\alpha_{i,j}$ is patient-specific so that our intensity function λ_i is personalized. We set $\kappa = \exp(\nu_2) + 1 > 1$ so that the intensity rises

to a “global peak” and then decreases: it would eventually approach to the “baseline level” $\exp(\mu)$ unless the next visitation happens and sets up a new intensity curve. For easy interpretation, we parameterize γ as $\gamma = \exp(\nu_2 - \nu_1)$ such that the “peak time” (i.e., when the peak of the intensity function occurs) can be easily computed as $\frac{\kappa-1}{\gamma} = \exp(\nu_1)$. Moreover, since the intensity level often depends on the clinical measurement (e.g., as in Figure 3(a), a higher creatinine level implies a higher intensity), we condition the parameter $\alpha_{i,j}$, which controls the peak intensity for patient i between time $t_{i,j}$ and $t_{i,j+1}$, on the clinical measurement taken at the j -th visitation:

$$\alpha_{i,j} = \frac{\xi}{1 + \exp((1, y_{i,j})\beta_\alpha)}.$$

This design reflects the human intuition that the time of “next visit” is usually determined based on the clinical measurement of “this visit.” Note that our design allows incorporating other covariates (i.e., measurements) by simply augmenting them to the vector $(1, y_{i,j})$. Figure 3(b) shows how the visitation intensity under our model is affected by the most recent creatinine level $y_{i,j}$ (and thus the magnitude parameter $\alpha_{i,j}$) given a specific set of parameter values.

Next, we model the dosage $d_{i,j}$ at the j -th visitation of patient i as the “mark” of the visitation event. Generally speaking, the physician would assign a dosage based on the patient’s current clinical measurement $y_{i,j}$ and potential risk factors $\mathbf{x}_i$. We assume the following dosage model reflecting this knowledge:

$$d_{i,j} = (1, y_{i,j}, \mathbf{x}_i)\beta_d + \zeta_{i,j}, \quad (3.2)$$

where $\zeta_{i,j} \stackrel{\text{i.i.d.}}{\sim} \text{Normal}(0, \sigma_d^2)$. This Gaussian error assumption is suitable for our dosage data, which has continuous values, but can be easily modified for other dosage/treatment types (e.g., generalized linear regression models for discrete treatment choices). In the kidney transplantation application, we assume that the clinical decision at time $t_{i,j}$ (i.e., assigning a dosage $d_{i,j}$ and scheduling the next visitation time $t_{i,j+1}$) is independent of the patient’s history conditional on the measured creatinine at the current visit $y_{i,j}$. This is reasonable in our application since the duration between two consecutive visitations for patients after kidney transplantation can be months or even longer. However, this assumption can be relaxed as in other popular DTR methods such as Q-learning (Clifton and Laber, 2020) if desired. For instance, the patient’s past creatinine levels can be included in the vector $(1, y_{i,j}, \mathbf{x}_i)$ of (3.2) to model the effect of past creatinine levels on the dosage. Thus, the probability density of the i -th patient’s sequence of visitations and assigned dosages $\mathbf{e}_{i,T_i} = \{(t_{i,0}, d_{i,0}), \dots, (t_{i,J_i}, d_{i,J_i})\}$ up to time T_i can be written as

$$\begin{aligned} & p(\mathbf{e}_{i,T_i} \mid \mathbf{y}_i, \mathbf{x}_i, \beta_v, \beta_d, \sigma_d^2) \\ &= \underbrace{\exp\left(-\int_0^{T_i} \lambda_i(t \mid \mathbf{y}_i, \beta_v) dt\right)}_{\text{Prob. of no visits at } t \in [0, T_i] \setminus \{t_{i,j}\}_{j=1}^{J_i}} \prod_{j=0}^{J_i} \underbrace{p(d_{i,j} \mid y_{i,j}, \mathbf{x}_i, \beta_d, \sigma_d^2)}_{\text{Prob. of dosage}} \prod_{j=1}^{J_i} \underbrace{\lambda_i(t_{i,j} \mid y_{i,j-1}, \beta_v)}_{\text{Prob. of a visit at } t_{i,j}}^{(3.1)}, \end{aligned} \quad (3.3)$$

where $\mathbf{y}_i = (y_{i,0}, \dots, y_{i,J_i})$, $\beta_v = \{\mu, \nu_1, \nu_2, \xi, \beta_\alpha\}$.

3.2 Linking clinical observations with clinical decisions

In this section, we introduce the proposed Bayesian joint model that links the submodel of the longitudinal measurements and time-to-event data to the MTPP of the clinical decisions by carefully designing parameter sharing in order to capture the mutual influence between clinical observations and decisions. Shortly in Section 4, we will leverage this joint model to optimize clinical decisions with the goal of maximizing patients' survival.

Our clinical observation model is composed of two submodels—a linear mixed effects model for longitudinal clinical measurements (e.g., creatinine levels) and a time-to-event model for patient survival (e.g., graft survival time after kidney transplantation). The two submodels are then connected by sharing random effects (Rizopoulos et al., 2014). Recall that $y_{i,j} = y_i(t_{i,j})$ denotes the longitudinal measurement value for patient i at j -th follow-up visitation at time $t_{i,j}$, $i = 1, \dots, I$, $j = 0, \dots, J_i$. Let $y_i^*(t)$ be the underlying true but unobserved longitudinal process at time $t \geq 0$. We assume

$$y_i(t) = y_i^*(t) + \epsilon_{i,j} = \mathbf{z}_i(t)\boldsymbol{\beta}_l + \mathbf{r}_i(t)\mathbf{b}_i + \epsilon_i(t), \quad (3.4)$$

where $\epsilon_i(t) \stackrel{\text{i.i.d.}}{\sim} \text{Normal}(0, \sigma_l^2)$ and $\mathbf{b}_i \sim \text{Normal}(\mathbf{0}, \boldsymbol{\Sigma}_b)$. The covariate vectors $\mathbf{z}_i(t)$ and $\mathbf{r}_i(t)$ are associated with fixed and random effects respectively:

$$\mathbf{z}_i(t) = (1, d_i(t), \mathbf{x}_i, t, t^2) \text{ and } \mathbf{r}_i(t) = (1, d_i(t), t),$$

where $d_i(t)$ at time t is the dosage assigned by the physician at the most recent visitation, i.e., $d_i(t) = d_{i,j}$ for $t \in (t_{i,j}, t_{i,j+1}]$. The temporal dependence of $\mathbf{z}$ and $\mathbf{r}$ on the dosage d captures the drug effect on the longitudinal measurements of interest: in the kidney transplantation application, it is supposed to capture the suppressive effect of tacrolimus on the creatinine level. Denote $\mathbf{d}_i = (d_{i,0}, \dots, d_{i,J_i})$, the probability of the observed sequence of creatinine measurements $\mathbf{y}_i$ is

$$p(\mathbf{y}_i \mid \mathbf{d}_i, \mathbf{x}_i, \boldsymbol{\beta}_l, \sigma_l^2, \mathbf{b}_i) = \prod_{j=1}^{J_i} p(y_{i,j} \mid t_{i,j}, d_{i,j-1}, \mathbf{x}_i, \boldsymbol{\beta}_l, \sigma_l^2, \mathbf{b}_i). \quad (3.5)$$

Next, we construct the time-to-event submodel depending on the underlying true longitudinal trajectory $y_i^*(t)$ and the MTPP that models clinical decisions. We consider a Weibull proportional hazards model as follows:

$$h_i(t) = \exp \left(- \left(\underbrace{\beta_{s1} y_i^*(t)}_{\text{longitudinal effect}} + \underbrace{\beta_{s2} d_i(t) + \beta_{s3} \text{Tox}_i(t)}_{\text{dosage effect}} + \underbrace{\beta_{s4} \alpha_i(t)}_{\text{visitation effect}} + h_0 \right) \omega t^{\omega-1} \right), \quad (3.6)$$

where ω is the shape parameter. If desired, more complex survival models can be explored, such as Cox proportional hazard models (Lin and Wei, 1989) and Bayesian nonparametric survival regression models (Xu et al., 2019). The dependence on $y_i^*(t)$ reflects the domain knowledge that the survival event is usually associated with the underlying health condition reflected by longitudinal measurements. The dosage effect

term in (3.6) measures the overall drug effect on the patient: $\beta_{s2}d_i(t)$ is the “instantaneous” effect while $\beta_{s3}\text{Tox}_i(t)$ is the “accumulated” effect:

$$\text{Tox}_i(t) = \int_0^t d_i(\tau) \eta_{tox} \exp(-(t - \tau)/\eta_{tox}) d\tau,$$

where the parameter η_{tox} controls the rate of the exponential weighting for the past dosages. In practice, the instantaneous effect is usually beneficial (e.g., tacrolimus reduces the likelihood of graft rejection or death) while the accumulated effect is often toxic (and that is why we name it Tox): e.g., a prolonged high dosage of tacrolimus might have adverse effects on kidneys, central nervous system, and gastrointestinal tract, thereby worsening a patient’s survival (Randhawa et al., 1997). We also link the survival submodel with the visitation model by defining $\alpha_i(t) = \alpha_{i,j}$ for $t \in (t_{i,j}, t_{i,j+1}]$ since a high visitation intensity (i.e., larger $\alpha_{i,j}$) typically implies a higher risk, e.g., graft failure and thus shorter expected survival time.

Recall that T_i and C_i denote the survival and censoring times for patient i , respectively. We assume $\tilde{T}_i = \min(T_i, C_i)$ and $\delta_i = \mathbb{1}(T_i \leq C_i)$. Denote $f_i(t)$ and $S_i(t)$ to be the corresponding density and survival functions of the hazard function (3.6): $S_i(t) = \exp(-\int_0^t h_i(u)du)$, $f_i(t) = h_i(t)S_i(t)$. We can write the survival likelihood for patient i as

$$\begin{aligned} p(\tilde{T}_i, \delta_i \mid \mathbf{y}_i, \mathbf{x}_i, \mathbf{e}_{i,\tilde{T}_i}, \boldsymbol{\beta}_l, \mathbf{b}_i, \boldsymbol{\beta}_s) &= f_i(\tilde{T}_i \mid \mathbf{y}_i, \mathbf{x}_i, \mathbf{e}_{i,\tilde{T}_i}, \boldsymbol{\beta}_l, \mathbf{b}_i, \boldsymbol{\beta}_s)^{\delta_i} \\ &\quad \times S_i(\tilde{T}_i \mid \mathbf{y}_i, \mathbf{x}_i, \mathbf{e}_{i,\tilde{T}_i}, \boldsymbol{\beta}_l, \mathbf{b}_i, \boldsymbol{\beta}_s)^{1-\delta_i}, \end{aligned} \quad (3.7)$$

where $\boldsymbol{\beta}_s = \{\omega, \beta_{s1}, \beta_{s2}, \beta_{s3}, \beta_{s4}, h_0, \eta_{tox}, \boldsymbol{\beta}_\alpha, \xi\}$.

In summary, we propose a joint model consisting of an MTPP for clinical decisions including follow-up visitation schedules and dosages, a linear mixed effects model for longitudinal clinical measurements, and a time-to-event model for the patient survival; they are inter-connected by sharing structures and parameters. The joint probability of the clinical observations and decisions can then factor as

$$\begin{aligned} &\prod_{i=1}^I p(\mathbf{y}_i, \mathbf{e}_{i,\tilde{T}_i}, \tilde{T}_i, \delta_i \mid \mathbf{x}_i, \boldsymbol{\beta}_l, \boldsymbol{\beta}_d, \boldsymbol{\beta}_v, \boldsymbol{\beta}_s, \mathbf{b}_i, \sigma_l^2, \sigma_d^2) \\ &\propto \prod_{i=1}^I \left(\underbrace{p(\mathbf{e}_{i,\tilde{T}_i} \mid \mathbf{y}_i, \mathbf{x}_i, \boldsymbol{\beta}_v, \boldsymbol{\beta}_d, \sigma_d^2)}_{(3.3)} \underbrace{p(\mathbf{y}_i \mid \mathbf{d}_i, \mathbf{x}_i, \boldsymbol{\beta}_l, \sigma_l^2, \mathbf{b}_i)}_{(3.5)} \underbrace{p(\tilde{T}_i, \delta_i \mid \mathbf{y}_i, \mathbf{x}_i, \mathbf{e}_{i,\tilde{T}_i}, \boldsymbol{\beta}_l, \mathbf{b}_i, \boldsymbol{\beta}_s)}_{(3.7)} \right). \end{aligned} \quad (3.8)$$

Note that (3.8) appear to have circulation of $\mathbf{d}_i \mid \mathbf{y}_i$ in (3.3) and $\mathbf{y}_i \mid \mathbf{d}_i$ in (3.5). However, it is a valid factorization because the dependencies are temporal: each dosage $d_{i,j}$ is conditional on the *current* measurement $y_{i,j}$ in $\mathbf{d}_i \mid \mathbf{y}_i$, while the *next* measurement $y_{i,j+1}$ is conditional on the current dosage $d_{i,j}$ in $\mathbf{y}_i \mid \mathbf{d}_i$.

We complete the model by imposing the following priors: $\boldsymbol{\beta}_d \sim \text{Normal}(\boldsymbol{\beta}_{d0}, \boldsymbol{\Sigma}_{\boldsymbol{\beta}_d})$, $\sigma_d^2 \sim \text{InverseGamma}(\pi_{d1}, \pi_{d2})$, $\boldsymbol{\beta}_l \sim \text{Normal}(\boldsymbol{\beta}_{l0}, \boldsymbol{\Sigma}_{\boldsymbol{\beta}_l})$, $\sigma_l^2 \sim \text{InverseGamma}(\pi_{l1}, \pi_{l2})$

for conjugacy. We assume a flat prior for Σ_b . When conjugacy is unattainable for the visitation and survival parameters, we assume $\beta_{s1}, \beta_{s2}, \beta_{s3}, \beta_{s4}, h_0 \sim \text{Normal}(\beta_{s0}, \sigma_{s0}^2)$, $\eta_{tox} \sim \text{Gamma}(\pi_{s1}, \pi_{s2})$, $\omega \sim \text{Gamma}(\pi_{s3}, \pi_{s4})$, $\mu, \nu_1, \nu_2 \sim \text{Normal}(\beta_{v0}, \sigma_{v0}^2)$, $\beta_\alpha \sim \text{Normal}(\beta_{\alpha0}, \Sigma_{\beta_\alpha})$, and $\xi \sim \text{Gamma}(\pi_{v1}, \pi_{v2})$. We carry out posterior inference using the Markov chain Monte Carlo (MCMC) sampler. The details are included in the Supplementary Material Section A (Hua et al., 2021).

4 Second step: optimize personalized clinical decision

In this section, we propose a policy gradient method using stochastic gradient descent (SGD) (Ruder, 2016) to optimize personalized clinical decision including scheduling a patient's follow-up visitations and prescribing dosages to maximize the patient's health outcome, e.g., the graft survival time in the kidney transplantation application.

Let $\theta = (\nu_1, \nu_2, \mu, \beta_d, \sigma_d^2)$ denote the set of “policy” parameters related to clinical decisions, i.e., the parameters that only appear in the conditional intensity function (3.1) and the mark distribution (3.2), which control patients' follow-up schedules and dosages at follow-up visitations. Let $\phi = (\beta_s, \mathbf{b}_i, \beta_l, \sigma_l^2)$ denote the set of “observation” parameters, i.e., all other parameters in the joint model (3.8). Recall that in Section 2, we define the expected reward for any future/hypothetical patient i with baseline covariates $\mathbf{x}_i$ to be:

$$G_i(\theta) = \int E_{(\mathbf{y}_i, T_i, \mathbf{e}_{i, T_i}) \sim p(\mathbf{y}_i, T_i, \mathbf{e}_{i, T_i} | \theta, \phi, \mathbf{x}_i)} [R_i] p(\phi | \mathcal{D}) d\phi. \quad (4.1)$$

The expectation is taken over all possible stochastic realizations of $(\mathbf{y}_i, T_i, \mathbf{e}_{i, T_i})$. In the kidney transplantation application, we define a personalized reward function R_i as the log-scaled median survival time to optimize patients' survival: $R_i = \log(\hat{T}_i)$, where $S_i(\hat{T}_i) = 0.5$. This reward provides computational and variance-reduction advantages over a naive choice of the survival time itself. If desired, other reward functions can be considered. For example, if a physician or patient would like to take into consideration the healthcare cost per visit, we could penalize the number of visitations in the reward function, e.g., $R_i = \log(\hat{T}_i) + \eta_0 C_i$, where η_0 is a tuning parameter and C_i is the number of visitations. Our goal is to find, for any patient i with baseline covariates $\mathbf{x}_i$, the optimal clinical decision, represented by θ_i , to maximize the expected reward $G_i(\theta)$ after integrating out the uncertainty in the longitudinal process and the survival distribution:

$$\text{maximize}_{p(\mathbf{e}_{i, T_i} | \theta)} G_i(\theta),$$

where $p(\mathbf{e}_{i, T_i} | \theta)$ is the probability density of the MTPP.

To find the optimal clinical decision parameter $\tilde{\theta}_i$ for patient i , we use stochastic gradient descent (SGD) (Robbins and Monro, 1951), i.e., $\theta_{i, m+1} = \theta_{i, m} + s_{i, m} \nabla_{\theta} G_i(\theta) |_{\theta=\theta_{i, m}}$, which requires computing the gradient of the expected reward: $\nabla_{\theta} G_i(\theta)$. As the expectation is taken over realizations of the joint distribution $p(\mathbf{y}_i, T_i, \mathbf{e}_{i, T_i} | \theta, \phi, \mathbf{x}_i)$, it is intractable to directly compute $\nabla_{\theta} G_i(\theta)$. Fortunately, we can indirectly compute this gradient by taking the expectation of the reward-weighted gradient of log-policy. Precisely,

Proposition 1. For any patient i with baseline covariates $\mathbf{x}_i$, given a joint distribution $p(\mathbf{y}_i, T_i, \mathbf{e}_{i,T_i} \mid \boldsymbol{\theta}, \boldsymbol{\phi}, \mathbf{x}_i)$, the gradient of the expected reward $G_i(\boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}$ is:

$$\nabla_{\boldsymbol{\theta}} G_i(\boldsymbol{\theta}) = \int E_{(\mathbf{y}_i, T_i, \mathbf{e}_{i,T_i}) \sim p(\mathbf{y}_i, T_i, \mathbf{e}_{i,T_i} \mid \boldsymbol{\theta}, \boldsymbol{\phi}, \mathbf{x}_i)} [R_i \nabla_{\boldsymbol{\theta}} \log p(\mathbf{e}_{i,T_i} \mid \mathbf{y}_i, \mathbf{x}_i, \boldsymbol{\phi}, \boldsymbol{\theta})] p(\boldsymbol{\phi} \mid \mathcal{D}) d\boldsymbol{\phi},$$

where $p(\mathbf{e}_{i,T_i} \mid \mathbf{y}_i, \mathbf{x}_i, \boldsymbol{\phi}, \boldsymbol{\theta})$ is the probability of the patient's sequence of visitations and assigned dosages in (3.3).

We leave the detailed proof of Proposition 1 to Supplementary Section B.

According to Proposition 1, in order to compute $\nabla_{\boldsymbol{\theta}} G_i(\boldsymbol{\theta})$, we first need to be able to sample $\mathbf{y}_i, T_i, \mathbf{e}_{i,T_i}$ from $p(\mathbf{y}_i, T_i, \mathbf{e}_{i,T_i} \mid \boldsymbol{\theta}, \boldsymbol{\phi})$ and calculate R_i from the generated samples. We sample the j -th follow-up visitation time $t_{i,j}$ and the survival time T_i using an inverse transform sampling method: first computing the cumulative distribution function (CDF) of the distribution, sampling a random number U from $\text{Uniform}(0, 1)$, and then inverting the CDF function at U to yield the visitation/survival time (Giesecke et al., 2011). If the j -th visitation time occurs before the survival time, i.e., $t_{i,j} < T_i$, we sample $y_{i,j}$ and $d_{i,j}$ from their respective distributions and continue to sample the $(j+1)$ -th visitation time and the survival time. We iteratively sample follow-up visitation times, survival times, longitudinal measurements, and dosages until the sampled survival event occurs before the next visitation time. After obtaining samples of $\mathbf{y}_i, T_i, \mathbf{e}_{i,T_i}$, we can easily compute R_i . We describe the sampling process for a general R_i in Algorithm 1. The algorithm details of sampling $\mathbf{y}_i, T_i, \mathbf{e}_{i,T_i}, R_i$ for the reward being the log median survival time are provided in Supplementary Section C.

Next we compute the gradient of the log-likelihood of the MTPP, $\nabla_{\boldsymbol{\theta}} \log p(\mathbf{e}_{i,T_i} \mid \mathbf{y}_i, \mathbf{x}_i, \boldsymbol{\phi}, \boldsymbol{\theta})$, using the parametrization defined in (3.3). The details are described in Supplementary Section D. Lastly, we integrate out $\boldsymbol{\phi}$ in computing $\nabla_{\boldsymbol{\theta}} G_i(\boldsymbol{\theta})$ using the Monte Carlo method since it is analytically intractable. Suppose that we have K MCMC draws from the posterior distribution of $\boldsymbol{\phi}$ and we denote the k -th draw as $\boldsymbol{\phi}^{(k)}$, then $\nabla_{\boldsymbol{\theta}} G_i(\boldsymbol{\theta})$ can be approximated as follows:

$$\nabla_{\boldsymbol{\theta}} G_i(\boldsymbol{\theta}) \approx \frac{\sum_{k=1}^K E_{(\mathbf{y}_i, T_i, \mathbf{e}_{i,T_i}) \sim p(\mathbf{y}_i, T_i, \mathbf{e}_{i,T_i} \mid \boldsymbol{\theta}, \boldsymbol{\phi}^{(k)}, \mathbf{x}_i)} [R_i \nabla_{\boldsymbol{\theta}} \log p(\mathbf{e}_{i,T_i} \mid \mathbf{y}_i, \mathbf{x}_i, \boldsymbol{\phi}^{(k)}, \boldsymbol{\theta})]}{K}. \quad (4.2)$$

To compute each term of the summation in the numerator of (4.2), we first sample $T_i, \mathbf{y}_i$, and $\mathbf{e}_{i,T_i}$ from $p(\mathbf{y}_i, T_i, \mathbf{e}_{i,T_i} \mid \boldsymbol{\theta}, \boldsymbol{\phi}^{(k)}, \mathbf{x}_i)$ using Algorithm 1 to compute R_i for each $\boldsymbol{\phi}^{(k)}$, then multiply the gradient of the log-probabilities of visitation times and dosages under the MTPP policy. We use an adaptive step size algorithm and choose the step size to be: $s_{i,m} = \frac{0.01}{\sqrt{\sum_{l=m-50}^{m-1} \nabla_{\boldsymbol{\theta}} G_i(\boldsymbol{\theta}_{i,l})^2}}$. The entire SGD algorithm for finding the optimal parameter $\tilde{\boldsymbol{\theta}}_i$ is described in Algorithm 2, where $\overline{G_i(\boldsymbol{\theta}_{i,m})}$ denotes the expected reward in iteration m . We select the optimal policy $\tilde{\boldsymbol{\theta}}_i$ to be the one yielding the highest expected reward across all iterations. Note that, in the step 7 of Algorithm 2, we subtract the average reward from each individual reward: this “baseline subtraction” trick significantly reduce the variance while still yielding an unbiased estimate of the gradient (Williams, 1992; Greensmith et al., 2004).

Algorithm 1 Sampling $\mathbf{y}_i, T_i, \mathbf{e}_{i,T_i}$ from the joint model and computing R_i .

Let $\mathbf{e}_{i,T_i} = \{(t_{i,0}, d_{i,0}), \dots, (t_{i,J_i}, d_{i,J_i})\}$, and $\mathbf{y}_i = (y_i(t_{i,1}), \dots, y_i(t_{i,J_i}))$ denote the simulated follow-up schedules, dosages, and longitudinal data over J_i visitations until the survival time, T_i for any patient n with covariates $\mathbf{x}_i$.

Input: $\boldsymbol{\theta}, \phi, \mathbf{x}_i, y_{i,0}$

Output: $\mathbf{y}_i, T_i, \mathbf{e}_{i,T_i}, R_i$

```

1: Initialize  $j \leftarrow 1$ , continue  $\leftarrow \mathbf{true}$ 
2:  $t_{i,0} \leftarrow 0$ 
3:  $y_i(0) \leftarrow y_{i,0}$ 
4:  $d_{i,0} \leftarrow \text{Normal}((1, y_i(0), \mathbf{x}_i)\boldsymbol{\beta}_d, \sigma_d^2)$ 
5:  $\mathbf{b}_i \leftarrow \text{Normal}(\mathbf{0}, \boldsymbol{\Sigma}_b)$ 
6: while continue do
7:    $U_v \leftarrow \text{Uniform}(0, 1)$ 
8:   Solve for  $t_{i,j} : 1 - \exp(-\int_{t_{i,j-1}}^{t_{i,j}} \lambda_i(x)dx) = U_v$ 
9:    $U_s \leftarrow \text{Uniform}(0, 1)$ 
10:  Solve for  $T_i : 1 - \exp(-\int_{t_{i,j-1}}^{T_i} h_i(x)dx) = U_s$ 
11:  if  $T_i > t_{i,j}$  then
12:     $\mathbf{z}_i(t_{i,j}) \leftarrow (1, d_i(t_{i,j-1}), \mathbf{x}_i, t_{i,j}, t_{i,j}^2)$ ,  $\mathbf{r}_i(t_{i,j}) \leftarrow (1, d_i(t_{i,j-1}), t_{i,j})$ 
13:     $y_i(t_{i,j}) \leftarrow \text{Normal}(\mathbf{z}_i(t_{i,j})\boldsymbol{\beta}_l + \mathbf{r}_i(t_{i,j})\mathbf{b}_i, \sigma_l^2)$ 
14:     $d_{i,j} \leftarrow \text{Normal}((1, y_i(t_{i,j}), \mathbf{x}_i)\boldsymbol{\beta}_d, \sigma_d^2)$ 
15:     $j \leftarrow j + 1$ 
16:  else
17:     $J_i \leftarrow j - 1$ , continue  $\leftarrow \mathbf{false}$ 
18:     $\mathbf{e}_{i,T_i} \leftarrow \{(t_{i,0}, d_{i,0}), \dots, (t_{i,J_i}, d_{i,J_i})\}$  and  $\mathbf{y}_i \leftarrow (y_i(t_{i,1}), \dots, y_i(t_{i,J_i}))$ 
19:    Compute  $R_i$ 
20:  end if
21: end while

```

5 Simulation study

To demonstrate the advantage of the proposed Bayesian joint model, we compared it to an alternative model that breaks the connection between longitudinal and survival processes. Furthermore, to illustrate the benefit of optimizing the personalized clinical decision, we compared the expected reward under the estimated optimal clinical decision to alternative strategies of scheduling follow-up visitations on a regular basis, e.g., every three months (Israni et al., 2014).

5.1 Simulation setup

We simulated a dataset mimicking the DIVAT dataset composed of longitudinal creatinine measurements, follow-up schedules, tacrolimus dosages, and survival events for $I = 500$ patients. We considered three baseline covariates in $\mathbf{x}_i$: donor age (AgeD), delayed graft function (DGF), and body mass index (BMI). DGF is a binary variable with 1 indicating that the patient used dialysis within the first week of the trans-

Algorithm 2 Stochastic Gradient Descent for optimizing θ for any patient i .

Input $\theta_0, \phi^{(k)}$ ($k = 1, \dots, K$), $\mathbf{x}_i, y_{i,0}$.
Output $\tilde{\theta}_i$

- 1: Initialize $\theta_{i,1} \leftarrow \theta_0$
- 2: **for** $m:=1$ **to** $M-1$ **do**
- 3: **for** $k:=1$ **to** K **do**
- 4: do Algorithm 1 ($\theta_{i,m}, \phi^{(k)}, \mathbf{x}_i, y_{i,0}$) to sample $\mathbf{e}_{i,T_i^{(k)}}^{(k)}$ and $\mathbf{y}_i^{(k)}$, and compute $R_i^{(k)}$.
- 5: **end for**
- 6: $\overline{G_i(\theta_{i,m})} \leftarrow \frac{\sum_{k=1}^K R_i^{(k)}}{K}$
- 7: $\nabla_{\theta} G_i(\theta_{i,m}) \leftarrow \frac{\sum_{k=1}^K (R_i^{(k)} - \overline{G_i(\theta_{i,m})}) \nabla_{\theta} \log p(\mathbf{e}_{i,T_i^{(k)}}^{(k)} | \mathbf{y}_i^{(k)}, \mathbf{x}_i, \phi^{(k)}, \theta_{i,m})}{K}$
- 8: $\theta_{i,m+1} \leftarrow \theta_{i,m} + s_{i,m} \nabla_{\theta} G_i(\theta_{i,m})$
- 9: **end for**
- 10: $m^* \leftarrow \arg \max_m \overline{G_i(\theta_{i,m})}$
- 11: $\tilde{\theta}_i \leftarrow \theta_{i,m^*}$

plant, 0 otherwise. For each patient, the donor age and BMI were generated from Normal(52.5, 15.8²) and Normal(24.3, 4.5²), respectively, and then standardized. Patients' delayed graft functions were generated from Bernoulli(0.4) independently. In the MTPP model for follow-up schedules, the simulated true parameters were set to be $\nu_1 = 2.5$, $\nu_2 = 1.5$, $\mu = -4.8$, $\xi = 2$, and $\beta_{\alpha} = (9.5, -1.5)^T$ so that a higher creatinine level results in a higher visitation intensity; for assigning dosages, the simulated true β_d was set to be $(1, 0.2, 0.15, 0.2, 0.15)^T$ and $\sigma_d = 0.3$. In modeling log-transformed longitudinal creatinine levels, the simulated true parameters were set to be $\beta_l = (5.3, 0.1, 0.3, 0.4, 0.25, -1 \times 10^{-4}, 3 \times 10^{-8})^T$, $\sigma_l = 0.1$, and

$$\Sigma_b = \begin{bmatrix} 0.04 & 0 & 0 \\ 0 & 0.0049 & 0 \\ 0 & 0 & 10^{-8} \end{bmatrix}.$$

Note that the last two terms in the simulated true β_l were small since the times were recorded in days. Patients' initial log-transformed creatinine levels right after transplantation $y_{i,0}$'s were independently generated from Normal(5, 0.1²). In the survival submodel (3.6), we assumed that the simulated true parameters were $h_0 = 5$, $\omega = 1.05$, $\beta_{s1} = 1$, $\beta_{s2} = 0.9$, $\beta_{s3} = -0.75$, $\beta_{s4} = -5$, and $\eta_{tox} = 50$. The censoring times C_i 's were independently generated from Weibull(3, 8000). Based on the proposed Bayesian joint model in Section 3, we generated the data $\mathbf{y}_i, \mathbf{e}_{i,\tilde{T}_i}, \tilde{T}_i, \delta_i$ for each patient i , $i = 1, \dots, I$.

The simulated dataset had a total of 14,395 follow-up visitations for 500 patients with a 10.8% censoring rate. The median survival time was 1,684 days with the shortest being 24 days and the longest being 10,016 days. Supplementary Figure S1 plots the simulated longitudinal creatinine levels and follow-up schedules with dosages for four randomly selected patients.

5.2 Results: model fitting

We applied the proposed Bayesian joint model to the simulated dataset. The hyper-parameters were set to be $\beta_{d0} = \beta_{l0} = \beta_{\alpha0} = \mathbf{0}$, $\Sigma_{\beta_d} = \Sigma_{\beta_l} = \Sigma_{\beta_\alpha} = 100^2 I$, $\pi_{d1} = \pi_{d2} = \pi_{l1} = \pi_{l2} = \pi_{s3} = \pi_{s4} = 0.01$, $\pi_{s1} = \pi_{s2} = 0.01$, $\beta_{s0} = \beta_{v0} = 0$, $\sigma_{s0}^2 = \sigma_{v0}^2 = 100^2$, $\pi_{v1} = 400$, $\pi_{v2} = 200$. We ran 20,000 MCMC iterations with an initial burn-in of 5,000 iterations and a thinning factor of 50. The convergence was assessed using R package *coda*, including traceplots of the post-burn-in MCMC samples for some randomly selected parameters (Supplementary Figure S2), showing no issues of non-convergence. We first report on the performance of the proposed joint model in terms of parameter estimation. Supplementary Figure S3 plots the 95% estimated credible intervals (CIs) for selected parameters, showing that all 95% CIs are centered around the simulated true values. As another metric of performance, we computed the mean squared error (MSE) taken as the averaged squared errors between the post-burn-in MCMC posterior samples and the simulated true values. Supplementary Table S1 summarizes the MSE and the standard deviation of squared errors, indicating that the proposed joint model can accurately estimate parameters.

As the proposed model represents the first effort in the literature to jointly model clinical decisions, longitudinal markers, and the survival event, there is no existing method we can compare with. To demonstrate the advantage of jointly modeling longitudinal creatinine levels and the survival event, we compared the proposed model with an alternative “separate longitudinal and survival (SLS)” model that breaks the connection between the longitudinal and survival submodels by replacing the process $y_i^*(t)$ with the observational data $y_i(t)$ in the hazard model (3.6). We first compared the two models by checking their model adequacy using the Watanabe-Akaike information criterion (WAIC) (Watanabe and Opper, 2010): the joint model has a WAIC value of 226,982 while the SLS model has a WAIC of 226,992, indicating that the proposed joint model fits data slightly better. Furthermore, we compared the two models in terms of parameter estimation. Table 1 reports the simulated true values of parameters in the survival submodel, and posterior means of these parameters under the joint model and the SLS model with 95% CIs, showing that the joint model estimates parameters more accurately.

	Truth	Joint posterior mean (95% CI)	SLS posterior mean (95% CI)
β_{s1}	1	1.1 (0.92,1.26)	1.19 (0.94,1.6)
β_{s2}	0.9	1.25 (0.74,1.95)	1.41 (0.63,2.3)
β_{s3}	-0.75	-0.92 (-1.62,-0.33)	-1.03 (-1.8,-0.18)
β_{s4}	-5	-5.01 (-5.51,-4.47)	-5.16 (-6.14,-4.56)
h_0	5	4.36 (3.44,5.35)	3.89 (1.6,5.22)
ω	1.05	1.06 (0.99,1.12)	1.06 (0.97,1.13)

Table 1: Parameter estimation under the joint and SLS models.

5.3 Results: personalized optimal clinical decision estimation

We applied the proposed policy gradient method in Section 4 to the simulated dataset to estimate the personalized optimal clinical decision that maximizes one patient's graft median survival time, i.e., $R_i = \log(\hat{T}_i)$, where $\hat{T}_i$ is the median survival time of any patient n . The starting parameter values θ_0 in Algorithm 2 were set to be the estimated posterior means of these parameters from posterior inference, which can be considered as the estimates of how physicians treated patients in the simulated data. Therefore, the goal of the optimization procedure is to improve physicians' current treatment strategy in terms of prolonging patients' survival.

We implemented Algorithm 2 with $M = 1000$ steps to estimate the personalized optimal parameter $\tilde{\theta}_i$ for two randomly selected patients, denoted as S1 and S2. Patient S1 had a DGF of 0, donor age of 54.2 years, and BMI of 24, while patient S2 had a DGF of 1, donor age of 37.4 years, and BMI of 24.8. Figure 4(a, b) plots the expected mean reward versus SGD iterations. Note that the SGD procedure in the context of reinforcement learning usually has high variability, which is alleviated by our variance reduction method in Algorithm 2 as shown in Figure 4. For patient S1, the expected mean reward increases from an initial value of 7.65 to its maximum in the SGD, 7.69, which corresponds to a predictive median survival time of 2,209 days, a 4.6% increase from its initial value 2,111. For patient S2, the expected mean reward goes from an initial value of 7.69 to a maximum at 7.76. This corresponds to the predictive median survival time increasing from 2,203 days to 2,383 days, an 8.2% improvement.

To further interpret the estimated optimal “policy” parameters for patients S1 and S2, we compared the initial parameter values of the SGD-posterior means, with the optimized values by the SGD in Table 2. Recall that the dosage model is $d_{i,j} = (1, y_{i,j}, \mathbf{x}_i)\beta_d + \zeta_{i,j}$. Denote $\beta_d = (\beta_{d1}, \beta_{d2}, \dots, \beta_{dL})^T$, where L is the dimension of β_d . Since $\mathbf{x}_i$ denotes the baseline covariate and does not change over time, we define the personalized dosage intercept to be $\tilde{\beta}_d = (1, \mathbf{x}_i)(\beta_{d1}, \beta_{d3}, \dots, \beta_{dL})^T$ so that optimizing β_d is equivalent to optimizing $(\beta_{d2}, \tilde{\beta}_d)$. As shown in Table 2, the optimized dosage parameters $\tilde{\beta}_d$ and β_{d2} for patient S1 were lower than the estimated posterior means, indicating that patient S1 would benefit from a lower dosage for the same creatinine level compared to the observed dosages. In contrast, the optimal $\tilde{\beta}_d$ and β_{d2} were higher than the posterior means for patient S2, indicating the preference for higher dosages. The optimal dosage errors, σ_d^2 , for both patients were significantly lower than the initial value, indicating that a lower variance in the dosing procedure would benefit patient survival. The optimal baseline visitation intensity μ and the peak time parameter ν_1 were both roughly the same as their posterior means, indicating that the simulated follow-up schedules were close to optimal. However, the visitation intensity shape parameter ν_2 increased from 1.464 to 1.778 and 2.008 for patients S1 and S2 respectively and thus implies a higher intensity around the peak time ν_1 : intuitively, the optimized policy learns to be more certain about the “optimal peak time.”

In addition, to illustrate the advantage of optimizing both follow-up schedules and dosages, we compared our results to alternative strategies based on regular visits. As studied in Israni et al. (2014), during the first year post-transplant, patients were most frequently seen every 1 month or 3 months, depending on their physicians. After the




Figure 4: Panels (a, b) plot the expected mean reward versus SGD iterations for two randomly selected patients S1 and S2. Panels (c, d) plot the density of the predictive median survival times under our method and the three alternative strategies for patients S1 and S2.

first year, stable patients were most frequently referred back between 4–6 months but the follow-up frequency was reported to vary from 0–12 months. We considered three alternative follow-up strategies: recommend patients to follow up every 1 month, 3 months, and 6 months. The dosages at follow-up visitations were still optimized in the same way as the proposed joint model with the policy gradient method. Figure 4(c, d) show the density plots of 100 realizations of the predictive median survival times under our method and the three alternative strategies for patients S1 and S2. Comparing the predictive median survival times under the three regular visitation strategies,

	$\tilde{\theta}_0$	$\tilde{\theta}_{S1}$	$\tilde{\theta}_{S2}$
β_d : personalized dosage intercept	S1: 0.864, S2:0.987	0.746	1.316
β_{d2} : dosage effect of creatinine	0.200	0.153	0.307
σ_d^2 : dosage error	0.0940	0.0217	0.00252
μ : baseline visitation intensity	-4.781	-4.821	-4.785
ν_1 : visitation intensity peak	2.512	2.416	2.519
ν_2 : visitation intensity shape	1.464	1.778	2.008

Table 2: Simulation: Stochastic Gradient Descent Optimal Parameter Results.

we can see that more frequent visitations yield longer median survival times. The optimized visitation schedule under the proposed method outperforms the three alternative strategies although it yields a similar overall visitation frequency with the strategy of “regular visits every 3 months” (not shown), highlighting the importance of optimizing visitation schedules based on longitudinal clinical measurements to prolong patients’ survival.

6 Application: DIVAT data analysis

We extracted data from Nantes University Hospital Centers in the DIVAT cohort (www.divat.fr), yielding a total of $I = 947$ patients who received a first or second renal graft transplanted from a living or heart-beating deceased donor between 2000 and 2014. All patients in the dataset received an initial maintenance therapy with tacrolimus and did not experience graft failure or death during hospitalization. Immediately after transplantation, several covariates as risk factors for graft failure were collected: donor age (AgeD), recipient age (AgeR), delayed graft function (DGF) defined as the indicator of the use of dialysis within the first week of transplant (1=used dialysis, 0=didn’t use dialysis), diabetes history (Diab) with 1 indicating the patient has a history of diabetes and 0 otherwise, type of donor (Type), and body mass index (BMI). There were two types of donors: donation after brain death but with heart beating (Type=1) and donation by a living donor (Type=0). Supplementary Table S2 summarizes patients’ characteristics at baseline immediately after transplantation. For each patient, longitudinal data were collected from the date of transplantation until the graft failure or being censored. At each follow-up visitation, the creatinine level and tacrolimus dosage were recorded. The next follow-up visitation time was determined by the physician.

6.1 Experimental results: model fitting

We first applied the proposed Bayesian joint model to the DIVAT data with $\mathbf{x}_i = (\text{AgeD}_i, \text{AgeR}_i, \text{DGF}_i, \text{BMI}_i, \text{Diab}_i, \text{Type}_i)$. The hyperparameters were set to the same as in the simulation study. We ran a total of 20,000 MCMC iterations with an initial burn-in of 5,000 iterations, and a thinning factor of 50. The convergence was assessed using R package *coda* and the trace plots for randomly selected parameters were shown in Supplementary Figure S4, showing no issues of non-convergence.

We plot the estimated posterior means with 95% CIs for some selected parameters in the dosage, longitudinal, and survival submodels in Figure 5. Figure 5(a) plots posterior means of the linear coefficient β_d with respect to the creatinine level and baseline covariates in the dosage model. DGF was negatively associated with the dosage, indicating that patients who used dialysis within the first week of transplant were likely to be assigned lower dosage levels. In contrast, BMI was positively associated with the dosage since bodyweight-based dosing of tacrolimus is the standard care for patients after transplantation (Andrews et al., 2017). Diabetes history was positively associated with the dosage. While the effect of diabetes on tacrolimus was not well characterized in the literature, Mendonza et al. (2007) showed that the time to maximum concentration of tacrolimus in the pharmacokinetics study was significantly longer in diabetics versus nondiabetics. Furthermore, donor type also increased the dosage level, indicating that patients who received kidney from a non-living donor were more likely to be assigned higher dosages compared to that from a living donor.

Figure 5(b) plots the estimated posterior means with 95% CIs for the fixed-effects regression coefficients with respect to the most recent tacrolimus dosage and baseline covariates in the longitudinal model (3.4). The dosage, donor age, DGF, BMI, and donor type were positively associated with the creatinine level, which agreed with findings in the literature (Katari et al., 1997; Gerchman et al., 2009; Foucher et al., 2016). In contrast, the recipient age was negatively associated with the creatinine level, suggesting that younger patients tend to have lower creatinine levels (Maraghi et al., 2016). Diabetes history also decreased the creatinine level. Hjelmseth et al. (2010) showed that a low creatinine was associated with type 2 diabetes in a cross-sectional study. The estimated posterior means and 95% CIs for selected survival submodel parameters are plotted in Figure 5(c). The posterior mean of the parameter corresponding to the tacrolimus dosage was positive while that corresponding to the toxicity was negative, suggesting that a higher tacrolimus drug reduces the hazard but the accumulated toxicity increases the hazard. These results were consistent with findings in Randhawa et al. (1997) and Böttiger et al. (1999), who reported nephrotoxicity caused by long-term high dosages of tacrolimus.

6.2 Experimental results: personalized optimal clinical decision estimation

Next, we applied the proposed policy gradient method to estimate the personalized optimal clinical decision in terms of maximizing a patient's median survival time. We initialized the parameters in Algorithm 2 by setting θ_0 to be their posterior means. Algorithm 2 was implemented with $M = 1000$ steps to estimate $\hat{\theta}_i$ for two randomly selected patients, denoted as R1 and R2. Patient R1 at transplantation was 60 years old with a BMI of 17, no history of diabetes, no DGF, and received donation from a 61-year-old non-living donor. Patient R2 at transplantation was 28 years old with a BMI of 25.5, no history of diabetes, no DGF, and received a kidney from a living 29-year-old donor. Patient R1 had an observed survival time of 1,527 days, while patient R2 had a censored survival time of 4,487 days. Supplementary Figure S5 plots the predictive median survival times across SGD iterations for the two patients. Patient R1's




Figure 5: Estimated posterior means and 95% CIs for parameters in the dosage, longitudinal, and survival submodels. The dosages and longitudinal measurements are in log-scale. The squares represent posterior means.

predictive median survival time increased from 1,793 to 1,895 days at the maximum, a 5.7% improvement; while patient R2's predictive median survival time increased from 5,191 to 5,628, an 8.4% gain.

To further interpret the estimated optimal parameters in clinical decisions, we compared their initial values with the optimized values in Table 3. In practice, these optimized parameters can be interpreted to guide clinical decisions. Patient R1's optimal dosage parameters, $\tilde{\beta}_d$ and β_{d2} , were higher than their posterior means, suggesting that

assigning a higher dosage level compared to what the physician actually did for the same creatinine level would improve his/her survival outcome. Specifically, the recommended dosage for patient R1 based on the optimized $\hat{\beta}_d$ and β_{d2} is the baseline of 2.788 plus an additional 0.076 units of tacrolimus for each unit of creatinine recorded at the most recent followup visitation. On the other hand, patient R2's optimal dosage parameters were both lower than the initial values, so lower dosage levels are recommended. The optimal dosage errors, σ_d^2 , for both patients were significantly lower than the initial value, meaning that the optimized policy is more certain about its dosing decisions so the variance is lower than the observed data. Therefore, a decrease in the optimized dosage error term can be implemented in practice by closely adhering to the mean dosage amount suggested by the optimized dosage parameters. The optimal baseline visitation intensities μ for both patients were lower than the initial value, indicating that they should be instructed to visit less often without the knowledge of their creatinine measurements. Their optimized visitation intensity peak times were lower than the posterior mean, indicating that they should be scheduled more frequent follow-ups when their creatinine levels are high. Furthermore, the visitation intensity shapes were significantly higher than the initial value so the optimized policy is more certain about the optimal peak time for visitation schedules.

	θ_0	θ_{R1}	θ_{R2}
$\hat{\beta}_d$: personalized dosage intercept	R1:2.367, R2:2.363	2.788	2.161
β_{d2} : dosage effect of creatinine	-0.038	0.076	-0.065
σ_d^2 : dosage error	0.111	0.035	0.0024
μ : baseline visitation intensity	-4.197	-4.617	-4.322
ν_1 : visitation intensity peak	1.479	1.123	1.311
ν_2 : visitation intensity shape	0.258	0.864	1.261

Table 3: DIVAT data: optimal parameters estimated by the policy-optimizing method.

6.3 Ablation study: optimizing time or dosage or both

Moreover, to demonstrate the benefit of optimizing the follow-up visitation schedules and dosages together, we compared the predictive median survival times under the non-optimized initial policy (Non-Opt.) with three versions of optimized policies: 1) only visitation schedules are optimized (Opt. Visits); 2) only dosages are optimized (Opt. Dosage); and 3) both visitation schedules and dosages are optimized (Opt. Both). Specifically, Non-Opt. used the parameters estimated from the proposed Bayesian joint model, mimicking what physicians did as collected in the DIVAT dataset; Opt. Visits used the optimized parameters from the SGD in the visitation model (3.1) and the non-optimized parameters in the dosage model (3.2); Opt. Dosage used the optimized parameters from the SGD in the dosage model and the non-optimized parameters in the visitation model; Opt. Both is the fully optimized model obtained in Section 6.2. Figure 6 plots boxplots for 100 realizations of the predictive median survival times under each of the four policies. Both patients benefit more from optimizing dosages relative to optimizing visitation schedule. The visitation schedule optimization accounts for more

improvement in prolonging the survival for patient R1 compared to R2 because, as shown in Table 3, there was a larger difference between the optimal parameter values (μ and ν_1) in the visitation model and their initial values for patient R1. The optimized visitation schedule for both patients, as we have discussed in Section 6.2, suggested slightly fewer visits overall, but more frequent visits when their creatinine levels are high. Comparing the four policies, we can see that optimizing both visitation schedules and dosages is clearly beneficial to these patients, thus empirically strengthening the motivation of our work.




Figure 6: The boxplots of the predictive median survival times under different policies of visitation schedules and dosages for patients R1 and R2.

7 Conclusion

In this work, we developed a two-step Bayesian approach to optimize clinical decisions with timing. Firstly, we proposed a Bayesian joint model for clinical observations (e.g., longitudinal measurements and survival time) and clinical decisions (e.g., follow-up visitation schedules and dosage assignments). The model components are connected by sharing certain structures and parameters in order to capture the mutual influence between the clinical observations and decisions. Moreover, we proposed a policy gradient method that optimized the personalized clinical decision for better survival, while parameter uncertainties in the clinical observation model are considered. Through simulation studies, we demonstrated that the optimized clinical decision obtained from the proposed approach yields longer predictive median survival times compared to scheduling follow-up visitations on a regular basis that is commonly used in caring for patients with chronic conditions nowadays. The analysis of the DIVAT data yields meaningful and interpretable results, showing that the proposed method has the potential to assist physicians' decisions on personalized treatment. In addition, we have built an R package

doct so that users can apply the proposed method to datasets in a similar setup that involves longitudinal decision making and an objective reward to optimize.

There are several potential extensions. Firstly, we consider one longitudinal measurement in the longitudinal process of the joint model. There could be other time-varying measurements affecting the clinical decision and survival. In our kidney transplantation application, besides creatinine levels, there are other longitudinal measurements recorded such as proteinuria, which represents having protein in the urine and can be an early sign of kidney disease. The proposed method can be extended to incorporate other longitudinal measurements by replacing the model in (3.4) with a multivariate mixed effects model (Chi and Ibrahim, 2006). Secondly, the proposed Bayesian joint framework that models both clinical decisions and observations relies on certain parametric assumptions that are suitable for our kidney transplantation application. However, these assumptions (e.g., the proportional hazard assumption in the survival submodel) may not hold for other medical applications. It will be straightforward to extend the proposed Bayesian model to more flexible models such as replacing the current prior distributions with a Bayesian nonparametric prior (e.g., the Dirichlet process, Ferguson (1973)), or considering other nonparametric models (e.g., neural networks). Thirdly, the proposed MTPP assumes stationarity in clinical decisions since it yields the same distribution over visitation timing and dosages given the same history. Although the stationarity is reasonable in the kidney transplantation application, our approach can be easily extended to allow non-stationary clinical decisions, by explicitly incorporating a time-varying term (e.g., $\exp(ct)$ where $c \in \mathbb{R}$ is a coefficient parameter) in (3.1). In addition, the policy gradient method is used to optimize personalized clinical decisions in this paper. Other optimization methods such as natural evolution strategies (Wierstra et al., 2014) will be explored in the second step of the proposed approach. Lastly, patients with chronic conditions may take multiple medicines, e.g., mycophenolate mofetil (an immunosuppressive drug) and steroids along with tacrolimus in our kidney transplantation application. Modeling the effects of multiple types of drugs (and their interactions with clinical observations) and learning their optimal dosage-assigning policies in the proposed optimization method will be an interesting and challenging research topic.

Supplementary Material

Supplement for “Personalized Dynamic Treatment Regimes in Continuous Time: A Bayesian Approach for Optimizing Clinical Decisions with Timing”
(DOI: [10.1214/21-BA1276SUPP](https://doi.org/10.1214/21-BA1276SUPP); .pdf).

References

- Aalen, O., Borgan, O., and Gjessing, H. (2008). *Survival and event history analysis: a process point of view*. Springer Science & Business Media. MR2449233. doi: <https://doi.org/10.1007/978-0-387-68560-1>. 853, 855, 856
- Andrews, L. M., de Winter, B. C., Tang, J.-T., Shuker, N., Bouamar, R., van Schaik, R. H., Koch, B. C., van Gelder, T., and Hesselink, D. A. (2017). “Overweight kidney

- transplant recipients are at risk of being overdosed following standard bodyweight-based tacrolimus starting dose.” *Transplantation direct*, 3(2). 869
- Arshad, A., Anderson, B., and Sharif, A. (2019). “Comparison of organ donation and transplantation rates between opt-out and opt-in systems.” *Kidney International*, 95(6): 1453–1460. doi: <https://doi.org/10.1016/j.kint.2019.01.036>. 850
- Bicalho, P. R., Requião-moura, L. R., Arruda, É. F., Chinen, R., Mello, L., Bertocchi, A. P. F., Naka, E. L., Tonato, E. J., and Pacheco-silva, A. (2019). “Long-Term Outcomes among Kidney Transplant Recipients and after Graft Failure: A Single-Center Cohort Study in Brazil.” *BioMed research international*. doi: <https://doi.org/10.1155/2019/7105084>. 850
- Böttiger, Y., Brattström, C., Tydén, G., Säwe, J., and Groth, C. G. (1999). “Tacrolimus whole blood concentrations correlate closely to side-effects in renal transplant recipients.” *British Journal of Clinical Pharmacology*, 48(3): 445–448. 869
- Butts, C. T. and Marcum, C. S. (2017). “A Relational Event Approach to Modeling Behavioral Dynamics.” In *Group Processes*, 51–92. Cham: Springer. 853, 856
- Carpenter, S. M., Menictas, M., Nahum-Shani, I., Wetter, D. W., and Murphy, S. A. (2020). “Developments in Mobile Health Just-in-Time Adaptive Interventions for Addiction Science.” *Current Addiction Reports*, 1–11. 852
- Chakraborty, B. (2013). *Statistical methods for dynamic treatment regimes*. Springer. MR3112454. doi: <https://doi.org/10.1007/978-1-4614-7428-9>. 851
- Chi, Y.-Y. and Ibrahim, J. G. (2006). “Joint models for multivariate longitudinal and multivariate survival data.” *Biometrics*, 62(2): 432–445. MR2227491. doi: <https://doi.org/10.1111/j.1541-0420.2005.00448.x>. 873
- Clifton, J. and Laber, E. (2020). “Q-Learning: Theory and Applications.” *Annual Review of Statistics and its Application*, 7: 279–301. MR4104194. doi: <https://doi.org/10.1146/annurev-statistics-031219-041220>. 850, 852, 858
- Ferguson, T. S. (1973). “A Bayesian analysis of some nonparametric problems.” *The annals of statistics*, 209–230. MR0350949. 873
- Foucher, Y., Blanche, P., and Buron, F. (2016). “A joint model for longitudinal and time-to-event data to better assess the specific role of donor and recipient factors on long-term kidney transplantation outcomes.” *European journal of epidemiology*, 31(5): 469–479. 869
- Gerchman, F., Tong, J., Utzschneider, K. M., Zraika, S., Udayasankar, J., McNeely, M. J., Carr, D. B., Leonetti, D. L., Young, B. A., de Boer, I. H., et al. (2009). “Body mass index is associated with increased creatinine clearance by a mechanism independent of body fat distribution.” *The Journal of Clinical Endocrinology & Metabolism*, 94(10): 3781–3788. 869
- Giesecke, K., Kakavand, H., and Mousavi, M. (2011). “Exact simulation of point processes with stochastic intensities.” *Operations Research*, 59(5): 1233–1245. MR2864337. doi: <https://doi.org/10.1287/opre.1110.0962>. 862

- Greensmith, E., Bartlett, P. L., and Baxter, J. (2004). “Variance reduction techniques for gradient estimates in reinforcement learning.” *Journal of Machine Learning Research*, 5(Nov): 1471–1530. MR2248025. doi: <https://doi.org/10.1162/jmlr.2003.4.7-8.1471>. 862
- Guan, Q., Reich, B. J., Laber, E. B., and Bandyopadhyay, D. (2019). “Bayesian Non-parametric Policy Search with Application to Periodontal Recall Intervals.” *Journal of the American Statistical Association*, 1–13. MR4143450. doi: <https://doi.org/10.1080/01621459.2019.1660169>. 852
- Hawkes, A. G. (1971). “Spectra of Some Self-Exciting and Mutually Exciting Point Processes.” *Biometrika*. MR0278410. doi: <https://doi.org/10.1093/biomet/58.1.83>. 857
- Hawkes, A. G. (2018). “Hawkes processes and their applications to finance: a review.” *Quantitative Finance*, 18(2): 193–198. MR3750729. doi: <https://doi.org/10.1080/14697688.2017.1403131>. 853
- Hjeltneseth, J., Røislien, J., Nordstrand, N., Hofsø, D., Hager, H., and Hartmann, A. (2010). “Low serum creatinine is associated with type 2 diabetes in morbidly obese women and men: a cross-sectional study.” *BMC endocrine disorders*, 10(1): 6. 869
- Hua, W., Mei, H., Zohar, S., and Giral, M. (2021). “Supplement for “Personalized Dynamic Treatment Regimes in Continuous Time: A Bayesian Approach for Optimizing Clinical Decisions with Timing”.” *Bayesian Analysis*. doi: <https://doi.org/10.1214/21-BA1276SUPP>. 861
- Israni, A., Dean, C. E., Salkowski, N., Li, S., Ratner, L. E., Rabb, H., Powe, N. R., and Kim, S. J. (2014). “Variation in structure and delivery of care between kidney transplant centers in the United States.” *Transplantation*, 98(5): 520. 851, 863, 866
- Jarl, J., Desatnik, P., Hansson, U. P., Go, K., and Gerdtham, U.-g. (2018). “Do kidney transplantations save money? A study using a before – after design and multiple register-based data from Sweden.” *Clinical Kidney Journal*, 11(2): 283–288. 850
- Kasiske, B. L., Zeier, M. G., Chapman, J. R., Craig, J. C., Ekberg, H., Garvey, C. A., Green, M. D., Jha, V., Josephson, M. A., Kiberd, B. A., Kreis, H. A., McDonald, R. A., Newmann, J. M., Obrador, G. T., Vincenti, F. G., Cheung, M., Earley, A., Raman, G., Abariga, S., Wagner, M., and Balk, E. M. (2010). “KDIGO clinical practice guideline for the care of kidney transplant recipients: A summary.” *Kidney International*, 77(4): 299–311. doi: <https://doi.org/10.1038/ki.2009.377>. 850
- Katari, S., Magnone, M., Shapiro, R., Jordan, M., Scantlebury, V., Vivas, C., Gritsch, A., McCauley, J., Starzl, T., Demetris, A., et al. (1997). “Clinical features of acute reversible tacrolimus (FK 506) nephrotoxicity in kidney transplant recipients.” *Clinical transplantation*, 11(3): 237. 869
- Laber, E. B., Lizotte, D. J., Qian, M., Pelham, W. E., and Murphy, S. A. (2014). “Dynamic treatment regimes: Technical challenges and applications.” *Electronic journal of statistics*, 8(1): 1225. MR3263118. doi: <https://doi.org/10.1214/14-EJS920>. 851

- Lamb, K. E. and Lodhi, S. (2011). “Long-Term Renal Allograft Survival in the United States: A Critical Reappraisal.” *American Journal of Transplantation*, 11(3): 450–462. 850
- Lin, D. Y. and Wei, L.-J. (1989). “The robust inference for the Cox proportional hazards model.” *Journal of the American Statistical Association*, 84(408): 1074–1078. MR1134495. 859
- Liu, H., Carlson, N. E., Grunwald, G. K., and Polotsky, A. J. (2018). “Modeling associations between latent event processes governing time series of pulsing hormones.” *Biometrics*, 74(2): 714–724. MR3825358. doi: <https://doi.org/10.1111/biom.12790>. 853, 856
- Luckett, D. J., Laber, E. B., Kahkoska, A. R., Maahs, D. M., Mayer-Davis, E., and Kosorok, M. R. (2019). “Estimating dynamic treatment regimes in mobile health using v-learning.” *Journal of the American Statistical Association*, 1–34. MR4107673. doi: <https://doi.org/10.1080/01621459.2018.1537919>. 851
- Maraghi, E., Foroushani, A. R., Younespour, S., Rostami, Z., Einollahi, B., Eshraghian, M. R., Akhoond, M. R., and Mohammad, K. (2016). “Longitudinal Assessment of Serum Creatinine Levels on Graft Survival After Transplantation: Joint Modeling Approach.” *Nephro-urology Monthly*, 8(4): 1–5. 869
- Mendonza, A. E., Zahir, H., Gohh, R. Y., and Akhlaghi, F. (2007). “Tacrolimus in diabetic kidney transplant recipients: pharmacokinetics and application of a limited sampling strategy.” *Therapeutic drug monitoring*, 29(4): 391–398. 869
- Murphy, S. A., Arjas, E., Jennison, C., Dawid, A. P., Cox, D. R., Senn, S., Cowell, R. G., Didelez, V., Gill, R. D., Kadane, J. B., and Robins, J. M. (2003). “Optimal dynamic treatment regimes.” *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 65(2): 331–366. 851, 855
- Naesens, M., Kuypers, D. R., and Sarwal, M. (2009). “Calcineurin inhibitor nephrotoxicity.” *Clinical Journal of the American Society of Nephrology*, 4(2): 481–508. 851
- Nahum-Shani, I., Smith, S. N., Spring, B. J., Collins, L. M., Witkiewitz, K., Tewari, A., and Murphy, S. A. (2018). “Just-in-time adaptive interventions (JITAIs) in mobile health: key components and design principles for ongoing health behavior support.” *Annals of Behavioral Medicine*, 52(6): 446–462. 852
- Nie, X., Brunskill, E., and Wager, S. (2021). “Learning when-to-treat policies.” *Journal of the American Statistical Association*, 116: 392–409. MR4227702. doi: <https://doi.org/10.1080/01621459.2020.1831925>. 852
- Orellana, L., Rotnitzky, A., and Robins, J. M. (2010). “Dynamic regime marginal structural mean models for estimation of optimal dynamic treatment regimes, part I: main content.” *The international journal of biostatistics*, 6(2). MR2602551. doi: <https://doi.org/10.2202/1557-4679.1200>. 850
- Randhawa, P. S., Starzl, T. E., and Demetris, A. J. (1997). “Tacrolimus (FK506)-Associated Renal Pathology.” *Advances in Anatomic Pathology*, 4(4): 265. 860, 869

- Rizopoulos, D., Hatfield, L. A., Carlin, B. P., Johanna, J. M., Izopoulos, D. R., Atfield, L. A. H., Arlin, B. P. C., and Akkenberg, J. J. M. T. (2014). “Combining Dynamic Predictions From Joint Models for Longitudinal and Time-to-Event Data Using Bayesian Model Averaging for Longitudinal and Time-to-Event Data Using.” *Journal of the American Statistical Association*, 109(508): 1385–1397. [MR3293598](#). doi: <https://doi.org/10.1080/01621459.2014.931236>. 859
- Robbins, H. and Monro, S. (1951). “A Stochastic Approximation Method.” *The Annals of Mathematical Statistics*, 22(3): 400–407. [MR0042668](#). doi: <https://doi.org/10.1214/aoms/1177729586>. 861
- Robins, J. (1986). “A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect.” *Mathematical modelling*, 7(9-12): 1393–1512. [MR0877758](#). doi: [https://doi.org/10.1016/0270-0255\(86\)90088-6](https://doi.org/10.1016/0270-0255(86)90088-6). 850
- Robins, J., Orellana, L., and Rotnitzky, A. (2008). “Estimation and extrapolation of optimal treatment and testing strategies.” *Statistics in medicine*, 27(23): 4678–4721. [MR2528576](#). doi: <https://doi.org/10.1002/sim.3301>. 852
- Ruder, S. (2016). “An overview of gradient descent optimization algorithms.” *arXiv preprint arXiv:1609.04747*, 1–14. URL <http://arxiv.org/abs/1609.04747>. 856, 861
- Shardell, M. and Ferrucci, L. (2018). “Joint mixed-effects models for causal inference with longitudinal data.” *Statistics in medicine*, 37(5): 829–846. [MR3760452](#). doi: <https://doi.org/10.1002/sim.7567>. 851
- Shibue, R. and Komaki, F. (2020). “Deconvolution of calcium imaging data using marked point processes.” *PLoS computational biology*, 16(3): e1007650. doi: <https://doi.org/10.1371/journal.pcbi.1007650>. 857
- Staatz, C., Taylor, P., and Tett, S. (2001). “Low tacrolimus concentrations and increased risk of early acute rejection in adult renal transplantation.” *Nephrology Dialysis Transplantation*, 16(9): 1905–1909. 851
- Sutton, R. S. and Barto, A. G. (2018). *Reinforcement learning: An introduction*. [MR3889951](#). 853
- Tabibian, B., Upadhyay, U., De, A., Zarezade, A., Schölkopf, B., and Gomez-Rodriguez, M. (2019). “Enhancing human learning via spaced repetition optimization.” *Proceedings of the National Academy of Sciences*, 116(10): 3988–3993. [MR3923855](#). doi: <https://doi.org/10.1073/pnas.1815156116>. 853
- Watanabe, S. and Opper, M. (2010). “Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory.” *Journal of machine learning research*, 11(12). [MR2756194](#). 865
- Wierstra, D., Schaul, T., Glasmachers, T., Sun, Y., Peters, J., and Schmidhuber, J. (2014). “Natural evolution strategies.” *The Journal of Machine Learning Research*, 15(1): 949–980. [MR3195335](#). 873

- Williams, R. J. (1992). "Simple statistical gradient-following algorithms for connectionist reinforcement learning." *Machine learning*, 8(3-4): 229–256. 862
- Xu, Y., Müller, P., Wahed, A. S., and Thall, P. F. (2016). "Bayesian nonparametric estimation for dynamic treatment regimes with sequential transition times." *Journal of the American Statistical Association*, 111(515): 921–950. MR3561917. doi: <https://doi.org/10.1080/01621459.2015.1086353>. 851
- Xu, Y., Thall, P. F., Hua, W., and Andersson, B. S. (2019). "Bayesian non-parametric survival regression for optimizing precision dosing of intravenous busulfan in allogeneic stem cell transplantation." *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 68(3): 809–828. MR3937475. 859
- Zhao, Y., Zeng, D., Socinski, M. A., and Kosorok, M. R. (2011). "Reinforcement learning strategies for clinical trials in nonsmall cell lung cancer." *Biometrics*, 67(4): 1422–1433. MR2872393. doi: <https://doi.org/10.1111/j.1541-0420.2011.01572.x>. 850, 852
- Zhao, Y.-Q., Zeng, D., Laber, E. B., Song, R., Yuan, M., and Kosorok, M. R. (2015). "Doubly robust learning for estimating individualized treatment with censored data." *Biometrika*, 102(1): 151–168. MR3335102. doi: <https://doi.org/10.1093/biomet/asu050>. 850
- Zhu, Y. and Li, S. (2018). "One-dimensional Poisson marked point process model and Its Random Characteristic Analysis in Haze Weather." *IOP Conference Series: Earth and Environmental Science*, 199(5). 857