BAYESIAN METHODS FOR DECISION MAKING IN BIOMEDICAL APPLICATIONS

by

William Hua

A dissertation submitted to Johns Hopkins University in conformity with the requirements for the degree of Doctor of Philosophy

Baltimore, Maryland

October, 2020

Abstract

This dissertation focuses on developing Bayesian survival analysis methodology for optimizing decision making and treatment methods in a variety of biomedical applications.

First, we develop a flexible Bayesian nonparametric regression model based on a dependent Dirichlet process and Gaussian process, DDP-GP, for optimizing precision dosing of intravenous busulfan in allogenic stem cell transplantation. Our analyses of a dataset of 151 patients identified optimal therapeutic dosage intervals that maximizes patient survival outcomes and varies substantively with age and complete remission status. Extensive simulations to evaluate the DDP-GP model in similar settings showed that its performance compares favorably to alternative methods. We provide an R package, `DDPGPSurv`, that implements the DDP-GP model for a wide range of survival regression analyses.

The second main contribution of this dissertation is the development of personalized dynamic treatment regimes (DTR) in continuous time. 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 building a generative model for a sequence of medical interventions with a marked temporal point process (MTPP)

and embedding this into a Bayesian joint framework where the other components model longitudinal medical measurements and time-to-event data. Moreover, we propose a policy gradient method to learn the personalized optimal clinical decision that maximizes patient survival by interacting the MTPP with the model on clinical observations while accounting for uncertainties in clinical observations. A signature application is to schedule follow-up visitations and assign a dosage at each visitation for patients after kidney transplantation. We demonstrate that the personalized decisions made by our method are interpretable and help improve patient survival, and provide an R package, `doct`, that broadly implements our framework.

Lastly, we introduce a Bayesian semiparametric model for learning biomarker trajectories and change points in Alzheimer's disease (AD). Through simulation and real data studies, we show that our model is able to reliably detect a pre-diagnosis longitudinal change point, evaluate the probability of AD progression, and account for heterogeneity by clustering subjects using longitudinal and diagnosis time data.

Primary Reader: Yanxun Xu

Secondary Reader: Mei-Cheng Wang

Acknowledgments

During the past five years, I've met many people that were crucial to my development at Hopkins and the writing of this dissertation. One of the most difficult and important parts of a PhD is choosing a compatible advisor, and I'm grateful to have Yanxun Xu as mine. Not only is she a brilliant researcher, but she also has a welcoming personality that makes collaboration easy and fun. She has been an extraordinary research mentor that cared about me every step of the way, provided career advice, and allowed me to freely explore my career interests through internships. I first came to Hopkins with limited statistical knowledge but she has helped me grow both intellectually and personally.

I am thankful to all of the faculty and staff in the Applied Mathematics and Statistics department for creating an incredible learning and research environment that feels like one big family. Special thanks to Kristin Bechtel, Ann Gibbins, and Seanice Tucker for all your help over the years. I also want to express my gratitude to Steve Hanke and Zheyu Wang for their insights during my Graduate Board Oral Exam, and Avanti Athreya, Carey Pribe, and Mei-Cheng Wang for serving on my dissertation defense committee.

My time at Hopkins has allowed me to meet many fellow students that have made the past few years fun and enjoyable. In particular, I am thankful to Lingyao, Andrew, Kevin, Tu, Wayne, and Joseph for their companionship throughout this

journey. Lastly, I would like to thank my parents for their unconditional love and support and for cultivating my passion for mathematics from a young age.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Bayesian Statistics

In Bayesian statistics, we formulate our probabilistic beliefs using prior and posterior distributions. A prior distribution represents to our beliefs about a certain quantity before we consider information from the data. We then use the data to update our beliefs and produce a posterior distribution. In the context of statistical modeling, we are often concerned with the prior and posterior distributions of model parameters, θ, that model data D. Formally, we compute the posterior distribution of parameters using Bayes' rule:

$$\overbrace{p(\theta|D)}^{\text{Posterior}} = \frac{\overbrace{p(D \mid \theta)}^{\text{Likelihood}}\overbrace{p(\theta)}^{\text{Prior}}}{\underbrace{p(D)}_{\text{Integrating factor}}}. \tag{1.1}$$

Since the integrating factor is a constant, we can ignore it and formulate the

posterior by considering the proportional relationship

$$\overbrace{p(\theta|D)}^{\text{Posterior}} \propto \overbrace{p(D \mid \theta)}^{\text{Likelihood}} \overbrace{p(\theta)}^{\text{Prior}}. \tag{1.2}$$

For some cases with a carefully chosen prior distribution, the posterior belongs to the same family of distributions as the prior, and they are referred to as conjugate distributions. In complex models with many parameters, the posterior distribution of each parameter cannot be derived in closed form and instead we rely on alternative Bayesian inference sampling techniques to approximate the posterior distributions $p(\theta \mid D)$. The primary class of algorithms to solve this issue are Markov chain Monte Carlo (MCMC) methods, which establish a Markov chain that has the target posterior distribution as its equilibrium distribution. One of the most commonly used MCMC methods is the Metropolis-Hastings algorithm, described as follows for a model with I parameters $(\theta_1, \cdots, \theta_I)$ and N iterations:

In Algorithm 1, the computation for the acceptance probability can be simplified by carefully selecting the proposal distributions $q(\theta_i^* \mid \theta_i^{(n-1)})$. For example, if the full conditional distribution $p(\theta_i^* \mid \theta_{-i}, D)$ and the prior $p(\theta_i^*)$ are conjugate, then we can choose the proposal distribution to be the full conditional, $q(\theta_i^* \mid \theta_i^{(n-1)}) = p(\theta_i^* \mid \theta_{-i}, D)$. This special case of Metropolis-Hastings is called the Gibbs sampling algorithm and the acceptance probability always equal to one.

In some complicated models, the full conditional and prior distributions for certain parameters are not conjugate due to multiple types of data. For example, let the dataset be composed of both survival and longitudinal data, denoted D_s and D_l, respectively. Then, the acceptance probability in Algorithm 1 for a shared parameter,

Algorithm 1 Metropolis Hastings

The following illustrates the Metropolis Hastings algorithm to approximate the target posterior distributions for each parameter, $p(\theta_1 \mid D), \cdots, p(\theta_I \mid D)$. θ_{-i} represents the most recent values of all other parameters excluding θ_i.

Initialize $\theta_i^{(0)}$ for $i = 1, \cdots, I$ at arbitrary value, i.e. sample from priors.
for n:=1 to N **do**
  for i:=1 to I **do**
    Generate proposal value $\theta_i^* \sim q_i(\theta_i^* \mid \theta_i^{(n-1)})$
    $A(\theta_i^{(n-1)}, \theta_i^*) \leftarrow \frac{p(\theta_i^* | \theta_{-i}, D)}{p(\theta_i^{(n-1)} | \theta_{-i}, D)} \frac{q(\theta_i^{(n-1)} | \theta_i^*)}{q(\theta_i^* | \theta_i^{(n-1)})} = \frac{p(D | \theta_{-i}, \theta_i^*) p(\theta_i^*)}{p(D | \theta_{-i}, \theta_i^{(n-1)}) p(\theta_i^{(n-1)})} \frac{q(\theta_i^{(n-1)} | \theta_i^*)}{q(\theta_i^* | \theta_i^{(n-1)})}$
    $U_i \leftarrow \text{Uniform}(0, 1)$
    if $U_i < A(\theta_i^{(n-1)}, \theta_i^*)$ **then**
      $\theta_i^{(n)} \leftarrow \theta_i^*$
    else
      $\theta_i^{(n)} \leftarrow \theta_i^{(n-1)}$
    end if
  end for
end for

θ_i, that models both both data types is

$$\begin{aligned} A(\theta_i^{(n-1)}, \theta_i^*) &= \frac{p(\theta_i^* \mid \theta_{-i}, D_s, D_l)}{p(\theta_i^{(n-1)} \mid \theta_{-i}, D_s, D_l)} \frac{q(\theta_i^{(n-1)} \mid \theta_i^*)}{q(\theta_i^* \mid \theta_i^{(n-1)})} \\ &= \frac{p(D_s \mid \theta_{-i}, \theta_i^*) p(D_l \mid \theta_{-i}, \theta_i^*) p(\theta_i^*)}{p(D_s \mid \theta_{-i}, \theta_i^{(n-1)}) p(D_l \mid \theta_{-i}, \theta_i^{(n-1)}) p(\theta_i^{(n-1)})} \frac{q(\theta_i^{(n-1)} \mid \theta_i^*)}{q(\theta_i^* \mid \theta_i^{(n-1)})}. \end{aligned}$$

For these complex datasets, it's infeasible to construct a model such that the full conditional distribution, $p(\theta_i^* \mid \theta_{-i}, D_s, D_l)$, and prior, $p(\theta_i^*)$, are conjugate when considering both types of data. However, the conditional distribution for one data type, $p(\theta_i^* \mid \theta_{-i}, D_l)$, can be conjugate with the respective prior. This conjugacy allows for a prudently selected proposal distribution $q(\theta_i^* \mid \theta_i^{(n-1)}) = p(\theta_i^* \mid \theta_{-i}, D_l)$ to

significantly simplify the acceptance probability:

$$A(\theta_i^{(n-1)}, \theta_i^*) = \frac{p(D_s \mid \theta_{-i}, \theta_i^*)}{p(D_s \mid \theta_{-i}, \theta_i^{(n-1)})}.$$

This trick not only allows for easier computation, but also results in a higher acceptance probability and improved mixing of the MCMC chain.

1.2 Survival Analysis

In many medical applications, we are interested in studying how certain factors affect the time to a medical event such as disease progression, organ failure, or patient death. Such medical events are often denoted "survival events" and the time to a survival event is referred to as the "survival time" or "event time." Survival analysis is used to evaluate the relationship between patient characteristics and survival times. Let T represent a survival time of interest for a patient and $f(t)$ as its corresponding probability density function (pdf). Then the cumulative distribution function (cdf) is $F(t) = P(T \leq t) = \int_0^T f(u)du$. In survival analysis, we are often interested in modeling the survival function of T, which is the probability of not experiencing the survival event at time t:

$$S(t) = P(T > t) = 1 - F(t) = \int_T^\infty f(u)du \tag{1.3}$$

Another way to model the survival time T is to use a hazard function, which measures the instantaneous rate of the survival event, given that the survival event has not yet been observed. Using this definition, we can derive the relationship

between the hazard function and the density and survival functions:

$$h(t) = \lim_{dt \to 0} \frac{P(t < T \leq t + dt \mid T > t)}{dt} = \lim_{dt \to 0} \frac{P(t < T \leq t + dt)}{dt P(T > t)} = \frac{f(t)}{S(t)} \tag{1.4}$$

Survival analysis data can be incomplete when certain patients do not observe a survival event. For example, let us consider a study of cancer patients where the survival event is death. When the study concludes at time C, some patients may still be alive and their survival event is unobserved. We refer to this missing data phenomenon as right-censoring, and the time at which the study ends is the censor time. The survival data for each patient is composed of $Y = \min(C, T)$ and an indicator function for censoring, $\delta = I(T \leq C)$. If the patient survival event is censored, the only information known is that the survival event is larger than the censor time. Thus, we can write the likelihood of the survival data (Y, δ) as:

$$L = f(Y)^{\delta} S(Y)^{1-\delta} \tag{1.5}$$

Chapter 2

Bayesian Nonparametric Survival Regression for Optimizing Precision Dosing of Intravenous Busulfan in Allogeneic Stem Cell Transplantation

2.1 Introduction

Allogeneic stem cell transplantation (allo-SCT) is an established treatment for various hematologic diseases, including acute myelogenous and lymphocytic leukemia and non-Hodgkins lymphoma. Intravenous (IV) busulfan has been established as a desirable component of the preparative regimen for allo-SCT, due to its absolute bioavailability and dosing accuracy, leading to improved patient survival (Andersson

et al., 2002; Bredeson et al., 2013; Copelan et al., 2013; Nagler et al., 2013; Wachowiak et al., 2011). The patient's busulfan systemic exposure (Bu-SE) represented by the area under the plasma concentration versus time curve, AUC, is crucial, as serious adverse events are associated with an AUC that is either very high or too low. Higher AUC values are associated with neurologic toxicity (grand mal seizures), hepatic veno-occlusive disease, mucositis, and/or gastro-intestinal toxicity (Dix et al., 1996; Geddes et al., 2008; Kontoyiannis et al., 2001; Ljungman et al., 1997). Lower AUC is associated with an increased likelihood of disease recurrence and thus shorter survival time (Andersson et al., 2017; Bartelink et al., 2009; McCune et al., 2002; Russell et al., 2013; Slattery et al., 1997).

Consequently, it is important to define an optimal AUC interval of busulfan exposure that maximizes survival while minimizing risk. Studies of fixed-dose oral busulfan regimens suggest that inter-individual variations in Bu-SE exposure may be as high as 10- to 20-fold. In contrast, IV delivery of busulfan is more consistent and reliable for controlling delivered dose (Andersson et al., 2000), and thus is better suited for obtaining optimal busulfan AUC intervals. Andersson et al. (2002) showed that an optimal interval of IV Bu-SE had AUC values approximately 950 to 1520 μMol-min from one representative dose in a typical 16-dose treatment course, or a total course AUC of 15,200 – 24,400 μMol-min, yielding longer survival times and lower toxicity rates compared with values outside this interval. More recently, Bartelink et al. (2016) reported that, in children and young adults, the optimal daily AUC range in a prototype 4-day Bu-based regimen was 78-101 mg*h/L (corresponding to a total course AUC of about 19,100 – 21,200 μMol-min), regardless of the type or stage of underlying disease and whether the patients were in complete remission (CR) if they had an underlying malignancy.

In this chapter, we account for patient heterogeneity to assess the joint impact of patient age, CR status, and AUC on patient survival, with the goal to determine covariate-specific optimal daily AUC intervals. To evaluate possible interactions between covariates and Bu-SE, and their association with treatment outcome, we analyzed a dataset of 151 patients who underwent allo-SCT for acute myelogenous leukemia (AML) and myelodysplastic syndrome (MDS). It has been demonstrated that many different comorbidity conditions may affect the patient's risk for developing complications with these procedures (Sorror et al., 2005, 2014). Additionally, there commonly is a correlation between the severity of comorbidities and patient age. Therefore, we analyzed the outcome of our patients using age as a continuous covariate. We also included the indicator of whether the patient was in CR or had active disease at time of allo-SCT, since patients transplanted in CR have, on average, more favorable outcomes (De Lima et al., 2004; Kanakry et al., 2014). Our goal was to find patient-specific optimal AUC ranges that maximize expected survival time given the patient's age and CR status. The results of this analysis may provide specific guidelines for so-called "personalized" or "precision" medicine in clinical practice.

Andersson et al. (2002) estimated the optimal AUC range by fitting a Cox proportional hazards regression model for overall survival (OS) time and smoothing a martingale residual plot, which showed that the hazard of death was approximately a quadratic function of log(AUC). Bartelink et al. (2016) used a fourth-order polynomial model to estimate the association between AUC and OS. Both of these methods assumed specific parametric distributions for survival time, and the latter analysis assumed a specific polynomial function for the relationship between AUC and OS. For our data set, Figure 2.1 shows histograms and estimated density plots of the patients' OS times in weeks, with and without a log transformation of OS. The figure

clearly presents a long-tailed distribution that might result from a mixture of several unknown distributions. Alternatively, the long-tailed distribution might be due to the fact patients who have survived at least four years from transplant are at risk of death from natural causes, rather than leukemia or transplant related causes. Consequently, the specific models and strong parametric assumptions made by Andersson et al. (2002) and Bartelink et al. (2016) may not be suitable to fit the current data set well. In particular, the proportional hazards assumption underlying the Cox model may not be valid. Even given a survival distribution that fits the data reasonably well, an additional problem is determining functional relationships between AUC, prognostic covariates, and the risk of death.

We present a flexible Bayesian nonparametric (BNP) survival regression model to estimate the relationship between survival time, AUC, and baseline covariates. Based on our analysis of the allo-SCT dataset, we determined personalized optimal AUC ranges based on patient' age and CR status. An important advantage of BNP models is that they often fit complicated data structures better than parametric model-based methods because BNP models can accurately approximate essentially any distribution or function, a property known as "full support." Another important advantage of BNP models is that they often identify unexpected structures in a dataset that cannot be seen using conventional statistical models and methods. BNP models have been used widely for survival analysis. Hanson and Johnson (2002) proposed a mixture of Polya tree priors in semiparametric accelerated failure time (AFT) models, while Gelfand and Kottas (2003) developed the corresponding Dirichlet process (DP) mixture approach. Zhou and Hanson (2018) presented a unified approach for modeling survival data by exploiting and extending the three most commonly-used semiparametric models: proportional hazards, proportional odds, and accelerated

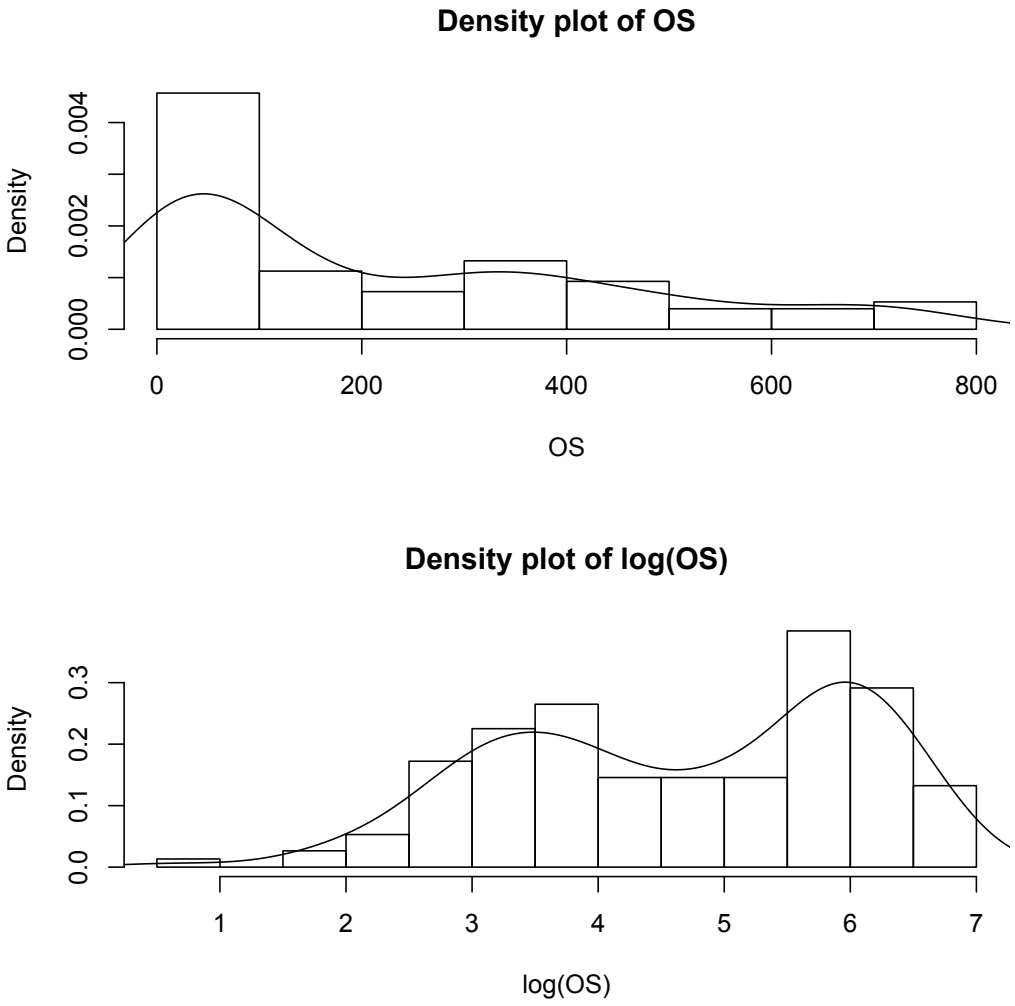

Figure 2.1: Histograms of overall survival time in weeks (top) and log overall survival time (bottom), with nonparametric density estimates.

failure time. Despite the flexibility of these approaches for modeling baseline survival distributions, they are restricted in the way that covariates may affect the baseline distribution. Fully nonparametric tree-based survival models have been developed, such as the use of random forests (Ishwaran et al., 2008) and Bayesian additive regression trees (Sparapani et al., 2016). De Iorio et al. (2009) proposed an unconstrained survival regression model with a dependent Dirichlet process (DDP) (MacEachern, 1999) prior in order to incorporate covariates in a naturally interpretable way. Xu et al. (2016) developed a DDP prior with Gaussian process as the base measure, the DDP-GP model, to evaluate overall survival (OS) times of complex dynamic treatment regimes including multiple transition times. However, a non-trivial limitation of the DDP-GP model in Xu et al. (2016) is that it gives the same fixed weight to all covariates, regardless of their numerical domains, when quantifying dependence between patients. Such a restriction may yield a sub-optimal posterior contraction rate, causing less accurate estimations (van der Vaart and van Zanten, 2011).

Building on the work in Xu et al. (2016), in this chapter we propose a flexible survival regression framework by formulating a DDP with a more general covariance function for the GP prior that includes an individual scale parameter for each covariate and additional hyperparameters for model flexibility and robustness. The proposed model provides easy-to-implement posterior inferences in settings where the proportional hazards assumption, specific parametric models such as AFT models, or semi-parametric models may not fit the data well. Currently, the R package `survival`, which is limited to such models, remains a standard tool for statistical and medical researchers. One of the main contributions of this chapter is to provide a new, easy-to-use R package, `DDPGPSurv`, that implements the proposed DDP-GP survival regression model for a broad range of survival analyses. A major goal is

that `DDPGPSurv` will become a new standard computational tool for implementing this generalized DDP-GP to conduct survival analysis in medical research.

The rest of the chapter is organized as follows. In Section 2.2 we review the motivating dataset. We present the DDP-GP survival regression model in Section 2.3. Section 2.4 gives a brief introduction to the R package `DDPGPSurv`. Extensive simulation studies with comparison to alternative methods are conducted in Section 2.5. We analyze our dataset in Section 6, and conclude with a brief discussion in Section 7.

2.2 Motivating Study

When total body irradiation was replaced with high-dose oral busulfan (Santos et al., 1983; Tutschka et al., 1987), it quickly became clear that unpredictable, often lethal, toxicities limited the use of a busulfan-based conditioning program. Several retrospective studies indicated an association between systemic drug exposure and clinical treatment outcome (Dix et al., 1996; Slattery et al., 1997). This spawned an interest in exploring pharmacokinetic dose guidance, but the erratic bioavailability of oral busulfan prevented its successful implementation in a prospective fashion. The advent of IV Busulfan, which guarantees complete bioavailability with absolute assurance for systemic dose delivery has changed this. Routine application of therapeutic dose guidance for IV busulfan in pre-transplant conditioning therapy now makes it possible to accurately deliver a predetermined systemic exposure dose in terms of AUC, thereby optimizing treatment. This is important, since in (myeloid) leukemia the cytotoxic drug dose and accurate dose delivery are associated with clinical treatment outcome (Andersson et al., 2002, 2017; Bartelink et al., 2016; De Lima

	Patients($n = 151$)
Age(years)	
≤ 25	12 (8%)
$26 - 35$	22 (15%)
$36 - 45$	32 (21%)
$46 - 55$	50 (33%)
≥ 56	35 (23%)
Sex	
Male	77 (51%)
Female	74 (49%)
In CR at transplantation	
Yes	80 (53%)
No	71 (47%)
AUC quantile	
10%	3,928
25%	4,328
50%	5,077
75%	5,754
90%	6,371

Table 2.1: Patient characteristics at time of transplantation.

et al., 2004; Russell et al., 2013). A question that has not yet been resolved satisfactorily is what optimal systemic exposure dose to target in an individual patient. To address this decisively, we have retrieved data in The University of Texas MD Anderson Cancer Center from 151 AML/MDS patients who received a standardized 4-day fludarabine-IV busulfan combination, with both agents administered based on body surface area. Pharmacological studies of busulfan were performed as an optional procedure, but the information was not used for busulfan dose-adjustments. The dataset includes overall survival (OS) times and the covariates age, CR status, and AUC. Table 1 summarizes the characteristics of the study population at baseline.

2.3 Probability Model

2.3.1 Dependent Dirichlet process-Gaussian process prior

Denote the log time to death by Y and censoring time on the log(time) domain by C, with $T = Y \wedge C$ the observed log time of the event or censoring, and $\delta = I(Y \leq C)$. Indexing patients by $i = 1, \cdots, n$, the observed outcome data for patient i are (T_i, δ_i), and we let $\boldsymbol{x}_i$ denote the baseline covariate vector, including age, CR status, and AUC.

We construct a Bayesian nonparametric (BNP) survival regression model for $F(Y \mid \boldsymbol{x})$, the distribution of $[Y \mid \boldsymbol{x}]$, as follows. We start with a model for a discrete random distribution $G(\cdot)$, then use a Gaussian kernel to extend this to a prior for a continuous random distribution, and finally we replace the kernel means by a regression structure to define the desired prior on $\{F(Y \mid \boldsymbol{x}),\ \boldsymbol{x} \in X\}$. The constructions of $G(\cdot)$ and $F(\cdot)$ are elaborated below, by way of a brief review of BNP models. See, for example, Müller and Mitra (2013) and Müller and Rodriguez (2013) for more extensive reviews.

First proposed by Ferguson (1973), the Dirichlet process (DP) prior has been used widely in Bayesian analyses as a prior model for random unknown probability distributions. A DP(α_0, G_0) involves a positive scaling parameter α_0 and a base probability measure G_0. A constructive definition of a DP is provided by Sethuraman (1994), the so-called "stick-breaking" construction, given by $G = \sum_{h=1}^{\infty} w_h \delta_{\theta_h}$, where $\theta_h \sim G_0$, and $w_h = v_h \prod_{l<h}(1 - v_l)$ with $v_h \sim \text{be}(1, \alpha)$. Here, $\delta_{\theta_h}(\cdot)$ denotes the Dirac delta function, which is equal to 1 at θ_h and is equal to 0 everywhere else. In many applications, the discrete nature of G is not appropriate. To deal with this, a DP mixture model extends the DP model by replacing each point mass $\delta(\theta_h)$ with

a continuous kernel, such as a DP mixture of normals $G = \sum_h w_h \; N(\theta_h, \sigma^2)$, where $N(\theta_h, \sigma^2)$ denotes a normal distribution with mean θ_h and standard deviation σ.

To include regression on covariates, MacEachern (1999), extended the DP mixture model by replacing each mean parameter θ_h in the sum with a function $\theta_h(\boldsymbol{x})$ of covariates $\boldsymbol{x}$. This is called a dependent Dirichlet process (DDP), obtained by assuming the regression model

$$F(y \mid \boldsymbol{x}) = \sum_{h=1}^{\infty} w_h \; N(y; \theta_h(\boldsymbol{x}), \sigma^2),$$

where one can specify a stochastic process prior for $\{\theta_h(\boldsymbol{x})\}$. As a default assumption MacEachern (1999) proposed a Gaussian process (GP) prior. Here, the GP is indexed by $\boldsymbol{x}$. Temporarily suppressing the subindex $_h$, a GP prior is characterized by the marginal distribution for any n-tuple $(\theta(\boldsymbol{x}_1), \ldots, \theta(\boldsymbol{x}_n))$ being a multivariate normal distribution with mean vector $(\mu(\boldsymbol{x}_1), \ldots, \mu(\boldsymbol{x}_n))$ and $(n \times n)$ covariance matrix with (i, j) element $C(\boldsymbol{x}_i, \boldsymbol{x}_j)$, for any set of $n \geq 1$ covariate vectors $\boldsymbol{x}_1, \cdots, \boldsymbol{x}_n$. We denote this model by $\theta(\boldsymbol{x}) \sim \text{GP}(\mu, C)$. Extensive reviews of the GP are given by MacKay (1999) and Rasmussen and Williams (2006).

In the context of modeling each patient's transition times between successive disease states in a dataset arising from multi-stage chemotherapy of acute leukemia, Xu et al. (2016) modeled $\{\theta_h(\boldsymbol{x})\} \sim \text{GP}(\mu_h(\cdot), C(\cdot, \cdot))$, $h = 1, 2, \ldots$ with $\mu_h(\boldsymbol{x}_i; \boldsymbol{\beta}_h) = \boldsymbol{x}_i \boldsymbol{\beta}_h$ and $C(\boldsymbol{x}_i, \boldsymbol{x}_\ell) = \exp\{-\sum_{d=1}^{D}(x_{id} - x_{\ell d})^2\} + \delta_{i\ell} J^2$. Here, D is the dimension of the covariate vector, with $\delta_{i\ell} = I(i = \ell) = 1$ if $i = \ell$ and 0 otherwise. The term J^2 is jitter added to provide numerical stability by avoiding singular covariance matrices, with a small value such as $J = 0.1$ typically used.

A non-trivial limitation of this covariance function is that it gives the same weight to all covariates, regardless of their numerical domains, when quantifying dependence

between patients. This implies that different covariates x_d and $x_{d'}$ contribute the same to the correlation between patients i and l as long as $(x_{id} - x_{\ell d})^2$ and $(x_{id'} - x_{\ell d'})^2$ are the same. Furthermore, without including hyperparameters with prior distributions in the covariance function, the posterior inference using a Gaussian process prior may yield a posterior contraction rate that is sub-optimal (van der Vaart and van Zanten, 2011). To avoid these limitations, we extend the DDP-GP model by including an additional scale parameter, λ_d, for each covariate x_d and also an overall multiplicative scale parameter σ_0^2 in the covariance function:

$$C(\boldsymbol{x}_i, \boldsymbol{x}_\ell) = \sigma_0^2 \ \exp\left\{ -\sum_{d=1}^{D} \frac{(x_{id} - x_{\ell d})^2}{\lambda_d^2} \right\} + \delta_{i\ell} J^2. \tag{2.1}$$

The multiplicative scale parameter σ_0^2 accounts for variability in the data that is not accommodated by the variance σ^2 of the normal component distributions.

The model can be summarized as

$$\begin{aligned} p(y_i \mid \boldsymbol{x}_i, F) &= F_{\boldsymbol{x}_i}(y_i) \\ \{F_{\boldsymbol{x}}\} &\sim \text{DDP-GP}\left\{\{\mu_h\}, C; \alpha, \{\boldsymbol{\beta}_h\}, \{\lambda_d^2\}, \ \sigma_0^2, \ \sigma^2\right\}. \end{aligned} \tag{2.2}$$

We use the acronym DDP-GP to refer to the proposed model with the DDP mixture of normals with this particular GP prior on the mean of the normal kernel. Thus,

$$F_{\boldsymbol{x}}(y) = \sum_{h=0}^{\infty} w_h \, \mathbb{N}(y; \ \theta_h(\boldsymbol{x}), \ \sigma^2) \text{ with } \{\theta_h(\boldsymbol{x})\} \sim \ \text{GP}(\mu_h(\cdot), C(\cdot, \cdot)), \tag{2.3}$$

where $h = 1, 2, \ldots$, $\mu_h(\boldsymbol{x}_i) = \boldsymbol{x}_i \boldsymbol{\beta}_h$, and $C(\cdot, \cdot)$ is defined in (2.1).

2.3.2 DDP-GP survival regression model

Denote the vector of all model parameters by Θ and the data by $\mathcal{D}_n = \{T_i, \delta_i, \boldsymbol{x}_i\}_{i=1}^n$. The likelihood function is the usual form

$$L(\Theta \mid \mathcal{D}_n) = \prod_{i=1}^{n} f_{\boldsymbol{x}_i}(t_i \mid \Theta)^{\delta_i} \{1 - F_{\boldsymbol{x}_i}(t_i \mid \Theta)\}^{1-\delta_i},$$

where $f_{\boldsymbol{x}}(\cdot)$ and $F_{\boldsymbol{x}}(\cdot)$ denote the density and cumulative distribution function of Y for an individual with covariates $\boldsymbol{x}$. Given the assumed DDP-GP prior on $F_{\boldsymbol{x}}(\cdot)$, shown in (2.2), we complete the model by assuming the priors $\boldsymbol{\beta}_h \sim N(\boldsymbol{\beta}_0, \Sigma_0)$, $1/\sigma^2 \sim \text{Gamma}(a_1, b_1)$, the precision parameter $\alpha \sim \text{Gamma}(a_2, b_2)$, $\sigma_0 \sim N(0, \tau_\sigma^2)$, and the covariate scale parameters $\lambda_d \sim \text{iid } N(0, \tau^2)$, $d = 1, \cdots, D$. Thus, the DDP-GP's hyperparameters are $\boldsymbol{\theta}^* = (\boldsymbol{\beta}_0, \Sigma_0, a_1, b_1, a_2, b_2, , \tau_\sigma^2, \tau^2)$.

To implement the DDP-GP model, one first must determine numerical values for the hyperparameters $\boldsymbol{\theta}^*$. We introduce default choices for fixing $\boldsymbol{\theta}^*$ in our `DDPGPSurv` package below, although users can define their own preferred values, if desired. We suggest using an empirical Bayes method to obtain $\boldsymbol{\beta}_0$ by fitting a normal distribution for patient response on the log scale, $\log(Y) \mid \boldsymbol{x} \sim N(\boldsymbol{x}\boldsymbol{\beta}_0, \widehat{\sigma}^2)$ and assuming Σ_0 to be a diagonal matrix with all diagonal values 10. Once an empirical estimate $\widehat{\sigma}^2$ of σ^2 is obtained, one can tune (a_1, b_1) so that the prior mean of σ^2 matches the empirical estimate and the variance equals 10 or a suitably large value to ensure a vague prior. The total mass parameter α in the stick-breaking construction determines the number of unique clusters in the underlying DP Polya urn scheme. Usually, the DP yields many small clusters, therefore changing the prior of α does not significantly alter posterior predictive inference, which we will use for estimating the survival function and optimal AUC ranges. We assume $a_2 = b_2 = 1$ to ensure a vague prior on α.

Lastly, we assume $\tau = \tau_\sigma = 10$ so that the ranges of λ_d's and σ_0 in the covariance function are large enough to cover variability in the data.

To obtain posterior inference for a DDP-GP survival regression model, we first marginalize (2.2) analytically with respect to the random probability measures $F_{\boldsymbol{x}}(\cdot)$. To do this, we first rewrite (2.3) equivalently as a hierarchical model with a set of new latent indicator variables γ_i as

$$(Y_i \mid \gamma_i = h, \boldsymbol{x}_i) \sim N(\theta_h(\boldsymbol{x}_i), \sigma^2) \text{ and } p(\gamma_i = h) = w_h, \tag{2.4}$$

for $i = 1, \cdots, n$. If clusters of patients are defined as $S_h = \{i : \widetilde{\theta}_i = \theta_h\}$, then the γ_i's are interpreted as cluster membership indicators. Posterior simulation makes use of these indicators and the vectors $\boldsymbol{\theta}_h = (\theta_h(\boldsymbol{x}_1), \ldots, \theta_h(\boldsymbol{x}_n))$. After marginalization with respect to $F_{\boldsymbol{x}}$, we are left with the marginal model for $\{\gamma_i,\ \theta_h(\boldsymbol{x}_i);\ i = 1, \ldots, n,\ h = 1, \ldots\}$. We implement posterior sampling based on the collapsed Gibbs sampler (Escobar and West, 1995) in the R package `DDPGPSurv`. Details of the MCMC computations are provided in Section 2.8.

2.3.3 Personalized optimal AUC range estimation

Let $\rho_n = (S_1, \ldots, S_H)$ denote the partition of the n patients, determined by the clusters induced by the γ_i's. A key advantage of the proposed BNP model is that we can easily write down the posterior predictive distribution of the outcome Y_{n+1} for a

future patient with covariate vector $\boldsymbol{x}_{n+1}$, given by

$$p(Y_{n+1} \mid \boldsymbol{x}_{n+1}, \mathcal{D}_n) = \sum_{\rho_n} p(\rho_n \mid \mathcal{D}_n) \int p(\Theta \mid \rho_n, \mathcal{D}_n)$$
$$\times \Big\{ \sum_{h=1}^{H+1} p(Y_{n+1} \mid n+1 \in S_h, \theta_h(\boldsymbol{x}_{n+1}), \Theta) p(n+1 \in S_h \mid \boldsymbol{x}_{n+1}, \mathcal{D}_n, \rho_n, \Theta) \Big\} d\Theta. \tag{2.5}$$

The innermost sum averages with respect to the cluster membership for the $(n+1)^{st}$ patient during the MCMC. The term $h = H+1$ corresponds to the case that this new patient may form his/her own singleton cluster. The posterior average with respect to $p(\rho_n \mid \mathcal{D}_n)$ and $p(\Theta \mid \rho_n, \mathcal{D}_n)$ is evaluated as an average over the MCMC sample.

For the IV busulfan allo-SCT data, $\boldsymbol{x}$ includes the key treatment variable AUC, which quantifies the patient's delivered dose of IV busulfan and thus may be targeted by the treating physician. From (2.5), based on our analysis of the IV busulfan data using the DDP-GP, we can use the predictive distribution to compute the optimal AUC for the future patient $n+1$ as that which maximizes expected log survival time,

$$\widehat{\mathrm{AUC}}_{n+1} = \mathrm{argmax}_{\mathrm{AUC}} E(Y_{n+1} \mid \boldsymbol{x}_{n+1}, \mathcal{D}_n), \tag{2.6}$$

where $\boldsymbol{x}_{n+1}$ includes patients' age, CR status, and AUC. The laboratory-based method for determining the median specific daily Bu-SE has about a 3% error when sampling is carried over 12-14 hours (or about 3-4 drug half-lives). However, if sampling is restricted to 4-6 hours (1.0 - 1.5 drug half-lives), as is done with many PK evaluation methods, the error increases to at least 6%. Based on these considerations, we decided to use optimal AUC +/- 10% as an reasonable interval for targeting, since it is not possible to detect any difference in covariate impact on outcome between patients

with AUC values falling within this narrow Bu-SE interval. Therefore, we define the optimal AUC interval for future patient $n+1$ as

$$\left[\ 0.9 \times \widehat{\mathrm{AUC}}_{n+1},\quad 1.1 \times \widehat{\mathrm{AUC}}_{n+1}\ \right],$$

bearing in mind that $\widehat{\mathrm{AUC}}_{n+1}$ depends on the patient's covariates $\boldsymbol{x}_{n+1}$.

2.4 R package: DDPGPSurv

One of the main contributions of this chapter is that we have developed an R package, `DDPGPSurv`, that implements the proposed DDP-GP model as a general tool for survival analysis. The functions in the package perform inference via MCMC simulations from the posterior distributions based on a DDP-GP prior using a collapsed Gibbs sampler (`mcmc_DDPGP`). The outputs from `mcmc_DDPGP` are then used as the inputs for the other functions to evaluate and plot the estimated posterior predicted density, survival, and hazard functions for new observations/patients. The package also includes a function for evaluating posterior mean survival for specified values of the covariate vector. For example, this allows the user to determine the optimal value of a specific covariate (with the other covariates fixed) that maximizes posterior expected survival time. The R package `DDPGPSurv` can be downloaded from `https://cran.r-project.org/web/packages/DDPGPSurv/index.html`.

Current standard methods for survival analysis typically involve the Kaplan-Meier (KM) estimator for unadjusted survival times with independent right censoring, accelerated failure time (AFT) models, or the Cox proportional hazards (PH) model. The KM estimator is a non-parametric statistic and is constructed using a finite number of conditional probabilities of survival at successive time intervals. To analyze

the effects of specific covariates using the KM estimator, the most common approach is simply to compute the KM for particular patient subsets that may be defined from $\boldsymbol{x}$, which reduces reliability. AFT regression models are fully parametric, which may be problematic if the baseline hazard function does not fit the specified AFT distribution. Comparisons between the DDP-GP and AFT models via simulations show that the DDP-GP is more robust, with much more accurate predictions across a range of various distributions (Weibull, lognormal, exponential). That is, if the distribution selected for the AFT model does not match the truth, the predictions will be inaccurate. The Cox model, which is semi-parametric, relies on the PH assumption, which states that the each covariate has a constant effect on the hazard function that does not vary over time. This assumption may not always be true, and it is not required by the DDP-GP model. Additionally, as with any BNP model, the DDP-GP accommodates irregularly shaped survival distributions, for example having multiple modes. Thus, the DDP-GP, implemented by the `DDPGPSurv` package, provides many advantages over these conventional methods, including robustness and accuracy across a wide range of possible distributions.

2.5 Simulation Studies

We conducted simulation studies to evaluate the DDP-GP model in terms of estimation of survival densities and optimal personalized AUC ranges, with the data simulated to mimic the structure of the allo-SCT dataset. We generated T=survival time, the covariates x_1 = age, x_2 = AUC, and x_3 = CR status for each patient, as follows. Let $LN(m, s)$ denote a log normal distribution with location and scale parameters m and s, and let $\boldsymbol{x}_i = (x_{i,1}, x_{i,2}, x_{i,3})$ denote the covariates for patient i.

Patients' ages and AUC values were sampled with replacement from the actual ages and AUC values in the allo-SCT dataset. We generated patients' CR statuses as independently and identically distributed (i.i.d.) binary variables from a Bernoulli(0.5). We simulated the Y_i's from a lognormal distribution, $Y_i \mid \boldsymbol{x}_i \sim LN(\mu(\boldsymbol{x}_i), \sigma_0^2)$, where the location parameter is the following function of $\boldsymbol{x}_i$,

$$\mu(\boldsymbol{x}_i) = 4 - 0.1x_{i,1} + 0.7x_{i,2} + 0.3x_{i,3} - 0.07x_{i,2}^2 - 0.1x_{i,1}x_{i,2} + 0.2x_{i,2}x_{i,3} - 0.18x_{i,1}x_{i,2}x_{i,3},$$

for $i = 1, \cdots, n$, and $\sigma_0 = 0.4$. We deliberately designed the form of $\mu(\boldsymbol{x}_i)$ based on clinical knowledge, including a quadratic term for AUC to reflect the fact that a Bu-SE that is either too high or too low is associated with shorter survival time. We also included interaction terms between AUC and covariates so that the relationship between survival and AUC may vary depending on each patient's age and CR status.

We considered two scenarios, one with $n = 200$ observations without censoring and the other with $n = 200$ and 25% censoring. For each scenario, $B = 100$ trials were simulated, and the proposed DDP-GP survival regression model was fit to each simulated dataset. The MCMC sampler was implemented for posterior inference and run for 5,000 iterations with an initial burn-in of 2,000 iterations, thinned by 10. We used the R package *coda* (Plummer et al., 2006) to check convergence, and both traceplots and autocorrelation plots (not shown) to check mixing of the Markov chain, which showed no convergence problems.

2.5.1 Survival density estimation

For simulated trials indexed by $b = 1, \cdots, B$, let $\overline{S}_b(t \mid \boldsymbol{x}) = p(Y_{n+1} \geq t \mid \boldsymbol{x}_{n+1} = \boldsymbol{x}, \mathcal{D}_n)$ denote the posterior predicted probability that a future patient $n + 1$ with

covariate $\boldsymbol{x}$ in trial b survives beyond time t. To estimate $\overline{S}_b(t \mid \boldsymbol{x})$ using our package `DDPGPSurv`, we first run the MCMC using the function `mcmc_DDPGP`. Then, the output from the function `mcmc_DDPGP` serves as the input to the function `DDPGP_Surv`, which returns the mean and 95% credible intervals for the survival function across the saved MCMC posterior samples. Using the empirical covariate distribution $\frac{1}{n}\sum_{i=1}^{n}\delta_{\boldsymbol{x}_i}$ to marginalize w.r.t. $\boldsymbol{x}_{n+1}$ and averaging across simulations, we get

$$\overline{S}(t) = \frac{1}{B}\sum_{b=1}^{B}\frac{1}{n}\sum_{i=1}^{n}\overline{S}_b(t \mid \boldsymbol{x}_i).$$

For comparators, we considered six alternative methods. First, we fit AFT regression models using either lognormal or Weibull distributions by assuming

$$\log(Y_i) = \beta_0 + \beta_1 x_{i,1} + \beta_2 x_{i,2} + \beta_3 x_{i,3} + \beta_4 x_{i,2}^2 + \beta_5 x_{i,1}x_{i,2} + \beta_6 x_{i,2}x_{i,3} + \beta_7 x_{i,1}x_{i,3} + \sigma\epsilon_i. \tag{2.7}$$

Here, assuming a normal distribution on ϵ_i implies that Y_i follows a lognormal distribution, while the extreme value distribution assumption on ϵ_i implies that Y_i follows a Weibull distribution. We also considered two flexible semiparametric survival methods that model the baseline survival using a Polya Trees (PT) prior (Hanson and Johnson, 2002) or a transformed Bernstein polynomials (TBP) prior (Zhou and Hanson, 2018), respectively. Both models were implemented in the R package `spBayesSurv` (Zhou et al., 2018) . We assumed the AFT regression model as the frailty model in both the PT and TBP methods with the same setup as in (2.7). Lastly, we compared the proposed DDP-GP model to two fully nonparametric survival models using random forests (RF) (Ishwaran et al., 2008) and Bayesian additive regression trees (BART) (Sparapani et al., 2016). We used the R packages `randomForestSRC` (Ish-

waran and Kogalur, 2007) and `BART` (McCulloch et al., 2018) to implement the RF method and the BART method, respectively.

Figure 2.2 compares $\overline{S}(\cdot)$ estimated under the DDP-GP model to the simulation truth,

$$S_0(t) = \frac{1}{B}\sum_{b=1}^{B}\frac{1}{n}\sum_{i=1}^{n} S_0(t \mid \boldsymbol{x}_i),$$

the maximum likelihood estimates (MLEs) obtained under each of the two AFT models, and the estimated survival curves under the PT, TBP, BF, and BART methods. In each scenario, the true curve is given as a solid black solid line and the posterior mean survival function under the DDP-GP model as a solid grat line with point-wise 95% posterior credible bands as two dotted gray lines. In both scenarios, the DDP-GP

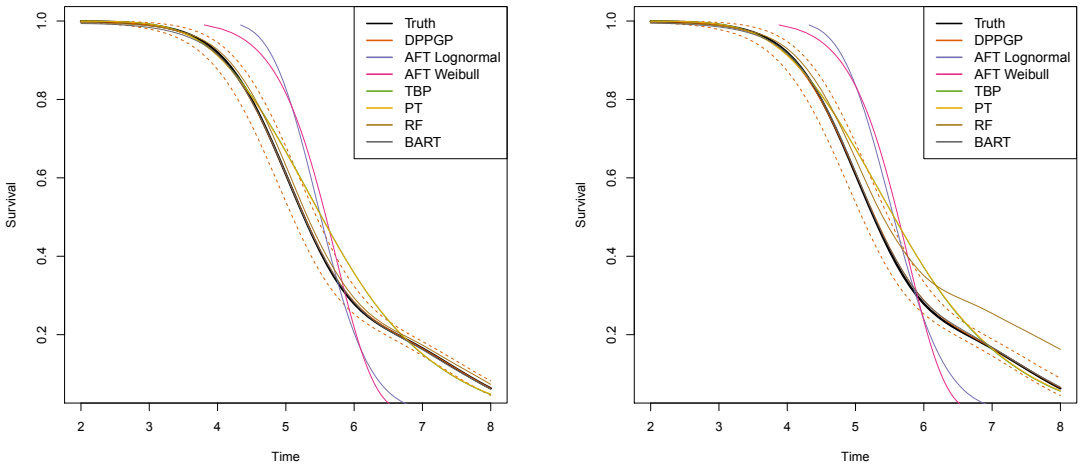

(a) $n = 200$ without censoring (b) $n = 200$ with 25% censoring

Figure 2.2: Survival function estimates for the simulated data, with survival time on the log scale. True survival functions are in black compared with the estimated posterior mean survival functions under the DDP-GP model with point-wise 95% credible bands as two dashed lines, for $n = 200$ (left) and $n = 200$ with 25% censoring (right). For comparators, we also show the survival function estimates under AFT regression models using the lognormal and Weibull distributions, TBP, PT, RF, and BART.

model based estimate, as well as the RF and BART methods, reliably recovered the shape of the true survival function, while the four other methods (AFT Lognormal, AFT Weibull, TBP, and PT) showed substantial bias.

2.5.2 Personalized optimal AUC estimation

For the simulated data, we next evaluated the ability of the DDP-GP survival regression model to estimate optimal personalized AUC ranges, computed by (2.6). This estimation can be performed by the function `DDPGP_meansurvival` in our R package `DDPGPSurv`. This function takes the output from `mcmc_DDPGP` and calculates the posterior mean survival times and 95% credible intervals for patients of interest. Figure 2.5 compares the simulated true optimal AUC and the optimal AUC range estimates for a 30-year old patient with two different CR statuses, under each of seven models: the DDP-GP model, and the lognormal or Weibull AFT models in (2.7), TBP, PT, RF, and BART. Since the RF and BART methods do not have closed-forms for mean survival times and only provide the estimated survival probabilities at the time points observed in the original data, we estimated the mean survival time as the area under the survival curve in the interval $(0, t_{\max})$, where $t_{\max}$ is the largest observed time point in the data. In Figure 2.5, the numbers in parentheses in the legend represent the simulated true optimal AUC and the optimal AUC range estimated by the DDP-GP survival regression model. The figure shows that the DDP-GP model accurately estimates the mean survival function and identifies the optimal AUC, with the simulated true AUC being in the estimated optimal AUC range. In contrast, the mean survival functions and the optimal AUC estimates given by the AFT models, PT, and TBP are considerably different from the simulation truth. For instance, when CR=No with 25% censoring, the AFT models with lognormal

or Weibull distributions estimate the optimal AUC to be 4.4 and 4.5, respectively, while the true AUC is 3.9. While the RF and BART methods are able to accurately estimate the survival function, the estimates of mean survival are biased, especially when the data are censored.

In summary, the DDP-GP is more robust than alternative methods in the sense that it can better fit the survival functions and more accurately estimate personalized optimal AUC ranges, even while only including the main effects ($\beta_0 + \beta_1 x_{i,1} + \beta_2 x_{i,2} + \beta_3 x_{i,3}$) in the mean of the Gaussian process prior. In contrast, the alternative parametric and semiparametric models which do not perform as well as the DDP-GP, include not only main effects but also quadratic terms and interactions between covariates as in (2.7). This illustrates an important advantage of the DDP-GP model. It allows one to include covariates as simple linear combinations, but still is able to identify quite general interactions that are not limited to conventional multiplicative interaction terms, such as $\beta_1 AUC \times age + \beta_2 AUC \times CR$, that typically are included in the linear components of conventional Cox or AFT models. Such a construction is extremely useful especially when the covariates are high-dimensional, in which case including all the interactions among covariates in the regression model is infeasible.

2.5.3 Additional Simulation Studies

In this additional simulation study, we evaluate whether the proposed DDP-GP model is stable under a setting with a larger number of covariates and show the necessity of the covariate-specific scale parameters λ_d's in the covariance function. For this purpose, we compared our DDP-GP model to a simpler version with only one single scale parameter, λ_{single}, called naive DDP-GP. The naive DDP-GP has an

identical setup except that the covariance function is modified to:

$$C(\boldsymbol{x}_i, \boldsymbol{x}_l) = \sigma_0^2 \ \exp\left\{ -\sum_{d=1}^{D} \left(\frac{x_{id} - x_{ld}}{\lambda_{single}} \right)^2 \right\} + \delta_{il} J^2.$$

Let $LN(m, s)$ denote a lognormal distribution with location and scale parameters m and s. We assume each patient i has a 10-dimensional covariate $\boldsymbol{x}_i = (x_{i,1}, x_{i,2}, x_{i,3}, x_{i,4}, x_{i,5}, x_{i,6}, x_{i,7}, x_{i,8}, x_{i,9}, x_{i,10})$. Patients' ages ($x_{i,1}$) and AUC ($x_{i,2}$) values were sampled with replacement from the ages and AUC's given in our IV busulfan allo-SCT dataset. We generated the patients' CR status ($x_{i,3}$) as a binary variable from a Bernoulli(0.5), which also mimics the IV busulfan dataset. The remaining seven covariates were independently sampled from a standard normal distribution $N(0, 1)$. We simulated the Y_i's from a lognormal distribution with location parameters that vary with $\boldsymbol{x}_i$, where the covariates $x_{i,9}$ and $x_{i,10}$ were not related to the response.

$$Y_i \mid \boldsymbol{x}_i \sim LN(\mu(\boldsymbol{x}_i), \sigma_0^2),$$

$$\mu(\boldsymbol{x}_i) = 4 - 0.1(x_{i,1} - x_{i,4} + 1) + 0.7(x_{i,2} + x_{i,5} - 1) + 0.3(x_{i,3} - x_{i,6}) - 0.07(x_{i,2}^2 + x_{i,7}^2)$$

$$-0.1 x_{i,1} x_{i,2}(x_{i,8} + 0.5) + 0.2 x_{i,2} x_{i,3} - 0.18 x_{i,1} x_{i,2} x_{i,3},$$

for $i = 1, \cdots, n$, and $\sigma_0 = 0.4$.

The survival density estimation was computed for two scenarios: one scenario with $n = 200$ observations without censoring and the other with $n = 200$ and 25% censoring. For each scenario, $B = 100$ trials were simulated, then the DDP-GP model and the naive DDP-GP model were fitted to each trial with 5,000 MCMC iterations and 2000 burn-in. From Figure 2.3, we can see that both models are able to recover the shape of the true survival function, with the naive DDP-GP model having a

slightly larger credible interval.

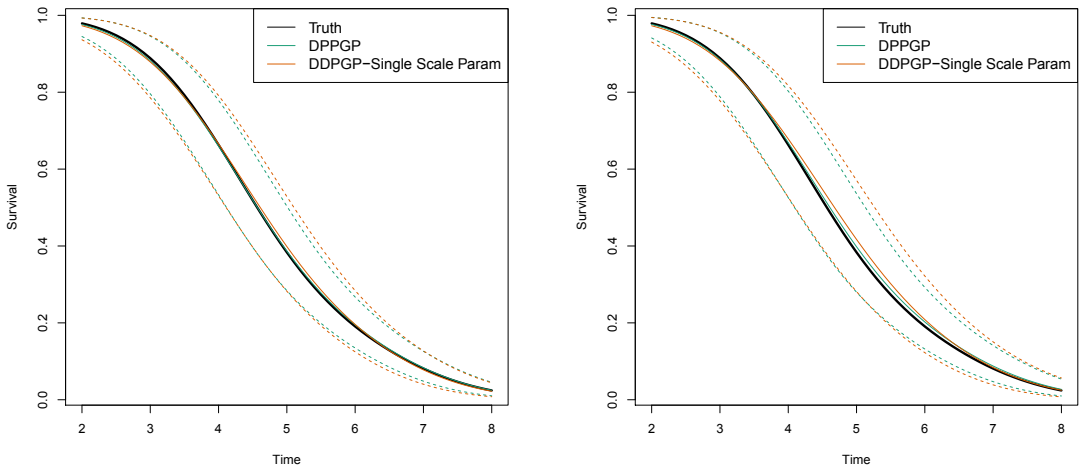

Figure 2.3: Survival function estimates for the simulated data, with survival time in log scale. True survival functions are in black and the estimated mean survival functions under the naive DDP-GP model and the DDP-GP model are in brown and green, respectively, with the point-wise 95% credible bands as the dotted lines, for $n = 200$ (left) and $n = 200$ with 25% censoring (right).

Figure 2.4 compares the simulated true optimal AUC and the optimal AUC range estimates under the DDP-GP model and the naive DDP-GP for a patient with the following covariate:

$$x_{i,1} = -1.14706 \text{ (30-year old)}$$
$$x_{i,3} = 0 \text{ (CR=No)}$$
$$x_{i,4} = x_{i,5} = x_{i,6} = x_{i,7} = 0$$
$$x_{i,8} = 0.5,\ x_{i,9} = x_{i,10} = 1.$$

In Figure 2.4, the numbers in parentheses in the legend represent the simulated true optimal AUC and the optimal AUC range estimated by the DDP-GP model. We can see that the proposed DDP-GP model recovers the true mean survival curve accurately under this complex setup with 10 covariates. In contrast, the naive DDP-GP model yields a poor estimate for the mean survival. Therefore, the proposed DDP-GP model is robust under a large number of covariates. Also, the inclusion of the covariate-specific scale parameters λ_d's in the model are necessary for accurate survival estimation and mean survival estimation.

2.6 IV busulfan Data Analysis

While an optimal AUC interval has been determined previously for use in all patients (Andersson et al., 2002), the underlying statistical analyses motivating this assume homogeneity, and thus do not allow the possibility that the optimal interval may vary non-trivially with patient characteristics. Here, we approach the problem differently by estimating mean survival time as a function of (age, CR status, AUC), allowing the possibility that the effect of AUC on survival may vary with age and CR.

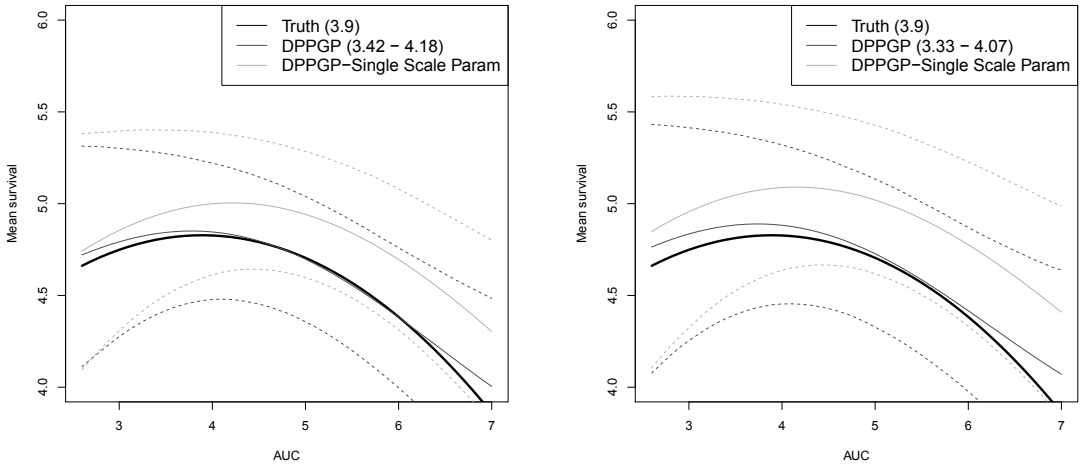

Figure 2.4: Optimal AUC estimation for the simulated data, with both survival time and AUC on the log scale. True mean survival functions versus AUC are in black and the estimated mean survival functions under the DDP-GP model are in dark gray with point-wise 95% credible bands as two dotted dark gray lines for $n = 200$ (left) and $n = 200$ with 25% censoring (right). Similarly, the estimated mean survival functions under the naive DDP-GP model are shown in light gray with point-wise 95% credible bands.

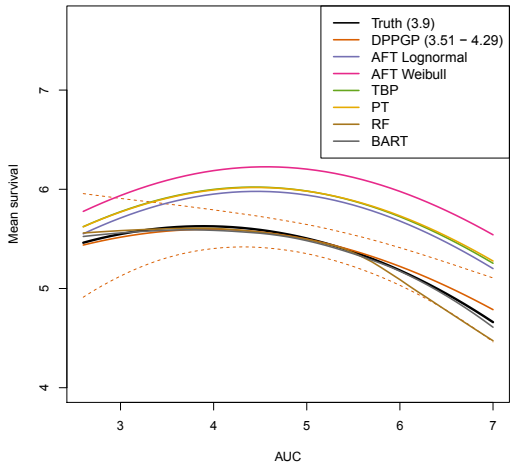

(a) $n = 200$ without censoring, CR=No

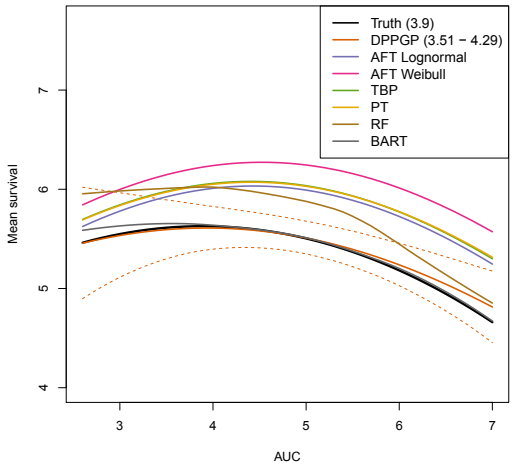

(b) $n = 200$ with 25% censoring, CR=No

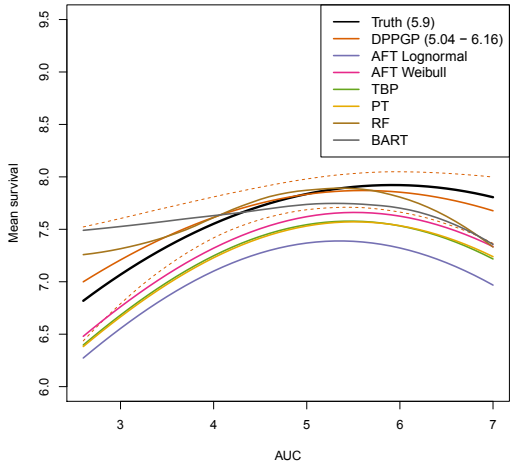

(c) $n = 200$ without censoring, CR=Yes

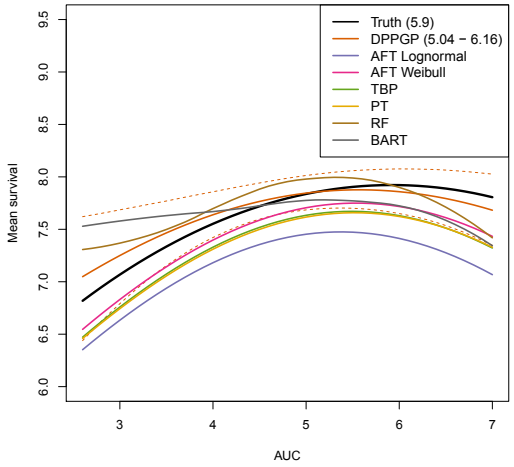

(d) $n = 200$ with 25% censoring, CR=Yes

Figure 2.5: Optimal AUC estimation for the simulated data, with both survival time and AUC on the log scale. True mean survival functions versus AUC are in black compared with estimated mean survival functions under the DDP-GP model with point-wise 95% credible bands as two dotted lines for $n = 200$ (left plots) and $n = 200$ with 25% censoring (right plots). For the comparators, we also show the mean survival function estimates under AFT regression models using the lognormal and Weibull distributions, TBP, PT, RF, and BART. The numbers in parentheses in the legend are the true and estimated optimal AUC values.

The AUC values in our analysis are in units of thousands of mean daily $\mu M/^*$min. An initial analysis of the IV Bu dataset using Kaplan-Meier estimates is given in Figure 2.6. We divided patients into four groups based on CR status and age, dichotomized as being above or below the median age of 49, and plotted their survival probabilities. Figure 2.6 illustrates the well-known fact that being in CR at transplant yields higher survival probabilities. Similarly, younger patients are also expected to have higher survival probabilities. The p-value obtained from the log rank test comparing the survival distributions between the four groups is significant, indicating that CR and age are important covariates for any survival regression model. The cut-off 49 for dichotomizing age was chosen for convenience, however, as is commonly done in survival analyses. In addition to loss of information about the joint effect of age and CR status on survival caused by dichotomizing age, the reliability of each Kaplan-Meier estimate is reduced because it is based on a subsample.

We fit the DDP-GP survival regression model to the allo-SCT dataset with 10,000 Gibbs sampler iterations and a burn-in of 5,000 iterations. The estimated posterior survival distributions with 95% credible intervals under the DDP-GP for patients with different CR statuses and ages 30, 40, 50, or 60, given AUC=5, are shown in Figure 2.7, respectively. For each (CR status, age) combination, the optimal AUC range is defined as the AUC value that maximizes estimated posterior mean survival, $\pm$ 10%. Given CR status and AUC, Figure 2.7 shows that the estimated posterior mean survival function decreases for older patients, agreeing with what was seen in the preliminary Kaplan-Meier estimates.

For the eight combinations of CR status and Age, we calculated predicted posterior mean survival time as a function of AUC, to address the primary goal of the analyses. These plots are given in Figure 2.8. Our analyses confirm the existence,

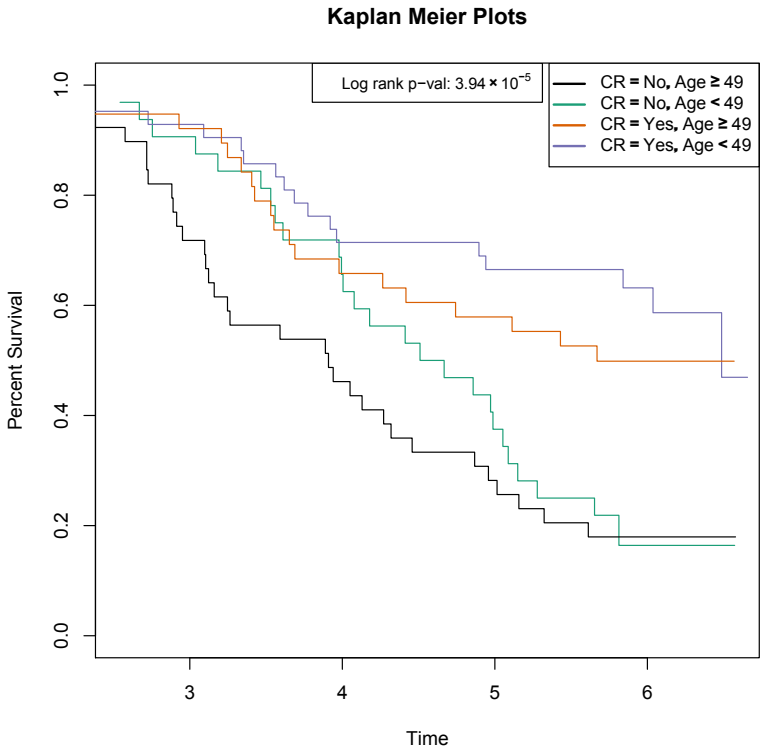

Figure 2.6: Kaplan Meier Plots. The time in weeks (log scale) versus probability of survival for four different groups are plotted. The p-value from the log rank test for comparison between the survival distributions between the four groups is given at the top of the figure.

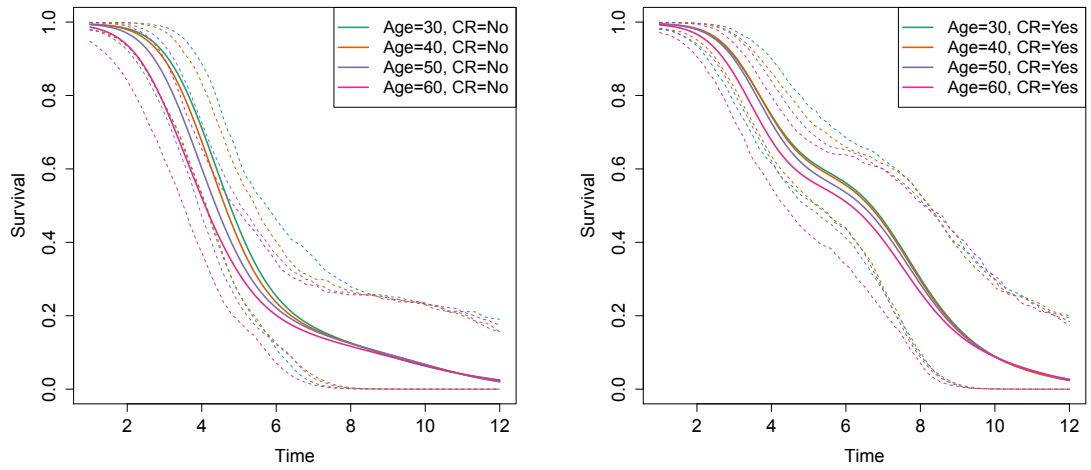

Figure 2.7: Estimated survival functions under the DDP-GP survival regression model for patients with different CR status (Yes or No) and ages (30, 40, 50, 60). The patients are assigned AUC=5. The dashed lines represent the point-wise 95% credible intervals for each survival curve.

for each combination of CR status and Age, of an optimal AUC range that yields higher expected survival times compared to an AUC that is either below or above the optimal range. A very important inference is that these optimal AUC ranges differ substantially between many of the (CR status, Age) combinations. This has extremely important therapeutic implications when choosing an individual patient's targeted AUC. For example, the optimal AUC interval for a patient not in CR with Age=50 is $4.7 \pm 0.47 = [4.23, 5.17]$ compared with the optimal interval $5.8 \pm 0.58 = [5.22, 6.38]$ for a patient in CR with Age=40. Since these intervals are disjoint, they suggest that these two patients should have very different targeted AUC values to maximize their expected survival times.

In contrast with our inferences, (Bartelink et al., 2016) concluded that CR status has a negligible effect on the optimal AUC. However, the results reported by Bartelink et al. (2016) were based on data from a large number of different medical

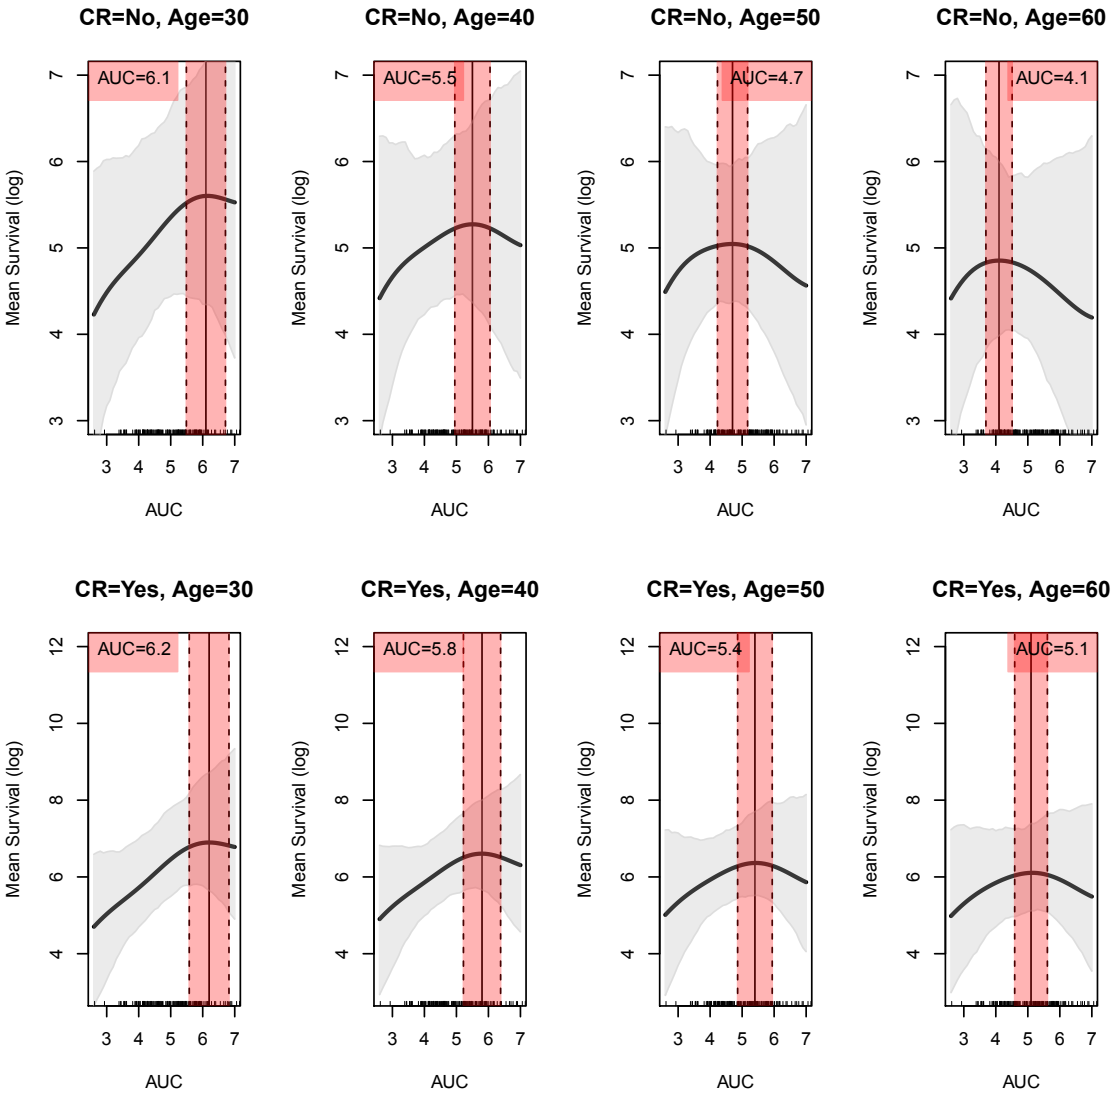

Figure 2.8: Mean log survival time estimates under the DDP-GP model, as a function of AUC, for each of eight (CR status, Age) combinations. The light gray area in each plot represents the 95% credible interval for estimated mean survival, and the tick marks on the horizontal axis (rug plot) indicate the AUC values for patients in the data set. The red area bounded by dashed lines represents the optimal AUC range, defined as the estimated mean ±10%.

centers, many different pretransplant conditioning regimens were used, the PK-data were obtained from different laboratories, with a very heterogeneous pediatric patient population having a large number of different diagnostic categories, including patients with malignant and non-malignant genetic disorders. In contrast, our analyses are based on a much more homogeneous dataset. Our results indicate that CR status is an important covariate, and that the optimal dose of AUC is higher for patients who are in CR at transplant. Furthermore, the increased optimal AUC for patients in CR at transplant versus patients not in CR is much larger in older patients, whereas these differences appear negligible in adolescents or young adults, similar to what was reported by Bartelink, et al. (2016). Our results also demonstrate that, across all ages, mean survival time for patients in CR is larger compared with those not CR.

To further illustrate how the optimal AUC ranges change with both CR and Age, we plotted the optimal AUC ranges as Age is varied continuously, for CR=Yes and CR=No, in Figure 2.9. The negative association between optimal AUC and Age is clearly shown by this figure. It also shows that, while CR status has virtually no effect on the optimal AUC interval for very young patients with Age ≤ 28, the optimal AUC for patients in CR at transplant is increasingly higher as Age increases, with the optimal intervals for CR = Yes versus CR = No becoming completely disjoint for patients above 55 years of age. Thus, the lower portions of the curves in Figure 2.9 for Age ≤ 28, agree with the conclusion of (Bartelink et al., 2016) for pediatric and adolescent patients, while the higher portions for Age > 28, provide news insights. Again, this demonstrates the importance of considering both CR status and Age when planning a targeted AUC for a patient with a diagnosis of AML or MDS.

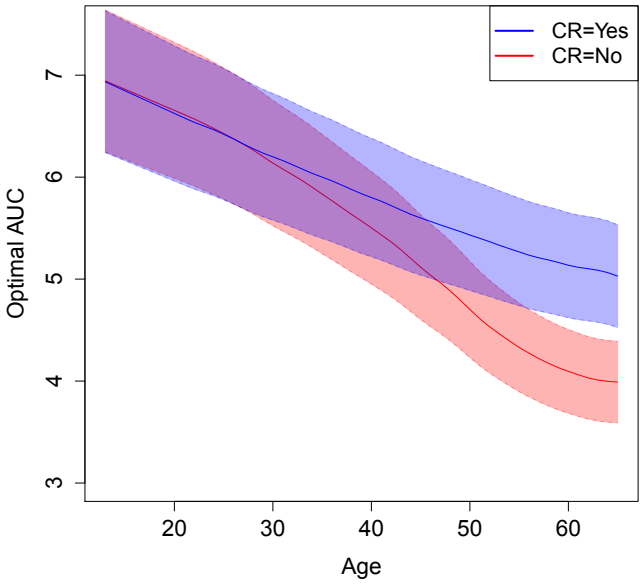

Figure 2.9: Optimal AUC ranges versus age given CR status. The blue and red lines represent the optimal AUC for CR=Yes and No, respectively. The optimal AUC ranges are represented by the shaded regions above and below the optimal AUC.

2.6.1 Comparison Models

We compare the mean survival estimates under the DDP-GP with those produced by the PT, TBP, RF, and BART, shown in Figures 2.10-2.13 respectively. For the semiparametric models (PT and TBP), we can see that the optimal AUC decreases as patient age increases, and is larger for CR=Yes comparing to CR=No. However, most combinations of CR status and Age do not present quadratic patterns, which do not agree with the clinical knowledge.

The nonparametric models (RF and BART) produce very noisy results, leading to difficult clinical interpretations. This may be caused by the limitation of the RF and BART methods mentioned in Section 5 of the main manuscript that they only provide the estimated survivals at the time points observed in the original data, leading to potential bias in estimating the mean survival time.

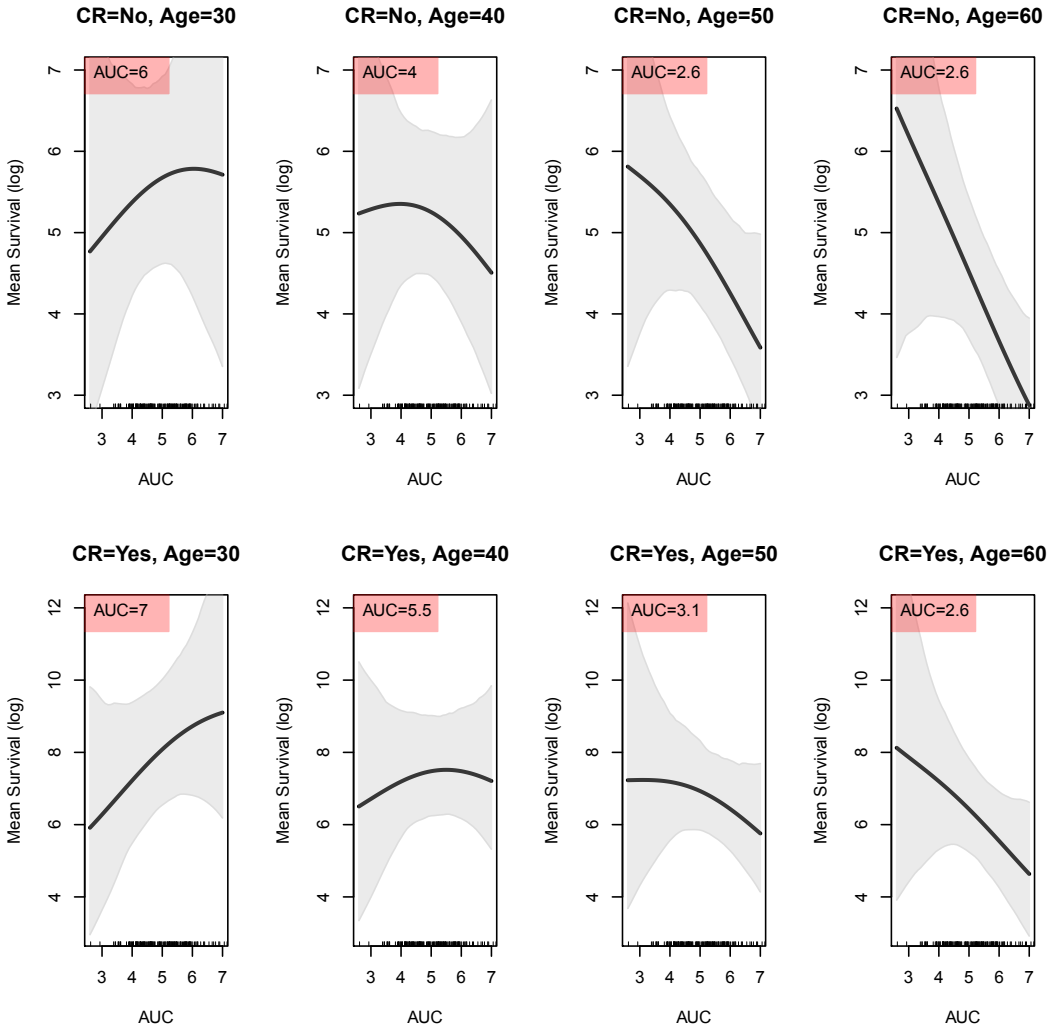

Figure 2.10: Mean survival time estimates under the PT survival model, as a function of AUC, for each of eight (CR status, Age) combinations. Mean survival time is in the log scale. The gray area in each plot represents the 95% credible interval for estimated mean survival, and the tick marks on the horizontal axis (rug plot) indicate the AUC values for patients in the data set.

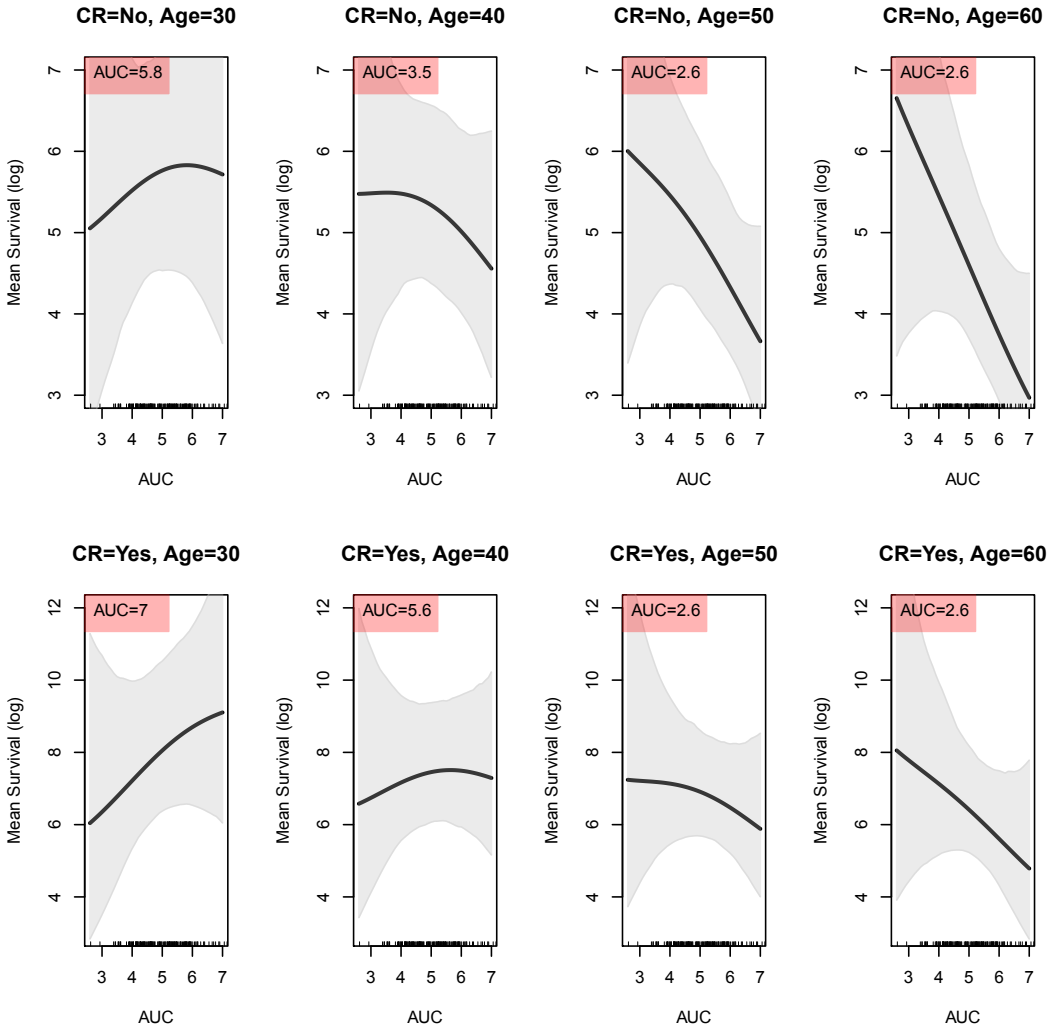

Figure 2.11: Mean survival time estimates under the TBP survival model, as a function of AUC, for each of eight (CR status, Age) combinations. Mean survival time is in the log scale. The gray area in each plot represents the 95% credible interval for estimated mean survival, and the tick marks on the horizontal axis (rug plot) indicate the AUC values for patients in the data set.

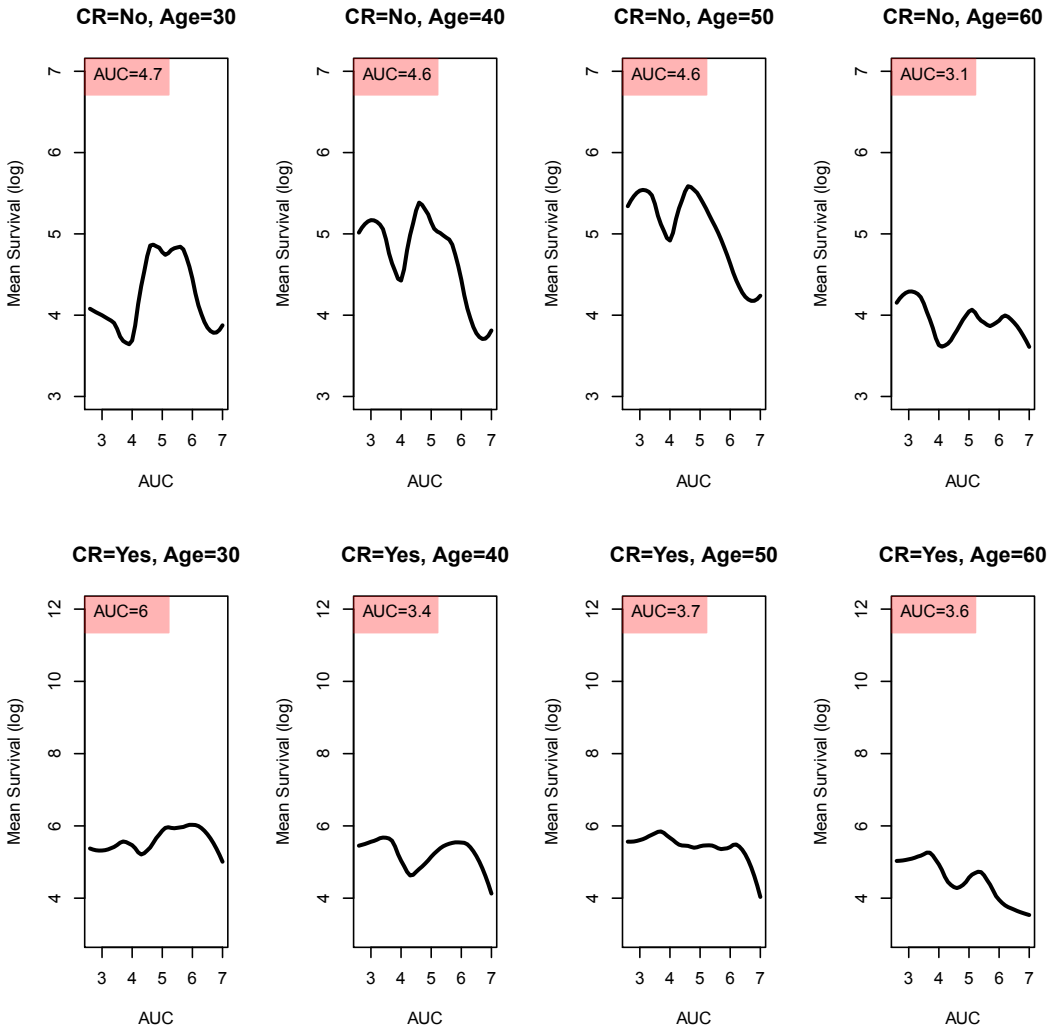

Figure 2.12: Mean survival time estimates under the RF survival model, as a function of AUC, for each of eight (CR status, Age) combinations. Mean survival time is in the log scale. The tick marks on the horizontal axis (rug plot) indicate the AUC values for patients in the data set.

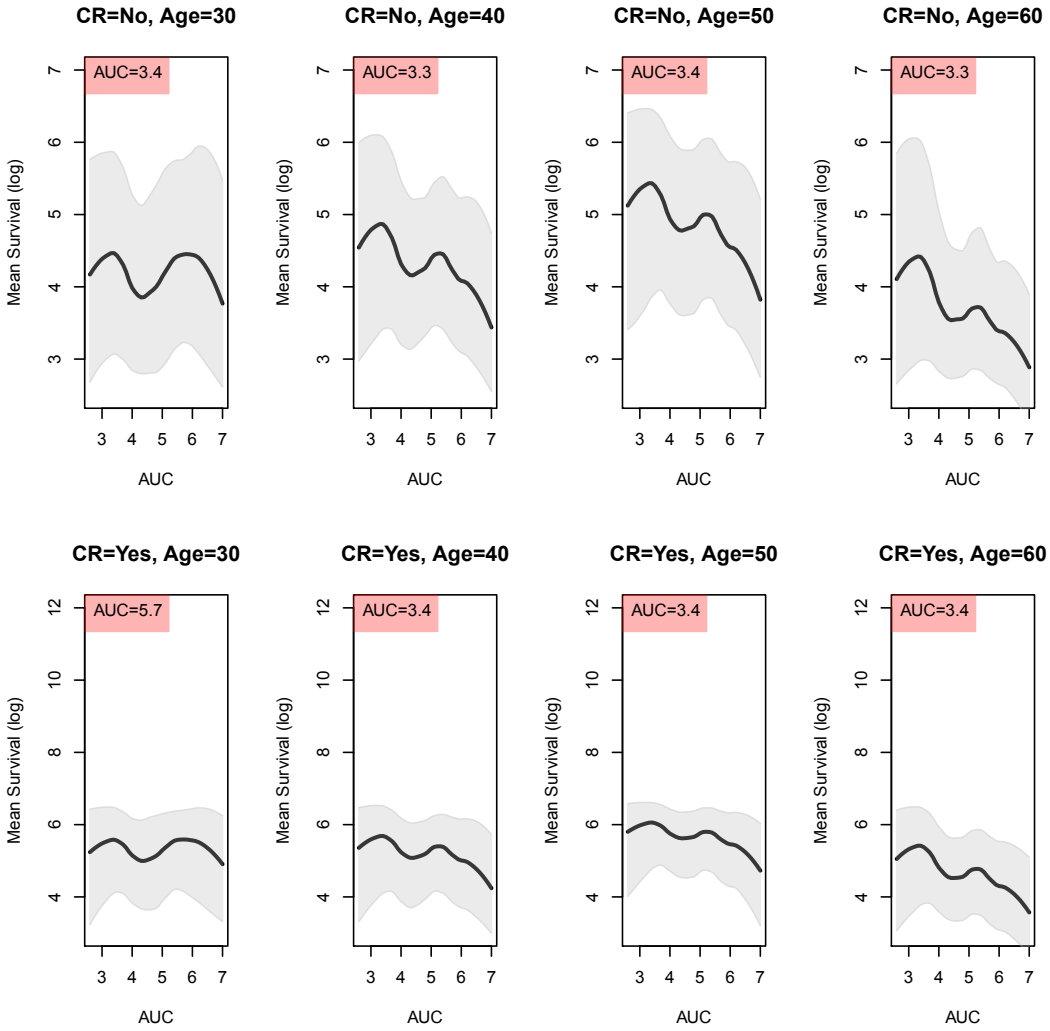

Figure 2.13: Mean survival time estimates under the BART survival model, as a function of AUC, for each of eight (CR status, Age) combinations. Mean survival time is in the log scale. The gray area in each plot represents the 95% credible interval for estimated mean survival, and the tick marks on the horizontal axis (rug plot) indicate the AUC values for patients in the data set.

2.7 Conclusions

We have proposed an extended Bayesian nonparametric DDP-GP model for survival regression having a generalized covariance structure, studied it by simulation, and applied it to estimate personalized optimal dose intervals for IV busulfan in allo-SCT for AML/MDS. Our simulations, constructed to mimic the dataset, show that the DDP-GP model provides more accurate survival function estimates and optimal AUC range estimates compared with conventional parametric to AFT models. Our analyses of the IV busulfan allo-SCT dataset identified optimal AUC intervals, varying with the patient's CR status and Age, that previously have not been known for this treatment. Our results may have profound therapeutic implications, since they provide a basis for personalized medicine by enabling physicians to prospectively assign an optimized therapeutic target interval for each patient based on his/her CR status and age.

More generally, we have developed an R package, `DDPGPSurv`, that implements the DDP-GP model for a broad range of survival regression analyses. While the DDP-GP is more complex than conventional survival regression models, its robustness and broad applicability make it an attractive methodology for survival analysis. The DDP-GP based data analysis reported here, while important in its own right, identified a nonlinear three-way interaction between age, CR status, and AUC in their joint effect on survival time, as shown by Figures 6 and 7. This pattern was identified despite the fact, noted above, that only the main effects were included in the mean of the Gaussian process prior via the linear term $\beta_0 + \beta_1 Age + \beta_2 CR + \beta_3 AUC$. This is because the DDP-GP is essentially a mixture model, hence it can identify complex patterns in the data that may be missed by conventional models. For the allo-SCT IV busulfan data, this may be related to the multi-modality of the survival time distri-

bution, seen in Figure 1. This illustrates the practical advantage that, when applying the DDP-GP, one need not guess or search for complex patterns in the linear term of the covariates, as is done routinely when applying conventional survival regression models.

2.8 Posterior Inference Details

The following section summarizes the generalized DDP-GP model and contains derivations for the Gibbs sampling steps.

2.8.1 Probability Model

For each patient i, $i = 1, \ldots, n$, denote Y_i and $\boldsymbol{x}_i$ to be the log survival time and patients' covariates. Let $T_i = \min(Y_i, C_i)$, where C_i is the censoring time for patient i. Denote $\delta_i = I_{\{Y_i \leq C_i\}}$. Let $f(y)$, $F(y)$ and $S(y)$ be the probability density function, the distribution function, and the survival function of Y, respectively.

Denote Θ the collection of all model parameters and $D_n = \{T_i, \delta_i, \boldsymbol{x}_i\}_{i=1}^n$, the likelihood function is

$$L(\Theta \mid D_n) = \prod_{i=1}^{n} f_{\boldsymbol{x}_i}(t_i \mid \Theta)^{\delta_i} S_{\boldsymbol{x}_i}(t_i \mid \Theta)^{1-\delta_i},$$

where $f_{\boldsymbol{x}}(\cdot)$ and $S_{\boldsymbol{x}}(\cdot)$ denote the density and the survival function of an individual with covariate $\boldsymbol{x}$.

We assume the Dependent Dirichlet Process (DDP) mixture model for F:

$$F(y_i \mid \boldsymbol{x}_i) = \sum_{h=1}^{\infty} w_h N(y_i; \theta_h(\boldsymbol{x}_i), \sigma^2),$$

where

$$w_h = v_h \prod_{h=1}^{\infty} \{1 - v_k\},$$

$$v_h \sim Beta(1, \alpha).$$

For each h, θ_h is sampled from a Gaussian Process:

$$\{\theta_h(\boldsymbol{x})\} \sim GP(\mu_h(\cdot), C(\cdot, \cdot))$$

$$\mu_h(\boldsymbol{x}_i; \beta_h) = \boldsymbol{x}_i \beta_h$$

$$C(\boldsymbol{x}_i, \boldsymbol{x}_l) = \sigma_0^2 \ \exp\left\{ -\sum_{d=1}^{D} \frac{(x_{id} - x_{ld})^2}{\lambda_d^2} \right\} + \delta_{il} J^2.$$

where J usually takes a small value (e.g., $J = 0.1$).

2.8.2 Collapsed Gibbs Sampler

We use $r_i = h$ to indicate that ith patient is assigned to hth cluster. Let H be the number of clusters, n_h be the number of patients in cluster h, and $X = (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)'$. Also, we denote $\boldsymbol{\theta}_h = (\theta_h(\boldsymbol{x}_i), \ldots, \theta_h(\boldsymbol{x}_n))$.

Since

$$\boldsymbol{\theta}_h \sim GP(X\boldsymbol{\beta}_h, C(\cdot, \cdot)), \quad \boldsymbol{\beta}_h \sim N(\boldsymbol{\beta}_0, \Sigma_0),$$

marginalizing $\boldsymbol{\beta}_h$ leads to

$$\boldsymbol{\theta}_h \sim GP(X\boldsymbol{\beta}_0, X\Sigma_0 X' + C)$$

1. Update r_i and $\boldsymbol{\theta}_i \mid \boldsymbol{\theta}_{-i}, \cdot$.

   - When $\delta_i = 1$,

$$p(r_i = h \mid \cdot) \propto \begin{cases} n_h^- \phi(t_i; \theta_h(\boldsymbol{x}_i), \sigma^2) & h \leq H^-, \\ \alpha \phi(t_i; \boldsymbol{u} X \boldsymbol{\beta}_0, \boldsymbol{u}(X \Sigma_0 X + C)\boldsymbol{u}' + \sigma^2) & h = H^- + 1. \end{cases}$$

Here $\boldsymbol{u} = (u_1, \ldots, u_n)$ with the element $u_i = 1$ and all other elements being 0; n_h^- is the number of patients assigned to the hth cluster after taking out the ith patient; H^- is the number of clusters after taking out the ith patient. ϕ refers to the normal distribution PDF, while Φ refers to the normal distribution CDF.

- When $\delta_i = 0$,

$$p(r_i = h \mid \cdot) \propto \begin{cases} n_h^-(1 - \Phi(C_i; \theta_h(\boldsymbol{x}_i), \sigma^2)) & h \leq H^-, \\ \alpha(1 - \Phi(C_i; \boldsymbol{u} X \boldsymbol{\beta}_0, \boldsymbol{u}(X \Sigma_0 X + C)\boldsymbol{u}' + \sigma^2)) & h = H^- + 1. \end{cases}$$

Once r_i has been sampled, we impute y_i from a left truncated normal distribution, that is p_{r_i} left truncated at C_i:

$$p_{r_i} \propto \begin{cases} N(\boldsymbol{\theta}_{r_i}, \sigma^2) & h \leq H^-, \\ N(0, \sigma^2 + A_i) & h = H^- + 1, \end{cases}$$

where $A = X \Sigma_0 X + C$ and A_i is the (i,i)-th element of A.

Then,

$$p(\boldsymbol{\theta}_i \mid r_i = h, \boldsymbol{\theta}_{-i}, \cdot) = \begin{cases} \delta_{\boldsymbol{\theta}_h} & h \leq H^- \\ p(\boldsymbol{\theta}_i \mid y_i, G_0) & h = H^- + 1. \end{cases}$$

Here

$$p(\boldsymbol{\theta}_i \mid y_i, G_0) \sim N((\frac{\boldsymbol{u}'\boldsymbol{u}}{\sigma^2} + A^{-1})^{-1}(\frac{\boldsymbol{u}'y_i}{\sigma^2} + A^{-1}X\boldsymbol{\beta}_0), (\frac{\boldsymbol{u}'\boldsymbol{u}}{\sigma^2} + A^{-1})^{-1}).$$

2. Update $\boldsymbol{\beta}_h$

$$\begin{aligned} p(\boldsymbol{\beta}_h \mid \cdot) &\propto p(\boldsymbol{\theta}_h \mid X, \boldsymbol{\beta}_h)p(\boldsymbol{\beta}_h) \\ &\propto exp\{-\frac{1}{2}(\boldsymbol{\beta}_h - \boldsymbol{\beta}_0)'\Sigma_0^{-1}(\boldsymbol{\beta}_h - \boldsymbol{\beta}_0) - \frac{1}{2}(\boldsymbol{\theta}_h - X\boldsymbol{\beta}_h)'C^{-1}(\boldsymbol{\theta}_h - X\boldsymbol{\beta}_h)\}. \end{aligned}$$

Therefore,

$$p(\boldsymbol{\beta}_h \mid \cdot) \sim N((\Sigma_0^{-1} + X'C^{-1}X)^{-1}(X'C^{-1}\boldsymbol{\theta}_h + \Sigma_0^{-1}\boldsymbol{\beta}_0), (\Sigma_0^{-1} + X'C^{-1}X)^{-1}).$$

3. Update $\boldsymbol{\theta}_h = (\theta_h(\boldsymbol{x}_i), \ldots, \theta_h(\boldsymbol{x}_n))$

$$p(\boldsymbol{\theta}_h \mid \cdot) \propto p(\boldsymbol{\theta}_h \mid \boldsymbol{\beta}_h) \prod_{i:r_i=h} f(y_i \mid \theta_h(\boldsymbol{x}_i)).$$

Let Y_h be the vector whose entries are the survival times for patients who are assigned to the hth cluster, i.e., $r_i = h$. Let $n_h = \#\{i : r_i = h\}$. Denote U to be an $n_h \times n$ matrix: if patient i is the jth element of cluster h, then $U_{ji} = 1$; otherwise 0. We have

$$\begin{aligned} p(\boldsymbol{\theta}_h \mid \cdot) &\propto exp\{-\frac{1}{2}[(\boldsymbol{\theta}_h - X\boldsymbol{\beta}_h)'C^{-1}(\boldsymbol{\theta}_h - X\boldsymbol{\beta}_h) + \frac{(Y_h - U\boldsymbol{\theta}_h)'(Y_h - U\boldsymbol{\theta}_h)}{\sigma^2}]\} \\ &\propto exp\{-\frac{1}{2}\boldsymbol{\theta}_h'(C^{-1} + \frac{U'U}{\sigma^2})\boldsymbol{\theta}_h + \boldsymbol{\theta}_h'(C^{-1}X\boldsymbol{\beta}_h + \frac{U'Y_h}{\sigma^2})\} \end{aligned}$$

Therefore,

$$p(\boldsymbol{\theta}_h \mid \cdot) \sim N((C^{-1} + \frac{U'U}{\sigma^2})^{-1}(C^{-1}X\boldsymbol{\beta}_h + \frac{U'Y_h}{\sigma^2}), (C^{-1} + \frac{U'U}{\sigma^2})^{-1}).$$

4. Update σ^2

$$\sigma^2 \sim IG(a_1 + \frac{n}{2}, b_1 + \frac{\sum_{h=1}^{H}\sum_{i:r_i=h}(y_i - \theta_h(\boldsymbol{x}_i))^2}{2}),$$

where the prior for σ^2 is $IG(a_1, b_1)$.

5. Update σ_0 and λ_d

$$p(\sigma_0 \mid \cdot) \quad \propto \quad \prod_h exp\{-\frac{1}{2}(\boldsymbol{\theta}_h - X\boldsymbol{\beta}_h)'C^{-1}(\boldsymbol{\theta}_h - X\boldsymbol{\beta}_h)\}p(\sigma_0).$$

Here C is a matrix with σ_0 as a parameter in each matrix element. Since there is no closed form for the posterior distribution of σ_0, we use a Metropolis-Hastings step to update σ_0. Similarly for λ_d, $d = 1, \cdots, D$. Here D represents the total number of covariates. i.e., the dimension of $\boldsymbol{x}_i$.

6. Update α

The prior for α is Gamma(a_2, b_2). We generate a latent variable η, conditioning on α and the number of patients,

$$\eta \mid \alpha, n \sim Beta(\alpha + 1, n).$$

Then,

$$\begin{aligned}\alpha \mid \eta, H \;\sim\; & \frac{a_2 + H - 1}{a_2 + H - 1 + n(b_2 - log\{\eta\})} Ga(a_2 + H, b_2 - log\{\eta\}) \\ & + \frac{n(b_2 - log\{\eta\})}{a_2 + H - 1 + n(b_2 - log\{\eta\})} Ga(a_2 + H - 1, b_2 - log\{\eta\}).\end{aligned}$$

2.8.3 Predictive Inference

$$\begin{aligned}p(y_{n+1} \mid \boldsymbol{x}_{n+1}, \sigma^2) \;=\; & \sum_{h=1}^{H} \frac{n_h}{\alpha + n} \int N(\boldsymbol{x}_{n+1}\boldsymbol{\beta}_h + aC^{-1}(\boldsymbol{\theta}_h - \boldsymbol{X}\boldsymbol{\beta}_h), \sigma^2 + \sigma_0^2 - aC^{-1}a') \\ & p(\boldsymbol{\theta}_h, \sigma^2 \mid D_n) d\sigma^2 d\boldsymbol{\theta}_h d\boldsymbol{\beta}_h \\ & + \frac{\alpha}{n + \alpha} \int N(\boldsymbol{x}_{n+1}\boldsymbol{\beta}_0, \sigma^2 + \sigma_0^2 + \boldsymbol{x}_{n+1}\Sigma_0\boldsymbol{x}'_{n+1}) p(\sigma^2 \mid D_n) d\sigma^2,\end{aligned}$$

where $a = (C(\boldsymbol{x}_1, \boldsymbol{x}_{n+1}), C(\boldsymbol{x}_2, \boldsymbol{x}_{n+1}), \cdots, C(\boldsymbol{x}_n, \boldsymbol{x}_{n+1}))$.

Chapter 3

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

3.1 Introduction

In biomedical applications involving long-term personalized care of patients with chronic health conditions (e.g., diabetes, HIV 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 re-

quired 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. This chapter develops a Bayesian joint framework consisting of a generative probabilistic model for clinical decisions with timing and a model for clinical observations (e.g., longitudinal clinical measurements and time-to-event data): these two models share certain structures and parameters in order to capture the mutual influence between the clinical observations and decisions. Furthermore, we propose an optimization method that allows the decision model, by interacting with the other parts of the joint framework, to learn to make the personalized optimal clinical decision at the right time. Such a joint model and the proposed optimization method will be useful in many biomedical applications. We elaborate on one signature application in section 3.1.1, explain why existing methods won't work well on it in section 3.1.2, and then give an overview of our method and its technical novelty in section 3.1.3.

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 (Bicalho et al., 2019; Lamb and Lodhi, 2011). 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 the 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 very 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 transplantation (DIVAT), provide us 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 3.1 presents one randomly selected patient's longitudinal creatinine levels, 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 were also slowly decreasing 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.

3.1.2 Why not use existing methods?

Although many statistical and machine learning DTR methods have been developed to optimize sequential clinical decisions (Chakraborty, 2013; Laber et al., 2014; Luckett et al., 2019; Murphy et al., 2003; Xu et al., 2016), 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 adjust other clinical decisions. 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

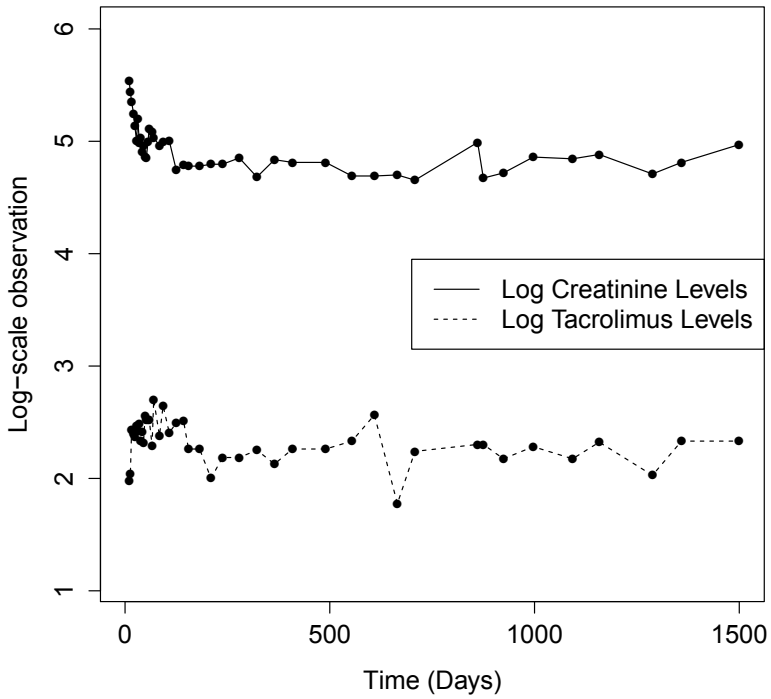

Figure 3.1: Example data for one patient's creatinine and tacrolimus levels on a log scale over time. The points represent actual visitations.

or achieved complete remission first then relapsed. 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 the treatment as given. Guan et al. (2019) attempted to optimize treatment schedules: they 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 3.1.1: 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.

3.1.3 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 model into a Bayesian joint framework that also models clinical observations; (3) we train it using a policy gradient method to generate personalized treatment schedules alongside other clinical decisions that would optimize patients' health outcomes.

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 3.1.1, 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 framework 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 3.1.1. 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).

Next, we fit the proposed Bayesian joint model on clinical observations and decisions to the data, and then 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 joint framework allows the learning to account for parameter uncertainties in the observation model. Figure 3.2 illustrates the proposed Bayesian joint framework and how its components interact. 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`.

The rest of the chapter is organized as follows. In section 3.2, we present the proposed Bayesian joint model consisting of the decision model (for visitation schedules and dosages) and the observation model (for clinical longitudinal measurements and patient survival). In section 3.3, we elaborate on our optimization procedure for the decision model. We evaluate the proposed method through simulation studies in section 3.4 and applying it to the DIVAT kidney transplantation dataset in section 3.5. Lastly, we conclude the chapter with a discussion in section 3.6.

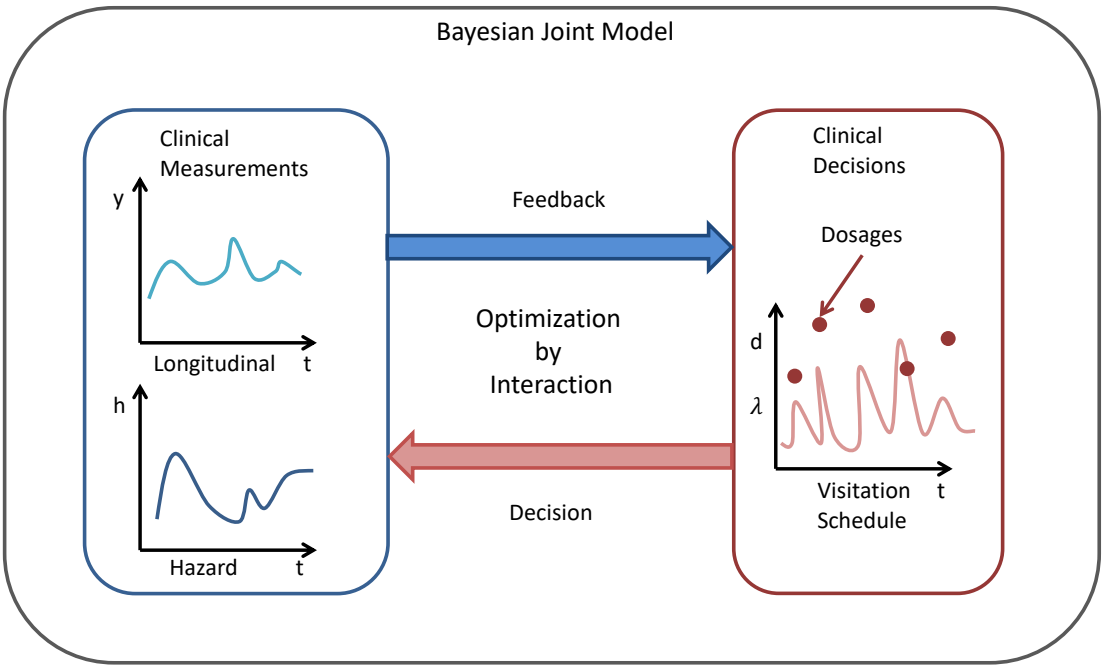

Figure 3.2: Illustration of the proposed method.

3.2 A Bayesian Joint Model

In this section, we describe the proposed Bayesian joint framework that models both clinical decisions and observations. In section 3.2.1, we introduce the clinical decision model for follow-up visitation schedules and dosages; in section 3.2.2, 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.

3.2.1 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., $\mathcal{H} = \{(t_0, d_0), (t_1, d_1), \ldots, (t_J, d_J)\}$ 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 intensity function $\lambda(t)$, namely $\lambda(t) = \lim_{dt \to 0} \frac{Pr\{\text{event happens in } [t, t+dt) | \mathcal{H}_t\}}{dt}$. 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 3.3(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 is correlated with 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. Its neural extensions (Du et al., 2016; Mei and Eisner, 2017) are flexible enough to fit complex data but unable to explicitly incorporate human expert knowledge.

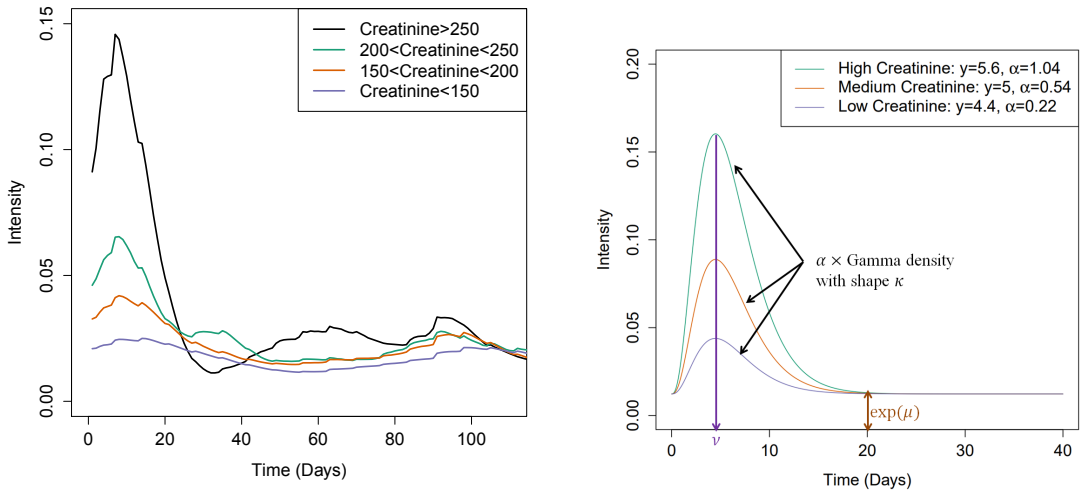

(a) Empirical intensity (b) The proposed intensity

Figure 3.3: 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 $k = 2$, $\boldsymbol{\beta}_\alpha = (10, -1.8)^T, \mu = -4.4$, $\nu_1 = 1.5$, 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. Denote $y_{i,j}$ to be the longitudinal clinical measurement of interest: in our kidney transplantation application, it is the logarithm of the creatinine level (μmol/l), for patient i at the j-th follow-up visitation occurring at time $t_{i,j}$ (days), $i = 1, \ldots, I$, $j = 0, \ldots, J_i$. Note here $t_{i,0} = 0$ denotes the transplantation date of patient i, and $y_{i,0}$ denotes the initial creatinine level. Our conditional intensity

function makes use of a Gamma density function as follows:

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

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 is often correlated with the clinical measurement (e.g., as in Figure 3.3(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:

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

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 3.3(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: in the kidney transplantation application, it is the logarithm of the tacrolimus level (ng/ml). Generally speaking, the physician would assign a dosage

based on the patient's current clinical measurement $y_{i,j}$ and potential risk factors $\boldsymbol{x}_i$. We assume the following dosage model reflecting this knowledge:

$$d_{i,j} = (1, y_{i,j}, \boldsymbol{x}_i)\boldsymbol{\beta}_d + \zeta_{i,j}, \tag{3.2}$$

where $\zeta_{i,j} \sim \text{Normal}(0, \sigma_d^2)$. In the kidney transplantation application, $\boldsymbol{x}_i$ includes baseline risk factors that would affect graft failure such as the patient's age when receiving the transplantation and the donor type. Thus, the probability of the i-th patient's sequence of visitations and assigned dosages $\boldsymbol{e}_{i,T_i} = \{(t_{i,0}, d_{i,0}), \ldots, (t_{i,J_i}, d_{i,J_i})\}$ up to time T_i can be written as

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

where $\boldsymbol{y}_i = (y_{i,0}, \ldots, y_{i,J_i})$, $\boldsymbol{\beta}_v = \{\mu, \nu_1, \nu_2, \xi, \boldsymbol{\beta}_\alpha\}$.

3.2.2 Modeling clinical observations

In this section, we introduce the clinical observation model of the Bayesian joint framework that handles longitudinal measurements and time-to-event data. In the kidney transplantation application, the longitudinal measurement is the creatinine level and the time-to-event data is the graft survival time. We will also show how it is linked to the MTPP model proposed in section 3.2.1 by carefully designing parameter sharing in order to capture the mutual influence between clinical observations and decisions. Shortly in section 3.3, we will leverage this joint framework 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 $t_{i,j}$, $i = 1, \ldots, I$, $j = 0, \ldots, J_i$. Let $y_i^*(t)$ be the underlying true but unobserved longitudinal process at time $t \geq 0$. We assume

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

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

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

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 $\boldsymbol{z}$ and $\boldsymbol{r}$ 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 $\boldsymbol{d}_i = (d_{i,0}, \ldots, d_{i,J_i})$, the probability of the observed sequence of creatinine measurements $\boldsymbol{y}_i$ is

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

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:

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

where ω is the shape parameter of the Weibull baseline hazard. 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 equation (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.

Let T_i and C_i denote the survival and censoring times for patient i, respectively. We observe only $\widetilde{T}_i = \min(T_i, C_i)$ and the censoring indicator $\delta_i = \mathbb{1}(T_i \leq C_i)$. 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

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

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

$$\begin{aligned} &\prod_{i=1}^{I} p(\boldsymbol{y}_i, \boldsymbol{e}_{i,\widetilde{T}_i}, \widetilde{T}_i, \delta_i \mid \boldsymbol{x}_i, \boldsymbol{\beta}_l, \boldsymbol{\beta}_d, \boldsymbol{\beta}_v, \boldsymbol{\beta}_s, \boldsymbol{b}_i, \sigma_l^2, \sigma_d^2) \\ &\propto \prod_{i=1}^{I} \Bigg(\underbrace{p(\boldsymbol{e}_{i,\widetilde{T}_i} \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{\beta}_v, \boldsymbol{\beta}_d, \sigma_d^2)}_{(3.11)} \underbrace{p(\boldsymbol{y}_i \mid \boldsymbol{d}_i, \boldsymbol{x}_i, \boldsymbol{\beta}_l, \sigma_l^2, \boldsymbol{b}_i)}_{(3.12)} \underbrace{p(\widetilde{T}_i, \delta_i \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{e}_{i,\widetilde{T}_i}, \boldsymbol{\beta}_l, \boldsymbol{b}_i, \boldsymbol{\beta}_s)}_{(3.13)} \Bigg). \end{aligned} \tag{3.8}$$

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 $\boldsymbol{\Sigma}_b$. When conjugacy is unattainable for the visitation and survival parameters, we assume $\beta_{s1}, \beta_{s2}, \beta_{s3}, \beta_{s4}, h_0 \sim \text{Normal}(\beta_{s0}, \sigma_{s0}^2)$, $\eta_{tox} \sim \text{Gamma}(\pi_{s1}, \pi_{s2})$, $\omega \sim \text{Gamma}(\pi_{s3}, \pi_{s4})$, $\mu, \nu_1, \nu_2 \sim \text{Normal}(\beta_{v0}, \sigma_{v0}^2)$, $\boldsymbol{\beta}_\alpha \sim \text{Normal}(\boldsymbol{\beta}_{\alpha 0}, \boldsymbol{\Sigma}_{\boldsymbol{\beta}_\alpha})$, 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 Section 3.7.

3.3 Optimize Personalized Clinical Decision

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

Let $\boldsymbol{\theta} = (\nu_1, \nu_2, \mu, \boldsymbol{\beta}_d, \sigma_d^2)$ 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 $\boldsymbol{\phi} = (\boldsymbol{\beta}_s, \boldsymbol{b}_i, \boldsymbol{\beta}_l, \sigma_l^2)$ denote the set of "observation" parameters, i.e., all other parameters in the joint model (3.8). Our goal is to find, for each patient i, the optimal "policy", i.e., the intensity function and mark distribution with the optimal parameter $\widetilde{\boldsymbol{\theta}}_i$ that maximizes the patient's expected survival time. Note here we have index i for $\widetilde{\boldsymbol{\theta}}_i$ since the optimal parameters may be different for patients with different baseline covariates, yielding personalized optimal clinical decision. We borrow the term "policy" from reinforcement learning (RL) since the setting is similar: in RL, the "policy" refers to the distribution from which an intelligent agent samples its actions and that distribution is optimized to achieve the highest expected reward (Kaelbling et al., 1996; Sutton and Barto, 2018). In the kidney transplantation application, we define a personalized reward function R_i as the log-scaled median survival time to optimize patients' survival: $R_i = \log(\widehat{T}_i)$, where $S_i(\widehat{T}_i) = 0.5$. If desired, other reward functions can be considered. For ex-

ample, 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., $R_i = \log(\widehat{T}_i) + \eta_0 C_i$, where η_0 is a tuning parameter and C_i is the number of visitations.

Without loss of generality, we assume that each patient i receives an arbitrary stochastic reward that is a function of the survival time: $R_i(T_i)$, which depends on the MTPP $\boldsymbol{e}_{i,T_i}$ for clinical decisions, the longitudinal process $\boldsymbol{y}_i$, and the survival time T_i. Formally, denote the expected reward for any patient i to be

$$G_i(\boldsymbol{\theta}) \;=\; \int E_{(\boldsymbol{y}_i, T_i, \boldsymbol{e}_{i,T_i}) \sim p(\boldsymbol{y}_i, T_i, \boldsymbol{e}_{i,T_i} | \boldsymbol{\theta}, \boldsymbol{\phi})}[R_i(T_i)] p(\boldsymbol{\phi} \mid \mathcal{D}) d\boldsymbol{\phi}, \tag{3.9}$$

where $p(\boldsymbol{y}_i, T_i, \boldsymbol{e}_{i,T_i} \mid \boldsymbol{\theta}, \boldsymbol{\phi})$ is the joint distribution of $(\boldsymbol{y}_i, T_i, \boldsymbol{e}_{i,T_i})$ in (3.8), $\mathcal{D}$ denotes the observed data, and $p(\boldsymbol{\phi} \mid \mathcal{D})$ is the posterior distribution of $\boldsymbol{\phi}$. The expectation is taken over all possible realizations of $(\boldsymbol{y}_i, T_i, \boldsymbol{e}_{i,T_i})$. We aim to find the optimal clinical decision, represented by $\widetilde{\boldsymbol{\theta}}_i$, to maximize the expected reward $G_i(\boldsymbol{\theta})$ for patient i after integrating out the uncertainty in the longitudinal process and the survival distribution:

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

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

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

Precisely,

Proposition 3.3.1 *For the i-th patient with baseline covariates $\boldsymbol{x}_i$, given a joint distribution $p(\boldsymbol{y}_i, T_i, \boldsymbol{e}_{i,T_i} \mid \boldsymbol{\theta}, \boldsymbol{\phi})$, the gradient of the expected reward $G_i(\boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}$ is:*

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

where $p(\boldsymbol{e}_{i,T_i} \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{\phi}, \boldsymbol{\theta})$ is the probability of the i-th patient's sequence of visitations and assigned dosages in (3.11).

Proof:

The expected reward function in Proposition 3.3.1 can be re-written as:

$$\begin{aligned}
& E_{(\boldsymbol{y}_i, T_i, \boldsymbol{e}_{i,T_i}) \sim p(\boldsymbol{y}_i, T_i, \boldsymbol{e}_{i,T_i} \mid \boldsymbol{\theta}, \boldsymbol{\phi})} [R_i(T_i)] = E_{J_i} \Big[E_{(\boldsymbol{y}_i, T_i, \boldsymbol{e}_{i,T_i}) \mid J_i} [R_i(T_i) \mid J_i] \Big] \\
= & \sum_J P(J_i = J) \times \Big(\prod_{j=1}^{J} \int_{t_{i,j}, d_{i,j}} \lambda_i(t_{i,j} \mid y_{i,j-1}, \boldsymbol{\theta}, \boldsymbol{\phi}) p(d_{i,j} \mid y_{i,j}, \boldsymbol{x}_i, \boldsymbol{\theta}, \boldsymbol{\phi}) \Big) \\
\times & \exp \Big(- \int_0^{T_i} \lambda_i(t \mid \boldsymbol{y}_i, \boldsymbol{\theta}, \boldsymbol{\phi}) dt \Big) \times \Big(\prod_{j=1}^{J} \int_{y_i(t_{i,j})} p(y_i(t_{i,j}) \mid d_i(t_{i,j-1}), \boldsymbol{x}_i, \boldsymbol{\phi})) \Big) \\
\times & \int_{T_i} p(T_i \mid \boldsymbol{x}_i, \boldsymbol{y}_i, \boldsymbol{e}_{i,T_i}, \boldsymbol{\phi}) R_i(T_i) dT_i \prod_{j=1}^{J} dt_{i,j} dd_{i,j} dy_i(t_{i,j}).
\end{aligned}$$

Here we take the expectation with respect to the event history given the number of events, and then take the expectation with respect to the number of events. Then,

we can compute the gradient of the expected reward as follows:

$$
\begin{aligned}
& \nabla_{\boldsymbol{\theta}} E_{(\boldsymbol{y}_i, T_i, \boldsymbol{e}_{i,T_i}) \sim p(\boldsymbol{y}_i, T_i, \boldsymbol{e}_{i,T_i} | \boldsymbol{\theta}, \boldsymbol{\phi})}[R_i(T_i)] \\
= & \sum_J \nabla_{\boldsymbol{\theta}} \Bigg(P(J_i = J) \times \Big(\prod_{j=1}^{J} \int_{t_{i,j}, d_{i,j}} \lambda_i(t_{i,j} \mid y_{i,j-1}, \boldsymbol{\theta}, \boldsymbol{\phi}) p(d_{i,j} \mid y_{i,j}, \boldsymbol{x}_i, \boldsymbol{\theta}, \boldsymbol{\phi}) \Big) \\
\times & \exp\Big(- \int_0^{T_i} \lambda_i(t \mid \boldsymbol{y}_i, \boldsymbol{\theta}, \boldsymbol{\phi}) dt \Big) \Bigg) \times \Big(\prod_{j=1}^{J} \int_{y_i(t_{i,j})} p(y_i(t_{i,j}) \mid d_i(t_{i,j-1}), \boldsymbol{x}_i, \boldsymbol{\phi})) \Big) \\
\times & \int_{T_i} p(T_i \mid \boldsymbol{x}_i, \boldsymbol{y}_i, \boldsymbol{e}_{i,T_i}, \boldsymbol{\phi}) R_i(T_i) dT_i \prod_{j=1}^{J} dt_{i,j} dd_{i,j} dy_i(t_{i,j}) \\
= & \sum_J \frac{\nabla_{\boldsymbol{\theta}} \left(P(J_i = J) \left(\prod_{j=1}^{J} \int_{t_{i,j}, d_{i,j}} \lambda_i(t_{i,j} | y_{i,j-1}, \boldsymbol{\theta}, \boldsymbol{\phi}) p(d_{i,j} | y_{i,j}, \boldsymbol{x}_i, \boldsymbol{\theta}, \boldsymbol{\phi}) \right) \exp\left(- \int_0^{T_i} \lambda_i(t | \boldsymbol{y}_i, \boldsymbol{\theta}, \boldsymbol{\phi}) dt \right) \right)}{P(J_i = J) \left(\prod_{j=1}^{J} \int_{t_{i,j}, d_{i,j}} \lambda_i(t_{i,j} | y_{i,j-1}, \boldsymbol{\theta}, \boldsymbol{\phi}) p(d_{i,j} | y_{i,j}, \boldsymbol{x}_i, \boldsymbol{\theta}, \boldsymbol{\phi}) \right) \exp\left(- \int_0^{T_i} \lambda_i(t | \boldsymbol{y}_i, \boldsymbol{\theta}, \boldsymbol{\phi}) dt \right)} \\
\times & P(J_i = J) \Big(\prod_{j=1}^{J} \int_{t_{i,j}, d_{i,j}} \lambda_i(t_{i,j} \mid y_{i,j-1}, \boldsymbol{\theta}, \boldsymbol{\phi}) p(d_{i,j} \mid y_{i,j}, \boldsymbol{x}_i, \boldsymbol{\theta}, \boldsymbol{\phi}) \Big) \\
\times & \exp\Big(- \int_0^{T_i} \lambda_i(t \mid \boldsymbol{y}_i, \boldsymbol{\theta}, \boldsymbol{\phi}) dt \Big) \times \Big(\prod_{j=1}^{J} \int_{y_i(t_{i,j})} p(y_i(t_{i,j}) \mid d_i(t_{i,j-1}), \boldsymbol{x}_i, \boldsymbol{\phi}) \Big) \\
\times & \int_{T_i} p(T_i \mid \boldsymbol{x}_i, \boldsymbol{y}_i, \boldsymbol{e}_{i,T_i}, \boldsymbol{\phi}) R_i(T_i) dT_i \prod_{j=1}^{J} dt_{i,j} dd_{i,j} dy_i(t_{i,j}) \\
= & \sum_J \nabla_{\boldsymbol{\theta}} \Bigg(\log \Bigg(P(J_i = J) \times \Big(\prod_{j=1}^{J} \int_{t_{i,j}, d_{i,j}} \lambda_i(t_{i,j} \mid y_{i,j-1}, \boldsymbol{\theta}, \boldsymbol{\phi}) p(d_{i,j} \mid y_{i,j}, \boldsymbol{x}_i, \boldsymbol{\theta}, \boldsymbol{\phi}) \Big) \\
\times & \exp\Big(- \int_0^{T_i} \lambda_i(t \mid \boldsymbol{y}_i, \boldsymbol{\theta}, \boldsymbol{\phi}) dt \Big) \Bigg) \Bigg) \times \Big(\prod_{j=1}^{J} \int_{y_i(t_{i,j})} p(y_i(t_{i,j}) \mid d_i(t_{i,j-1}), \boldsymbol{x}_i, \boldsymbol{\phi}) \Big) \\
\times & \int_{T_i} p(T_i \mid \boldsymbol{x}_i, \boldsymbol{y}_i, \boldsymbol{e}_{i,T_i}, \boldsymbol{\phi}) R_i(T_i) dT_i \prod_{j=1}^{J} dt_{i,j} dd_{i,j} dy_i(t_{i,j}) \\
= & E_{(\boldsymbol{y}_i, T_i, \boldsymbol{e}_{i,T_i}) \sim p(\boldsymbol{y}_i, T_i, \boldsymbol{e}_{i,T_i} | \boldsymbol{\theta}, \boldsymbol{\phi})}[R_i(T_i) \nabla_{\boldsymbol{\theta}} \log p(\boldsymbol{e}_{i,T_i} \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{\theta}, \boldsymbol{\phi})],
\end{aligned}
$$

where

$$\begin{aligned} & \log p(\boldsymbol{e}_{i,T_i} \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{\theta}, \boldsymbol{\phi}) \\ = & \sum_{j=1}^{J} \Big(\log \lambda_i(t_{i,j}) + \log\big(\text{Normal}(d_{i,j}; (1, y_{i,j}, \boldsymbol{x}_i)\boldsymbol{\beta}_d, \sigma_d^2)\big) \Big) - \int_0^{T_i} \lambda_i(t) dt. \end{aligned}$$

■

According to Proposition 3.3.1, in order to compute $\nabla_{\boldsymbol{\theta}} G_i(\boldsymbol{\theta})$, we first need to be able to sample $\boldsymbol{y}_i, T_i, \boldsymbol{e}_{i,T_i}$ from $p(\boldsymbol{y}_i, T_i, \boldsymbol{e}_{i,T_i} \mid \boldsymbol{\theta}, \boldsymbol{\phi})$ and calculate $R_i(T_i)$ 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 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 $\boldsymbol{y}_i, T_i, \boldsymbol{e}_{i,T_i}$, we can easily compute $R_i(T_i)$. We describe the sampling process for a general R_i in Algorithm 2 and for the reward being the log median survival time in Algorithm 3. Comparing these two algorithms reveals computational and variance-reduction advantages of selecting the log median survival time as the reward.

Next we compute the gradient of the log-likelihood of the MTPP, $\nabla_{\boldsymbol{\theta}} \log p(\boldsymbol{e}_{i,T_i} \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{\phi}, \boldsymbol{\theta})$, using the parametrization defined in (3.11). The details are described in Section 3.8. Lastly, we integrate out $\boldsymbol{\phi}$ in computing $\nabla_{\boldsymbol{\theta}} G_i(\boldsymbol{\theta})$ using the Monte

Algorithm 2 Sampling $\boldsymbol{y}_i, T_i, \boldsymbol{e}_{i,T_i}$ from the joint model and computing $R_i(T_i)$

Using the superscript s to denote the simulated data, $\boldsymbol{e}^s_{i,T_i^s} = \{(t^s_{i,0}, d^s_{i,0}), \ldots, (t^s_{i,J_i^s}, d^s_{i,J_i^s})\}$ and $\boldsymbol{y}_i^s = (y_i^s(t^s_{i,1}), \ldots, y_i^s(t^s_{i,J_i^s}))$ denote the simulated follow-up schedules, dosages, and longitudinal data over J_i^s visits until the survival time, T_i^s.

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

Output: $\boldsymbol{y}_i^s$, T_i^s, $\boldsymbol{e}^s_{i,T_i^s}$, $R_i(T_i^s)$

Initialize $j \leftarrow 1$, continue $\leftarrow$ **true**
$t^s_{i,0} \leftarrow 0$
$y_i^s(0) \leftarrow y_{i,0}$
$d^s_{i,0} \leftarrow \text{Normal}((1, y_i^s(0), \boldsymbol{x}_i)\boldsymbol{\beta}_d, \sigma_d^2)$
while continue **do**
  $U_v \leftarrow \text{Uniform}(0,1)$
  Solve for $t^s_{i,j}$: $1 - \exp(-\int_{t^s_{i,j-1}}^{t^s_{i,j}} \lambda_i(x)dx) = U_v$
  $U_s \leftarrow \text{Uniform}(0,1)$
  Solve for T_i^s : $1 - \exp(-\int_{t^s_{i,j-1}}^{T_i^s} h_i(x)dx) = U_s$
  if $T_i^s > t^s_{i,j}$ **then**
    $\boldsymbol{z}_i^s(t^s_{i,j}) \leftarrow (1, d_i^s(t^s_{i,j-1}), \boldsymbol{x}_i, t^s_{i,j}, {t^s_{i,j}}^2)$, $\boldsymbol{r}_i^s(t^s_{i,j}) \leftarrow (1, d_i^s(t^s_{i,j-1}), t^s_{i,j})$
    $y_i^s(t^s_{i,j}) \leftarrow \text{Normal}(\boldsymbol{z}_i^s(t^s_{i,j})\boldsymbol{\beta}_l + \boldsymbol{r}_i^s(t^s_{i,j})\boldsymbol{b}_i, \sigma_l^2)$
    $d^s_{i,j} \leftarrow \text{Normal}((1, y_i^s(t^s_{i,j}), \boldsymbol{x}_i)\boldsymbol{\beta}_d, \sigma_d^2)$
    $j \leftarrow j + 1$
  else
    $J_i^s \leftarrow j - 1$, continue $\leftarrow$ **false**
    $\boldsymbol{e}^s_{i,T_i^s} \leftarrow \{(t^s_{i,0}, d^s_{i,0}), \ldots, (t^s_{i,J_i^s}, d^s_{i,J_i^s})\}$ and $\boldsymbol{y}_i^s \leftarrow (y_i^s(t^s_{i,1}), \ldots, y_i^s(t^s_{i,J_i^s}))$
    Compute $R_i(T_i^s)$
  end if
end while

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:

$$\nabla_{\boldsymbol{\theta}} G_i(\boldsymbol{\theta}) \approx \frac{\sum_{k=1}^{K} E_{(\boldsymbol{y}_i, T_i, \boldsymbol{e}_{i,T_i}) \sim p(\boldsymbol{y}_i, T_i, \boldsymbol{e}_{i,T_i} | \boldsymbol{\theta}, \boldsymbol{\phi}_k)} [R_i(T_i) \nabla_{\boldsymbol{\theta}} \log p(\boldsymbol{e}_{i,T_i} \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{\phi}_k, \boldsymbol{\theta})]}{K}. \tag{3.10}$$

To compute each term of the summation in the numerator of (3.10), we first sample

Algorithm 3 Sampling $\boldsymbol{y}_i, \boldsymbol{e}_{i,\widehat{T}_i}$ from the joint model and computing $R_i = \log(\widehat{T}_i)$

Using the superscript s to denote the simulated data, $\boldsymbol{e}^s_{i,\widehat{T}_i} = \{(t^s_{i,0}, d^s_{i,0}), \ldots, (t^s_{i,J^s_i}, d^s_{i,J^s_i})\}$ and $\boldsymbol{y}^s_i = (y^s_i(t^s_{i,1}), \ldots, y^s_i(t^s_{i,J^s_i}))$ denote the simulated follow-up schedules, dosages, and longitudinal data over J^s_i visits until the median survival time, $\widehat{T}_i$. The simulated survival and hazard functions are denoted $S^s_i(t)$ and $h^s_i(t)$, where $S^s_i(t) = \exp(-\int_0^t h^s_i(u)du)$

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

Output: $\boldsymbol{y}^s_i$, $\widehat{T}_i$, $\boldsymbol{e}^s_{i,\widehat{T}_i}$, R_i

Initialize $j \leftarrow 1$, continue $\leftarrow$ **true**
$t^s_{i,0} \leftarrow 0$
$y^s_i(0) \leftarrow y_{i,0}$
$d^s_{i,0} \leftarrow \text{Normal}((1, y^s_i(0), \boldsymbol{x}_i)\boldsymbol{\beta}_d, \sigma^2_d)$
while continue **do**
  $U_v \leftarrow \text{Uniform}(0, 1)$
  Solve for $t^s_{i,j}$: $1 - \exp(-\int_{t^s_{i,j-1}}^{t^s_{i,j}} \lambda_i(x)dx) = U_v$
  $S^s_i(t_{i,j}) \leftarrow \exp(-\int_0^{t^s_{i,j}} h^s_i(u)du)$
  if $S^s_i(t^s_{i,j}) > 0.5$ **then**
    $\boldsymbol{z}^s_i(t^s_{i,j}) \leftarrow (1, d^s_i(t^s_{i,j-1}), \boldsymbol{x}_i, t^s_{i,j}, {t^s_{i,j}}^2)$, $\boldsymbol{r}^s_i(t^s_{i,j}) \leftarrow (1, d^s_i(t^s_{i,j-1}), t^s_{i,j})$
    $y^s_i(t^s_{i,j}) \leftarrow \text{Normal}(\boldsymbol{z}^s_i(t^s_{i,j})\boldsymbol{\beta}_l + \boldsymbol{r}^s_i(t^s_{i,j})\boldsymbol{b}_i, \sigma^2_l)$
    $d^s_{i,j} \leftarrow \text{Normal}((1, y^s_i(t^s_{i,j}), \boldsymbol{x}_i)\boldsymbol{\beta}_d, \sigma^2_d)$
    $j \leftarrow j + 1$
  else
    $J^s_i \leftarrow j - 1$, continue $\leftarrow$ **false**
    $\boldsymbol{e}^s_{i,\widehat{T}_i} \leftarrow \{(t^s_{i,0}, d^s_{i,0}), \ldots, (t^s_{i,J^s_i}, d^s_{i,J^s_i})\}$ and $\boldsymbol{y}^s_i \leftarrow [y^s_i(t^s_{i,1}), \ldots, y^s_i(t^s_{i,J_i})]$
    Solve for $\widehat{T}_i$: $\exp(-\int_0^{\widehat{T}_i} h^s_i(u)du) = 0.5$
    $R_i \leftarrow \log(\widehat{T}_i)$
  end if
end while

T_i, $\boldsymbol{y}_i$, and $\boldsymbol{e}_{i,T_i}$ from $p(\boldsymbol{y}_i, T_i, \boldsymbol{e}_{i,T_i} \mid \boldsymbol{\theta}, \boldsymbol{\phi}_k)$ using Algorithm 2 to compute $R_i(T_i)$ for each $\boldsymbol{\phi}_k$, then multiply the gradient of the log-probabilities of visitation times and dosages under the MTPP policy. The entire SGD algorithm for finding the optimal parameter $\widetilde{\boldsymbol{\theta}}_i$ is described in Algorithm 4, where $\overline{G_i(\boldsymbol{\theta}_{i,m})}$ denotes the expected reward in iteration m. Note that, in the step 7 of Algorithm 4, 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 (Greensmith et al., 2004; Williams, 1992).

Algorithm 4 Stochastic Gradient Descent for optimizing $\boldsymbol{\theta}$ for patient i

Input $\boldsymbol{\theta}_0$, $\boldsymbol{\phi}_k$ $(k = 1, \dots K)$, $\boldsymbol{x}_i$, $y_{i,0}$.
Output $\widetilde{\boldsymbol{\theta}}_i$

Initialize $\boldsymbol{\theta}_{i,1} \leftarrow \boldsymbol{\theta}_0$
for m:=1 **to** M-1 **do**
  for k:=1 **to** K **do**
    do Algorithm 2($\boldsymbol{\theta}_{i,m}$, $\boldsymbol{\phi}_k$, $\boldsymbol{x}_i$, $y_{i,0}$) to sample $R_{i,k}$, $\boldsymbol{e}_{i,T^s_{i,k}}^{s,k}$, and $\boldsymbol{y}_i^{s,k}$.
  end for
  $\overline{G_i(\boldsymbol{\theta}_{i,m})} \leftarrow \frac{\sum_{k=1}^{K} R_{i,k}}{K}$
  $\nabla_{\boldsymbol{\theta}} G_i(\boldsymbol{\theta}_{i,m}) \leftarrow \frac{\sum_{k=1}^{K}(R_{i,k} - \overline{G_i(\boldsymbol{\theta}_{i,m})}) \nabla_{\boldsymbol{\theta}} \log p(\boldsymbol{e}_{i,T^s_{i,k}}^{s,k} | \boldsymbol{y}_i^{s,k}, \boldsymbol{x}_i, \boldsymbol{\phi}_k, \boldsymbol{\theta}_{i,m}))}{K}$
  $\boldsymbol{\theta}_{i,m+1} \leftarrow \boldsymbol{\theta}_{i,m} + s_{i,m} \nabla_{\boldsymbol{\theta}} G_i(\boldsymbol{\theta}_{i,m})$
end for
$m^* \leftarrow \arg\max_m \overline{G_i(\boldsymbol{\theta}_{i,m})}$
$\widetilde{\boldsymbol{\theta}}_i \leftarrow \boldsymbol{\theta}_{i,m^*}$

3.4 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).

3.4.1 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 $\boldsymbol{x}_i$: 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 transplant, 0 otherwise. For each patient, the donor age and BMI were generated from Normal$(52.5, 15.8^2)$ and Normal$(24.3, 4.5^2)$, 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 $\boldsymbol{\beta}_\alpha = (9.5, -1.5)^T$ so that a higher creatinine level results in a higher visitation intensity; for assigning dosages, the simulated true $\boldsymbol{\beta}_d$ 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 $\boldsymbol{\beta}_l = (5.3, 0.1, 0.3, 0.4, 0.25, -1 \times 10^{-4}, 3 \times 10^{-8})^T$, $\sigma_l = 0.1$, and $\boldsymbol{\Sigma}_b = \begin{bmatrix} 0.04 & 0 & 0 \\ 0 & 0.0049 & 0 \\ 0 & 0 & 10^{-8} \end{bmatrix}$. Note that the last two terms in the simulated true $\boldsymbol{\beta}_l$ 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$(5, 0.1^2)$. In the survival submodel (3.6), we assumed that the simulated true parameters were $h_0 = 5$, $\omega = 1.05$, $\beta_{s1} = 1$, $\beta_\alpha = -5$, $\beta_{s2} = 0.9$, $\beta_{s3} = -0.75$, 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 3.2, we generated the data $\boldsymbol{y}_i, \boldsymbol{e}_{i,\widetilde{T}_i}, \widetilde{T}_i, \delta_i$ for each patient i, $i = 1, \ldots, I$.

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. Figure 3.4 plots the simulated longitudinal creatinine levels and follow-up schedules with dosages for four randomly

selected patients.

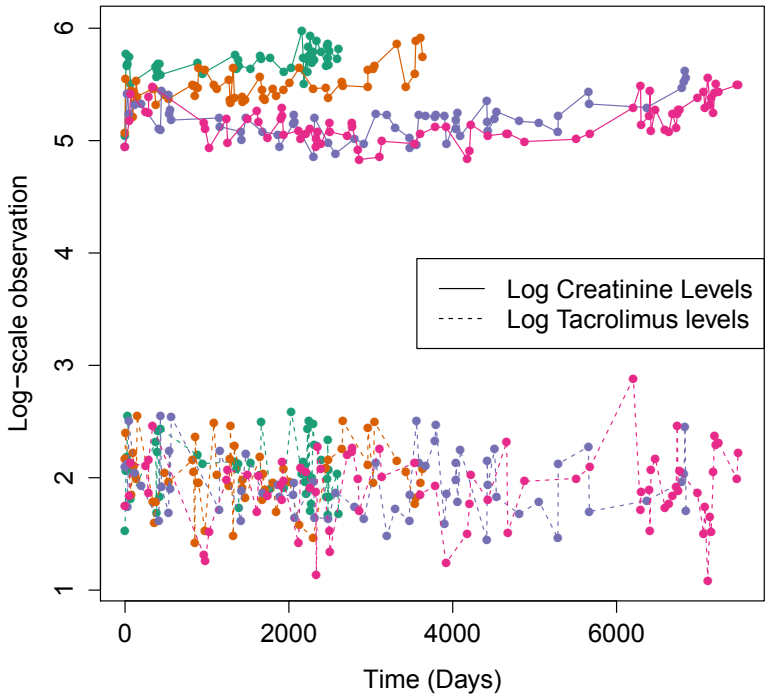

Figure 3.4: Longitudinal measurements (solid lines) and dosages (dashed lines) over time for four randomly selected patients from the simulated dataset. The points represent the visitation times.

3.4.2 Results: model fitting

We applied the proposed Bayesian joint model to the simulated dataset. The hyperparameters were set to be $\boldsymbol{\beta}_{d0} = \boldsymbol{\beta}_{l0} = \boldsymbol{\beta}_{\alpha 0} = \mathbf{0}$, $\boldsymbol{\Sigma}_{\boldsymbol{\beta}_d} = \boldsymbol{\Sigma}_{\boldsymbol{\beta}_l} = \boldsymbol{\Sigma}_{\boldsymbol{\beta}_\alpha} = 100^2 I$, $\pi_{d1} = \pi_{d2} = \pi_{l1} = \pi_{l2} = \pi_{s3} = \pi_{s4} = 0.01$, $\pi_{s1} = \pi_{s2} = 0.01$, $\beta_{s0} = \beta_{v0} = 0$, $\sigma^2_{s0} = \sigma^2_{v0} = 100^2$, $\pi_{v1} = 400$, $\pi_{v2} = 200$. 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 (Figure 3.5), showing no issues of non-convergence. We first report on the performance of the proposed joint model in terms of parameter estimation. Figure 3.6 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. Table 3.1 summarizes the MSE and the standard deviation of squared errors, indicating that the proposed joint model can accurately estimate parameters.

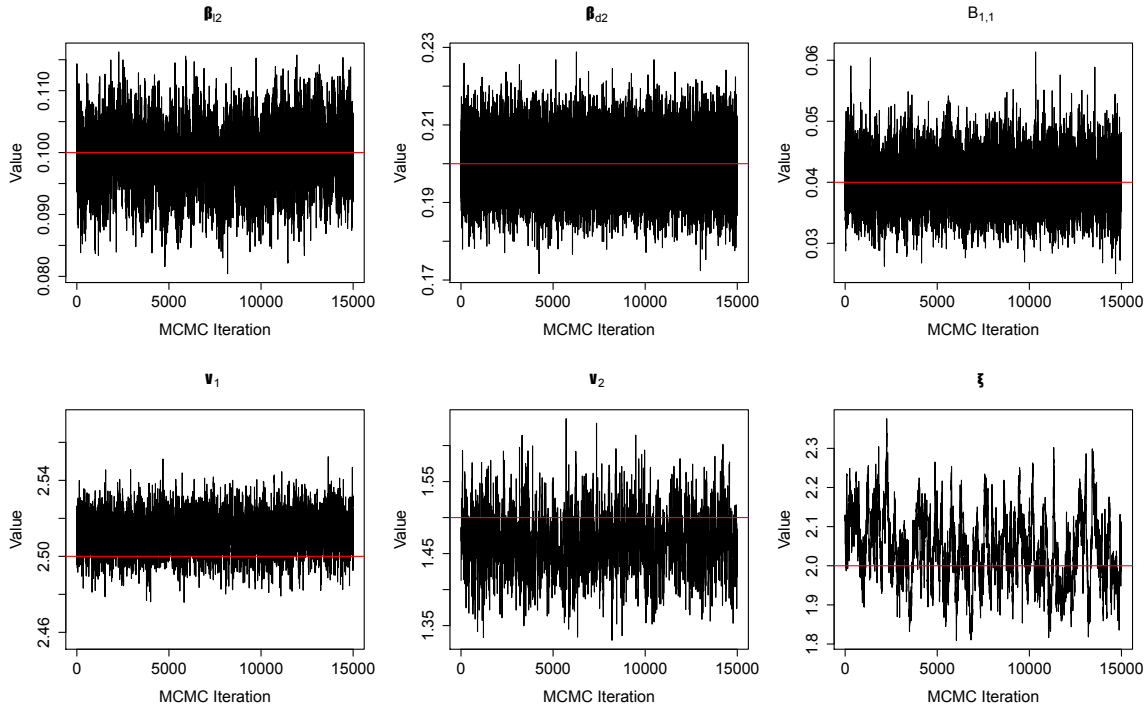

Figure 3.5: Post burn-in trace plots for selected parameters in the simulation study. β_{l2} is the coefficient associated with the dosage in $\boldsymbol{\beta}_l$ and β_{d2} is the coefficient associated with longitudinal measurements in $\boldsymbol{\beta}_d$. $B_{1,1}$ represents the variance in the patient-specific random effect for longitudinal measurements. The red line represents the parameter's simulated truth.

As the proposed model represents the first effort in the literature to jointly model

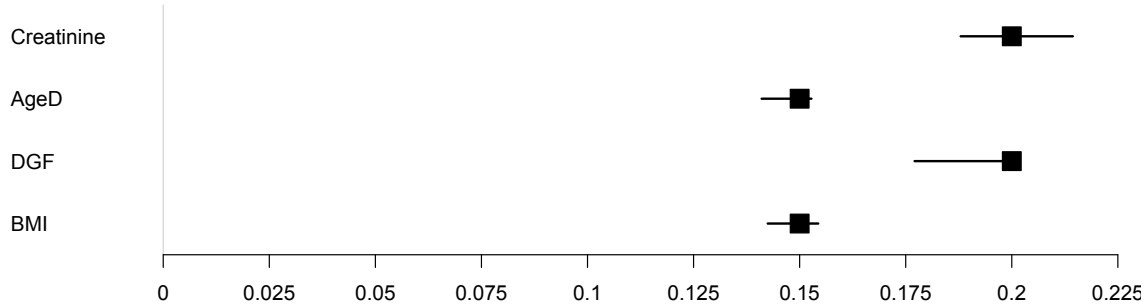

(a) Dosage parameters

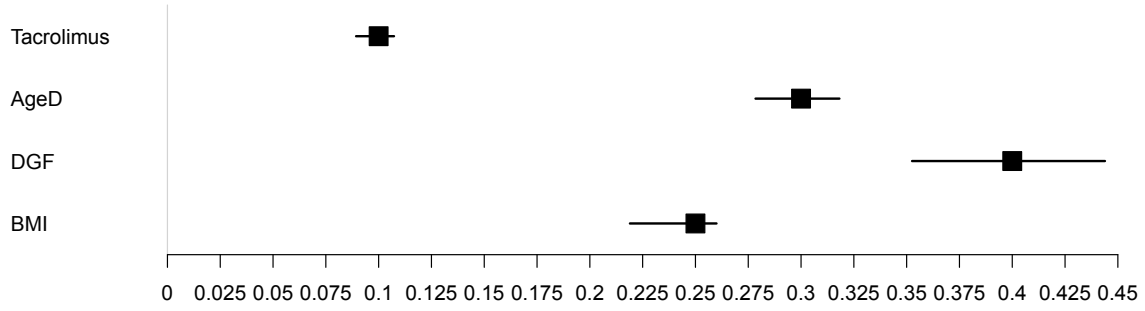

(b) Longitudinal parameters

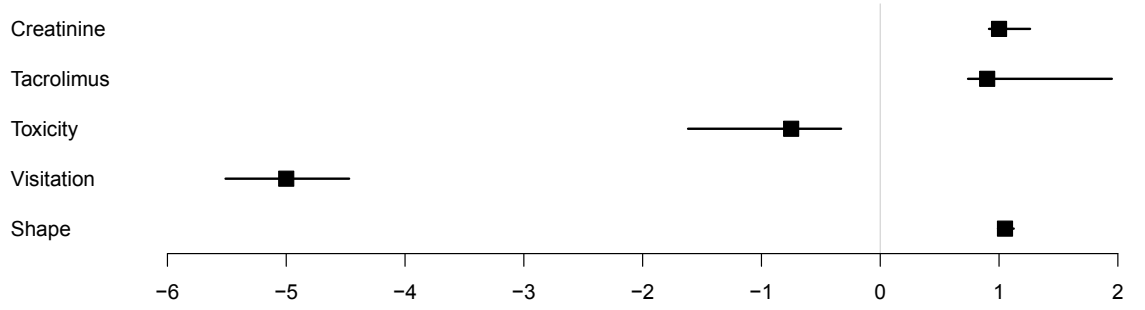

(c) Survival parameters

Figure 3.6: 95% credible intervals for parameters in the dosage, longitudinal, and survival submodels. The dosage and longitudinal values are in log-scale. The squares represent the simulated true values.

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 (SIS)" model that breaks the

	Truth	MSE	SD of squared errors
β_{s1}	1	0.0192	0.0220
β_{s2}	0.9	0.256	0.348
β_{s3}	-0.75	0.168	0.2194
β_{s4}	-5	0.0688	0.0922
h_0	5	0.678	0.704
ω	1.05	1.34×10^{-3}	1.57×10^{-3}
σ_l^2	0.01	6.31×10^{-8}	6.25×10^{-8}
σ_d^2	0.09	1.69×10^{-5}	9.34×10^{-6}
μ	-4.8	5.12×10^{-4}	4.99×10^{-4}
ξ	2	8.99×10^{-3}	1.32×10^{-2}
ν_1	2.5	2.49×10^{-4}	3.11×10^{-4}
ν_2	1.5	3.01×10^{-3}	3.64×10^{-3}

Table 3.1: MSE and standard deviation of squared errors for randomly selected parameters in the simulation study.

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 widely applicable information criterion (WAIC) (Gelman et al., 2014): the joint model has a WAIC value of 226,982 while theSLS 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 3.2 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_{α}	-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 3.2: Parameter estimation under the joint and SLS models.

3.4.3 Results: personalized optimal clinical decision estimation

We applied the proposed policy gradient method in Section 3.3 to the simulated dataset to estimate the personalized optimal clinical decision that maximizes one patient's graft median survival time, i.e., $R_i = \log(\widehat{T}_i)$, where $\widehat{T}_i$ is the median survival time of patient i. The starting parameter values $\boldsymbol{\theta}_0$ in Algorithm 4 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 4 with $M = 1000$ steps to estimate the personalized optimal parameter $\widetilde{\boldsymbol{\theta}}_i$ 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 3.7(a, b) plots the expected mean reward versus SGD iterations. 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 3.3. Recall that the dosage model is $d_{i,j} = (1, y_{i,j}, \boldsymbol{x}_i)\boldsymbol{\beta}_d + \zeta_{i,j}$. Denote $\boldsymbol{\beta}_d = (\beta_{d1}, \beta_{d2}, \ldots, \beta_{dL})^T$, where L is the dimension of $\boldsymbol{\beta}_d$. Since $\boldsymbol{x}_i$ denotes the baseline covariate and does not change over time, we define the personalized dosage intercept to be $\widetilde{\beta}_d = (1, \boldsymbol{x}_i)(\beta_{d1}, \beta_{d3}, \ldots, \beta_{dL})^T$ so that optimizing $\boldsymbol{\beta}_d$ is equivalent to optimizing $(\beta_{d2}, \widetilde{\beta}_d)$. As shown in Table 3.3, the optimized dosage parameters $\widetilde{\beta}_d$ 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 $\widetilde{\beta}_d$ 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."

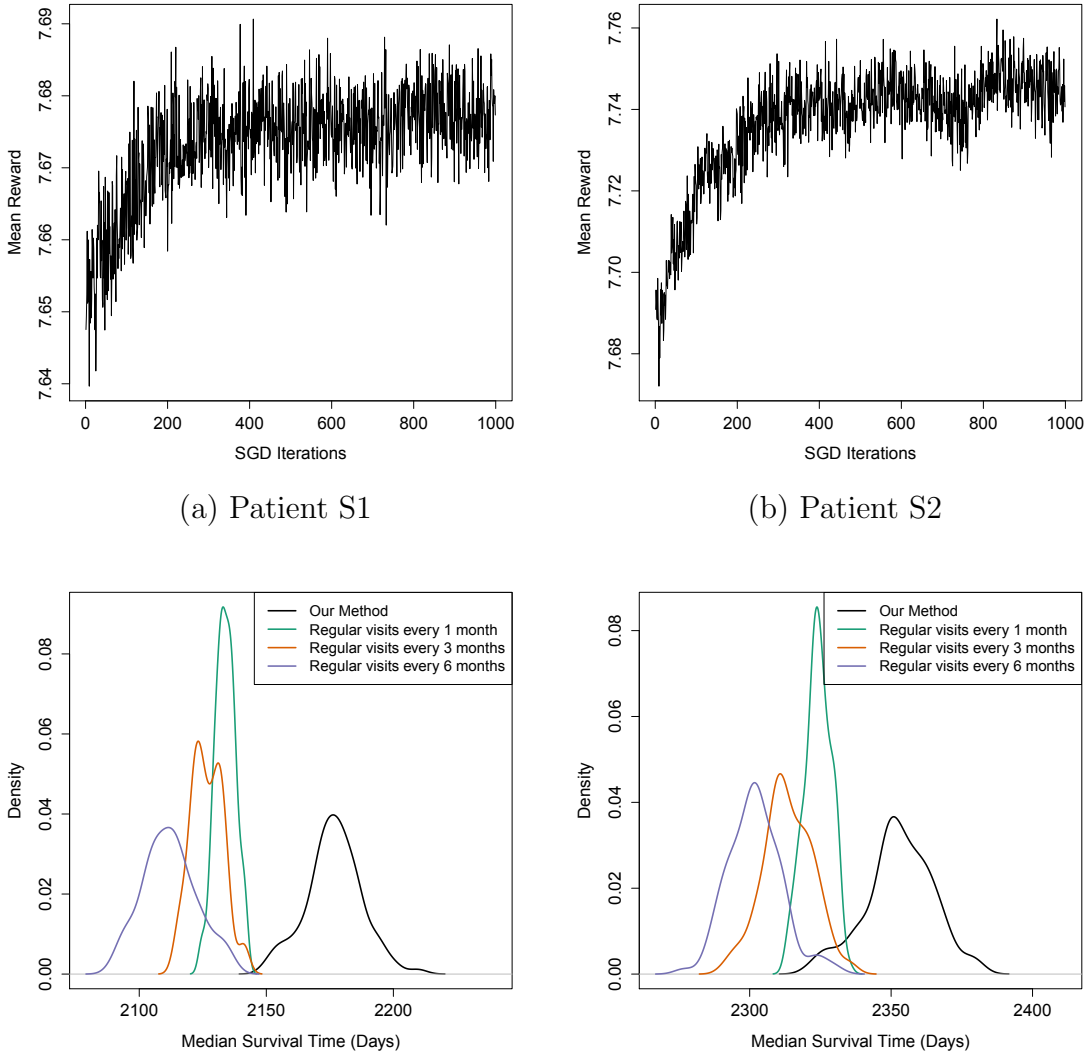

(a) Patient S1

(b) Patient S2

(c) Patient S1

(d) Patient S2

Figure 3.7: 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.

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

Table 3.3: Simulation: Stochastic Gradient Descent Optimal Parameter Results

	$\widetilde{\boldsymbol{\theta}}_0$	$\widetilde{\boldsymbol{\theta}}_{S1}$	$\widetilde{\boldsymbol{\theta}}_{S2}$
$\widetilde{\beta}_d$: 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

were most frequently seen every 1 month or 3 months, depending on their physicians. After the 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 3.7(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, 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.

3.5 Application: DIVAT Data Analysis

We extracted data from Nantes University Hospital Centers in the DIVAT cohort (www.divat.fr), yielding a total of $N = 947$ 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 baseline 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). Table 3.4 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.

3.5.1 Experimental results: model fitting

We first applied the proposed Bayesian joint model to the DIVAT data with $\boldsymbol{x}_i = (\text{AgeD}_i, \text{AgeR}_i, \text{DGF}_i, \text{BMI}_i, \text{Diab}_i, \text{Type}_i)$. 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

Donor age (years)	
Mean ± SD	52.5 ± 15.8
Median	54
Receipient age (years)	
Mean ± SD	51.1 ± 14.3
Median	52
Body mass index (BMI)	
Mean ± SD	24.3 ± 4.5
Median	23.7
Delayed graft function, n(%)	
Yes	329 (34.7%)
No	618 (65.3%)
Diabetes history, n(%)	
Yes	140 (14.5%)
No	807 (85.5%)
Type of donor, n(%)	
Yes	800 (84.5%)
No	147 (15.5%)

Table 3.4: Patient characteristics at baseline immediately after transplantation.

assessed using R package *coda* and the trace plots for randomly selected parameters were shown in Figure 3.8, 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 3.9. Figure 3.9(a) plots posterior means of the linear coefficient $\boldsymbol{\beta}_d$ 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 posi-

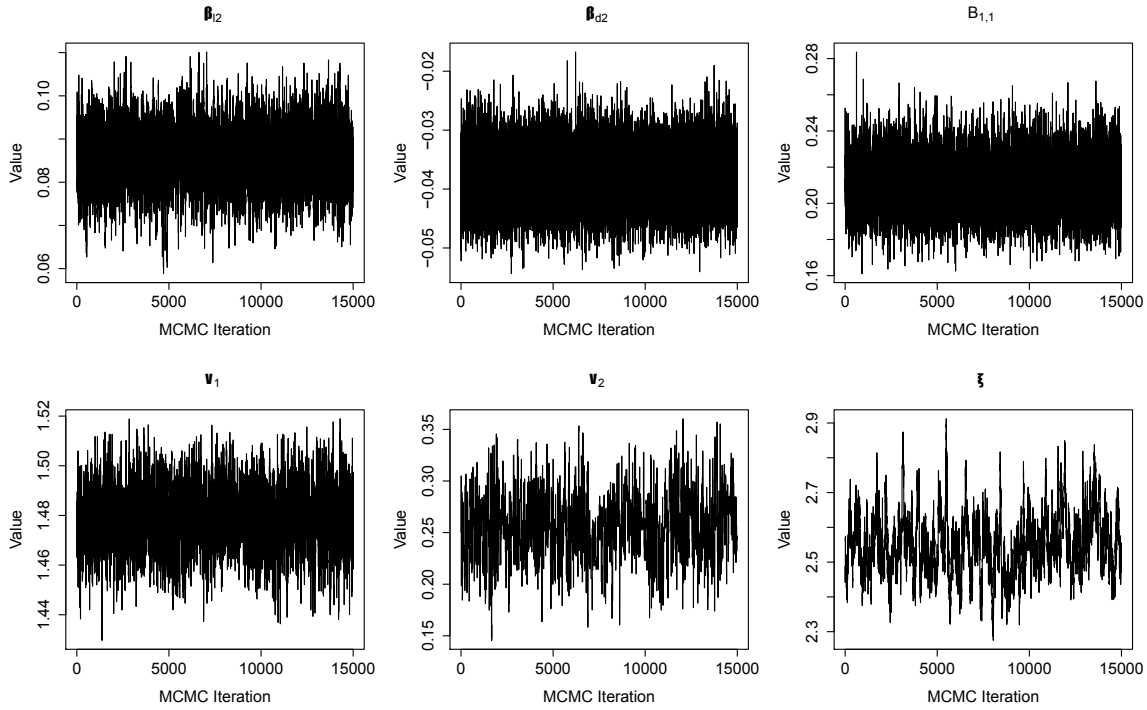

Figure 3.8: Post burn-in trace plots for randomly selected parameters in the DIVAT data analysis. β_{l2} is the coefficient associated with the dosage in $\boldsymbol{\beta}_l$ and β_{d2} is the coefficient associated with longitudinal measurements in $\boldsymbol{\beta}_d$. $B_{1,1}$ represents the variance in the patient-specific random effect for longitudinal measurements.

tively 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 3.9(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

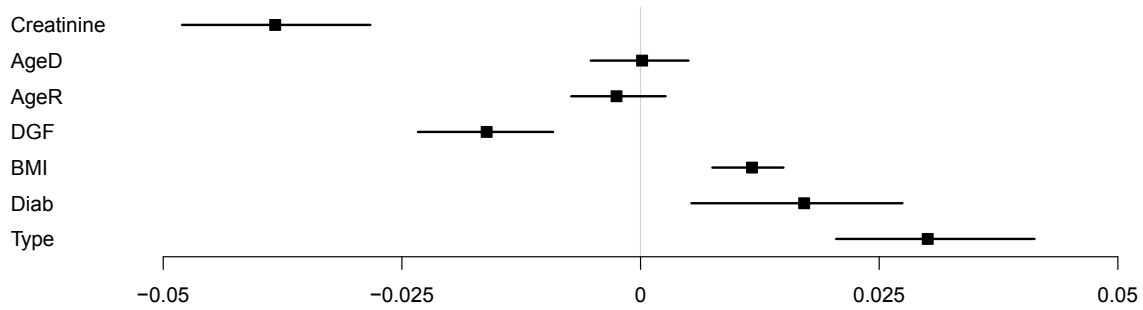

(a) Dosage parameters

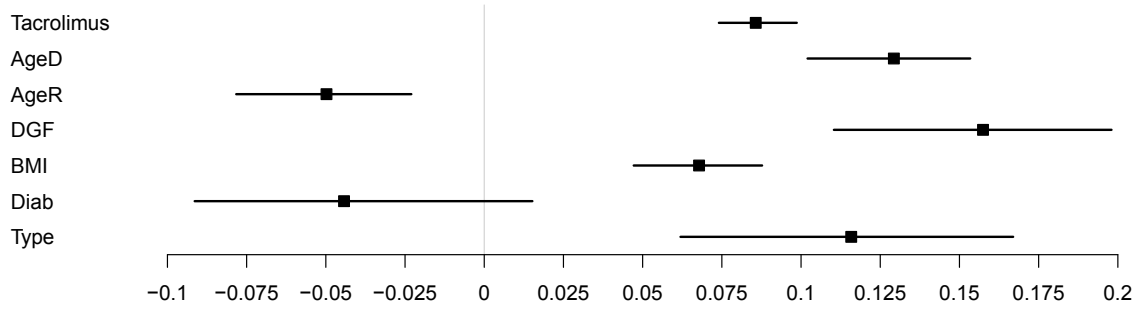

(b) Longitudinal parameters

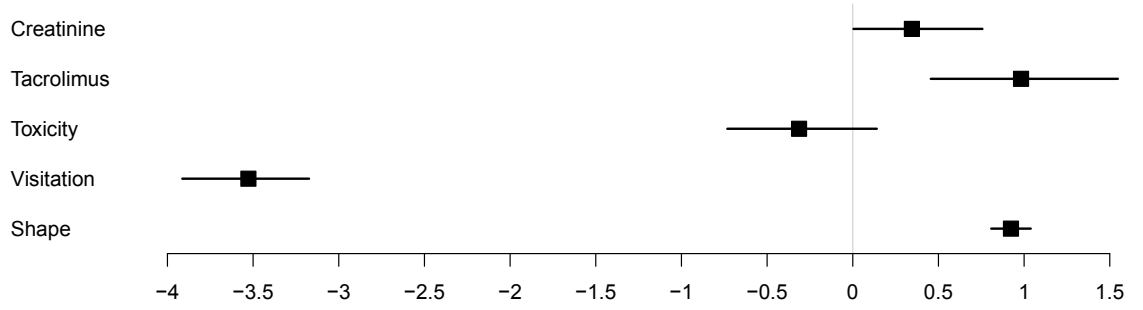

(c) Survival parameters

Figure 3.9: 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.

findings in the literature (Foucher et al., 2016; Gerchman et al., 2009; Katari et al., 1997). 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. Hjelmeseth 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 3.9(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.

3.5.2 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 4 by setting $\boldsymbol{\theta}_0$ to be their posterior means. Algorithm 4 was implemented with $M = 1000$ steps to estimate $\widetilde{\boldsymbol{\theta}}_i$ 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. Figure 3.10 plots the predictive median survival times across SGD iterations for the two patients. Patient R1's 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.

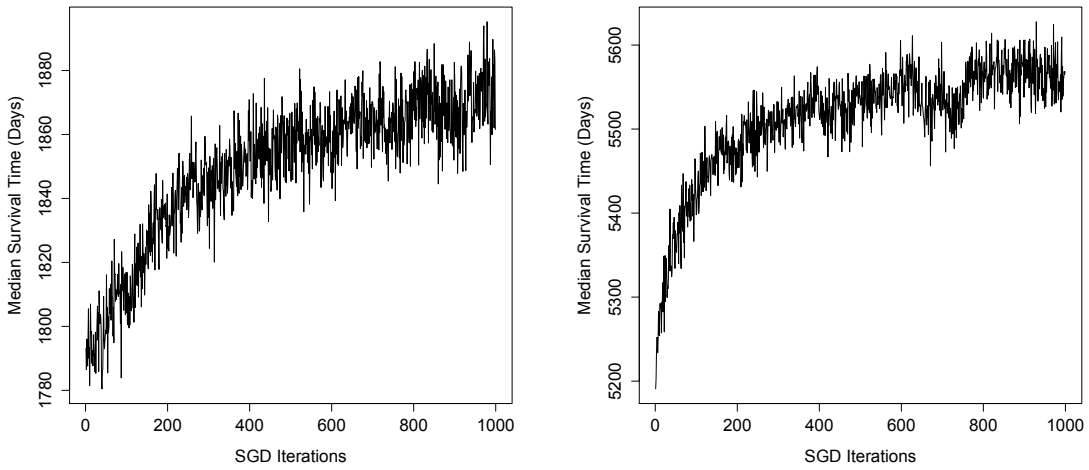

(a) Patient R1 (b) Patient R2

Figure 3.10: The expected mean reward versus SGD iterations for two randomly selected patients R1 and R2.

To further interpret the estimated optimal parameters in clinical decisions, we compared their initial values with the optimized values in Table 3.5. Patient R1's optimal dosage parameters, $\widetilde{\beta}_d$ and β_{d2}, were higher than their posterior means, suggesting that assgining a higher dosage level compared to what the physician actually did for the same creatinine level would improve his/her survival outcome. 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. 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.

Table 3.5: DIVAT data: optimal parameters estimated by the policy-optimizing method.

	$\widetilde{\boldsymbol{\theta}}_0$	$\widetilde{\boldsymbol{\theta}}_{R1}$	$\widetilde{\boldsymbol{\theta}}_{R2}$
$\widetilde{\beta}_d$: 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

3.5.3 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 3.5.2 which used the optimized parameters in both the visitation and dosage models. Figure 3.11 plots boxplots for 100 realizations of the predictive median survival times under each of the four policies. The visitation schedule optimization accounts for more improvement in prolonging the survival for patient R1 compared to patient R2 because, as shown in Table 3.5, 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 3.5.2, suggested slightly fewer visits overall, but more frequent visits when their creatinine levels are high. Comparing Opt. Visits vs. Non-Opt. and Opt. Both vs. Opt. Dosage, we can see that optimizing treatment schedules is clearly beneficial to these patients, thus empirically strengthening the motivation of our work. In summary, this analysis reveals that optimizing both visitation schedules and dosages is necessary to maximize patients' survival.

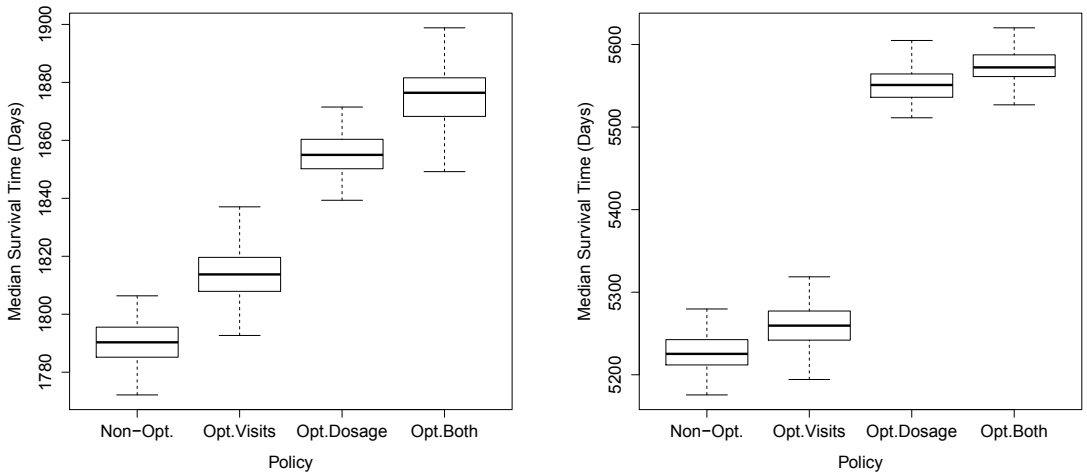

(a) Patient R1 (b) Patient R2

Figure 3.11: The boxplots of the predictive median survival times under different policies of visitation schedules and dosages for patients R1 and R2.

3.6 Conclusion

In this work, we developed a Bayesian method that jointly models the 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 in the Bayesian framework. 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, patients may be

heterogeneous, meaning that patients in different subgroups may respond differently to the treatment or have different disease progression and survival patterns depending on their clinical characteristics. We can extend the proposed Bayesian joint model to account for patients' heterogeneity by adding index i to all parameters in (3.8) and considering a Bayesian nonparametric prior, such as the Dirichlet process (Ferguson, 1973). 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.

3.7 MCMC Sampling Details

3.7.1 Joint Model Summary

Below we first summarize the proposed Bayesian joint model before presenting the MCMC details. The joint probability of clinical observations and decisions factors as:

$$\prod_{i=1}^{I} p(\boldsymbol{y}_i, \boldsymbol{e}_{i,\widetilde{T}_i}, \widetilde{T}_i, \delta_i \mid \boldsymbol{x}_i, \boldsymbol{\beta}_l, \boldsymbol{\beta}_d, \boldsymbol{\beta}_v, \boldsymbol{\beta}_s, \boldsymbol{b}_i, \sigma_l^2, \sigma_d^2)$$

$$\propto \prod_{i=1}^{I} \Bigg(\underbrace{p(\boldsymbol{e}_{i,\widetilde{T}_i} \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{\beta}_v, \boldsymbol{\beta}_d, \sigma_d^2)}_{(3.11)} \underbrace{p(\boldsymbol{y}_i \mid \boldsymbol{d}_i, \boldsymbol{x}_i, \boldsymbol{\beta}_l, \sigma_l^2, \boldsymbol{b}_i)}_{(3.12)} \underbrace{p(\widetilde{T}_i, \delta_i \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{e}_{i,\widetilde{T}_i}, \boldsymbol{\beta}_l, \boldsymbol{b}_i, \boldsymbol{\beta}_s)}_{(3.13)} \Bigg).$$

Here

$$
\begin{aligned}
& p(\boldsymbol{e}_{i,T_i} \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{\beta}_v, \boldsymbol{\beta}_d, \sigma_d^2) \\
= \; & \underbrace{\exp\Big(- \int_0^{T_i} \lambda_i(t \mid \boldsymbol{y}_i, \boldsymbol{\beta}_v) dt \Big)}_{\text{Prob. of no visits at } t \in [0,T_i] \setminus \{t_{i,j}\}_{j=1}^{J_i}} \prod_{j=0}^{J_i} \underbrace{p(d_{i,j} \mid y_{i,j}, \boldsymbol{x}_i, \boldsymbol{\beta}_d, \sigma_d^2)}_{\text{Prob. of dosage}} \prod_{j=1}^{J_i} \underbrace{\lambda_i(t_{i,j} \mid y_{i,j-1}, \boldsymbol{\beta}_v)}_{\text{Prob. of a visit at } t_{i,j}},
\end{aligned}
\tag{3.11}
$$

where $\boldsymbol{y}_i = (y_{i,0}, \ldots, y_{i,J_i})$, $\boldsymbol{\beta}_v = \{\mu, \nu_1, \nu_2, \xi, \boldsymbol{\beta}_\alpha\}$.

$$
p(\boldsymbol{y}_i \mid \boldsymbol{d}_i, \boldsymbol{x}_i, \boldsymbol{\beta}_l, \sigma_l^2, \boldsymbol{b}_i) = \prod_{j=1}^{J_i} p(y_{i,j} \mid t_{i,j}, d_{i,j-1}, \boldsymbol{x}_i, \boldsymbol{\beta}_l, \sigma_l^2, \boldsymbol{b}_i). \tag{3.12}
$$

$$
\begin{aligned}
p(\widetilde{T}_i, \delta_i \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{e}_{i,\widetilde{T}_i}, \boldsymbol{\beta}_l, \boldsymbol{b}_i, \boldsymbol{\beta}_s) \;\; = \;\; & f_i(\widetilde{T}_i \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{e}_{i,\widetilde{T}_i}, \boldsymbol{\beta}_l, \boldsymbol{b}_i, \boldsymbol{\beta}_s)^{\delta_i} \\
& \times S_i(\widetilde{T}_i \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{e}_{i,\widetilde{T}_i}, \boldsymbol{\beta}_l, \boldsymbol{b}_i, \boldsymbol{\beta}_s)^{1-\delta_i},
\end{aligned}
\tag{3.13}
$$

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

3.7.2 MCMC Sampling Steps

1. Update $\boldsymbol{\beta}_l$: The prior for $\boldsymbol{\beta}_l$ is Normal$(\boldsymbol{\beta}_{l0}, \boldsymbol{\Sigma}_{\boldsymbol{\beta}_l})$, so we have

$$
p(\boldsymbol{\beta}_l|\cdot) \propto p(\boldsymbol{\beta}_l) \prod_{i=1}^{I} \Big(p(\widetilde{T}_i, \delta_i \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{e}_{i,\widetilde{T}_i}, \boldsymbol{\beta}_l, \boldsymbol{b}_i, \boldsymbol{\beta}_s) \prod_{j=1}^{J_i} p(y_{i,j} \mid d_{i,j-1}, \boldsymbol{x}_i, \boldsymbol{\beta}_l, \sigma_l^2, \boldsymbol{b}_i) \Big).
$$

We use the Metropolis-Hastings algorithm by considering the following proposal distribution:

$$q(\boldsymbol{\beta}_l|\cdot) \propto p(\boldsymbol{\beta}_l) \prod_{i=1}^{I} \prod_{j=1}^{J_i} p(y_{i,j} \mid d_{i,j-1}, \boldsymbol{x}_i, \boldsymbol{\beta}_l, \sigma_l^2, \boldsymbol{b}_i)$$

$$\propto \text{Normal}(C_{\boldsymbol{\beta}_l} \sum_{i=1}^{I} \sum_{j=1}^{J_i} z_{i,j}^T (y_{i,j} - r_{i,j}\boldsymbol{b}_i)/\sigma_l^2, C_{\boldsymbol{\beta}_l}),$$

where $C_{\boldsymbol{\beta}_l} = \left(\boldsymbol{\Sigma}_{\boldsymbol{\beta}_l}^{-1} + \sum_{i=1}^{I} \sum_{j=1}^{J_i} z_{i,j} z_{i,j}^T/\sigma_l^2 \right)^{-1}$. Note that this proposal distribution is the posterior of $\boldsymbol{\beta}_l$ if we were to ignore the survival submodel. The Metropolis-Hastings acceptance rate is:

$$a(\boldsymbol{\beta}_l^*, \boldsymbol{\beta}_l) = \min(1, \frac{p(\boldsymbol{\beta}_l^*|\cdot) q(\boldsymbol{\beta}_l|\cdot)}{p(\boldsymbol{\beta}_l|\cdot) q(\boldsymbol{\beta}_l^*|\cdot)}) = \min(1, \frac{\prod_{i=1}^{I} p(\widetilde{T}_i, \delta_i \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{e}_{i,\widetilde{T}_i}, \boldsymbol{\beta}_l^*, \boldsymbol{b}_i, \boldsymbol{\beta}_s)}{\prod_{i=1}^{I} p(\widetilde{T}_i, \delta_i \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{e}_{i,\widetilde{T}_i}, \boldsymbol{\beta}_l, \boldsymbol{b}_i, \boldsymbol{\beta}_s)}).$$

2. Update $\boldsymbol{b}_i$: The prior for $\boldsymbol{b}_i$ is Normal$(\mathbf{0}, \boldsymbol{\Sigma}_b)$, so we have

$$p(\boldsymbol{b}_i|\cdot) \propto p(\boldsymbol{\beta}_l) \prod_{i=1}^{I} \Big(p(\widetilde{T}_i, \delta_i \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{e}_{i,\widetilde{T}_i}, \boldsymbol{\beta}_l, \boldsymbol{b}_i, \boldsymbol{\beta}_s) \prod_{j=1}^{J_i} p(y_{i,j} \mid d_{i,j-1}, \boldsymbol{x}_i, \boldsymbol{\beta}_l, \sigma_l^2, \boldsymbol{b}_i) \Big).$$

We use the Metropolis-Hastings algorithm by considering the following proposal distribution:

$$q(\boldsymbol{b}_i|\cdot) \propto p(\boldsymbol{b}_i) \prod_{i=1}^{I} \prod_{j=1}^{J_i} p(y_{i,j} \mid d_{i,j-1}, \boldsymbol{x}_i, \boldsymbol{\beta}_l, \sigma_l^2, \boldsymbol{b}_i)$$

$$\propto \text{Normal}(C_{\boldsymbol{b}_i} \sum_{i=1}^{I} \sum_{j=1}^{J_i} z_{i,j}^T (y_{i,j} - z_{i,j}\boldsymbol{\beta}_l)/\sigma_l^2, C_{\boldsymbol{b}_i}),$$

where $C_{\boldsymbol{b}_i} = \left(\boldsymbol{\Sigma}_b^{-1} + \sum_{i=1}^{I} \sum_{j=1}^{J_i} r_{i,j} r_{i,j}^T / \sigma_l^2 \right)^{-1}$. Note that this proposal distribution is the posterior of $\boldsymbol{b}_i$ if we were to ignore the survival submodel. The Metropolis-Hastings acceptance rate is:

$$a(\boldsymbol{b}_i^*, \boldsymbol{b}_i) = \min(1, \frac{p(\boldsymbol{b}_i^*|\cdot) q(\boldsymbol{b}_i|\cdot)}{p(\boldsymbol{b}_i|\cdot) q(\boldsymbol{b}_i^*|\cdot)}) = \min(1, \frac{p(\widetilde{T}_i, \delta_i \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{e}_{i,\widetilde{T}_i}, \boldsymbol{\beta}_l, \boldsymbol{b}_i^*, \boldsymbol{\beta}_s)}{p(\widetilde{T}_i, \delta_i \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{e}_{i,\widetilde{T}_i}, \boldsymbol{\beta}_l, \boldsymbol{b}_i, \boldsymbol{\beta}_s)}).$$

3. Update σ_l^2: The prior for σ_l^2 is InverseGamma(π_{l1}, π_{l2}), so we have

$$\begin{aligned} p(\sigma_l^2|\cdot) &\propto p(\sigma_l^2) \prod_{i=1}^{I} p(\boldsymbol{y}_i|\boldsymbol{y}_i^*, \sigma_l^2) \\ &\propto \text{InverseGamma}(\pi_{l1} + \sum_{i=1}^{I} \frac{J_i}{2}, \pi_{l2} + \sum_{i=1}^{I} \sum_{j=1}^{J_i} \frac{(y_{i,j} - y_{i,j}^*)^2}{2}). \end{aligned}$$

4. Update $\boldsymbol{\Sigma}_b$: We assume a flat prior for $\boldsymbol{\Sigma}_b$ is flat, so we have

$$p(\boldsymbol{\Sigma}_b|\cdot) \propto p(\boldsymbol{\Sigma}_b) \prod_{i=1}^{I} p(\boldsymbol{b}_i|\boldsymbol{\Sigma}_b) \propto \text{InverseWishart}(I, \sum_{i=1}^{I} \boldsymbol{b}_i \boldsymbol{b}_i^T).$$

5. Update $\beta_s^* \in \{\beta_{s1}, \beta_{s2}, \beta_{s3}, \beta_{s4}, h_0\}$: since their posterior distributions are not in closed-form, these parameters are updated using the Metropolis-Hastings. The prior for β_s^* is Normal$(\beta_{s0}, \sigma_{s0}^2)$, so we have

$$p(\beta_s^*|\cdot) \propto p(\beta_s^*) \prod_{i=1}^{I} p(\widetilde{T}_i, \delta_i \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{e}_{i,\widetilde{T}_i}, \boldsymbol{\beta}_l, \boldsymbol{b}_i, \boldsymbol{\beta}_s).$$

6. Update η_{tox}: We use the Metropolis-Hastings to update η_{tox}. The prior for η_{tox}

is Gamma(π_{s1}, π_{s2}), so we have

$$p(\eta_{tox}|\cdot) \propto p(\eta_{tox}) \prod_{i=1}^{I} p(\widetilde{T}_i, \delta_i \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{e}_{i,\widetilde{T}_i}, \boldsymbol{\beta}_l, \boldsymbol{b}_i, \boldsymbol{\beta}_s).$$

7. Update ω: We use the Metropolis-Hastings to update ω. The prior for ω is Gamma(π_{s3}, π_{s4}), so we have

$$p(\omega|\cdot) \propto p(\omega) \prod_{i=1}^{I} p(\widetilde{T}_i, \delta_i \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{e}_{i,\widetilde{T}_i}, \boldsymbol{\beta}_l, \boldsymbol{b}_i, \boldsymbol{\beta}_s).$$

8. Update $\beta_v^* \in \{\mu, \nu_1, \nu_2\}$: We use the Metropolis-Hastings algorithm. The prior for β_v^* is Normal($\beta_{v0}, \sigma_{v0}^2$).

$$p(\beta_v^*|\cdot) \propto p(\beta_v^*) \prod_{i=1}^{I} p(\boldsymbol{e}_{i,T_i} \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{\beta}_v, \boldsymbol{\beta}_d, \sigma_d^2).$$

9. Update $\boldsymbol{\beta}_\alpha$: We use the Metropolis-Hastings algorithm. The prior for $\boldsymbol{\beta}_\alpha$ is Normal($\boldsymbol{\beta}_{\alpha 0}, \boldsymbol{\Sigma}_{\boldsymbol{\beta}_\alpha}$).

$$p(\boldsymbol{\beta}_\alpha|\cdot) \propto p(\boldsymbol{\beta}_\alpha) \prod_{i=1}^{I} \prod_{j=1}^{J_i} p(\boldsymbol{e}_{i,T_i} \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{\beta}_v, \boldsymbol{\beta}_d, \sigma_d^2) p(\widetilde{T}_i, \delta_i \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{e}_{i,\widetilde{T}_i}, \boldsymbol{\beta}_l, \boldsymbol{b}_i, \boldsymbol{\beta}_s).$$

10. Update ξ: We use the Metropolis-Hastings algorithm. The prior for ξ is Gamma(π_{v1}, π_{v2}).

$$p(\boldsymbol{\beta}_\alpha|\cdot) \propto p(\boldsymbol{\beta}_\alpha) \prod_{i=1}^{I} \prod_{j=1}^{J_i} p(\boldsymbol{e}_{i,T_i} \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{\beta}_v, \boldsymbol{\beta}_d, \sigma_d^2) p(\widetilde{T}_i, \delta_i \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{e}_{i,\widetilde{T}_i}, \boldsymbol{\beta}_l, \boldsymbol{b}_i, \boldsymbol{\beta}_s).$$

11. Update $\boldsymbol{\beta}_d$: The prior for $\boldsymbol{\beta}_d$ is Normal$(\boldsymbol{\beta}_{d0}, \boldsymbol{\Sigma}_{\boldsymbol{\beta}_d})$, so we have

$$\begin{aligned} p(\boldsymbol{\beta}_d|\cdot) &\propto p(\boldsymbol{\beta}_d)\prod_{i=1}^{I}\prod_{j=1}^{J_i} p(d_{i,j} \mid y_{i,j}, \boldsymbol{x}_i, \boldsymbol{\beta}_d, \sigma_d^2) \\ &\propto \text{Normal}(C_{\boldsymbol{\beta}_d}\sum_{i=1}^{I}\sum_{j=1}^{J_i}\boldsymbol{a}_{i,j}^T d_{i,j}/\sigma_d^2, C_{\boldsymbol{\beta}_d}), \end{aligned}$$

where $\boldsymbol{a}_{i,j} = (1, y_{i,j}, \boldsymbol{x}_i)$ and $C_{\boldsymbol{\beta}_d} = \left(\Sigma_{\boldsymbol{\beta}_d}^{-1} + \sum_{i=1}^{I}\sum_{j=1}^{J_i}\boldsymbol{a}_{i,j}\boldsymbol{a}_{i,j}^T/\sigma_d^2\right)^{-1}$.

12. Update σ_d^2: The prior for σ_d^2 is InverseGamma(π_{d1}, π_{d2}), so we have

$$\begin{aligned} p(\sigma_d^2|\cdot) &\propto p(\sigma_d^2)\prod_{i=1}^{I}\prod_{j=1}^{J_i} p(d_{i,j} \mid y_{i,j}, \boldsymbol{x}_i, \boldsymbol{\beta}_d, \sigma_d^2) \\ &\propto \text{InverseGamma}(\pi_{d1} + \sum_{i=1}^{I}\frac{J_i}{2}, \pi_{d2} + \sum_{i=1}^{I}\sum_{j=1}^{J_i}\frac{(\boldsymbol{d}_{i,j} - \boldsymbol{a}_{i,j}\boldsymbol{\beta}_d)^2}{2}). \end{aligned}$$

3.8 Gradient Computation Details

To calculate $\nabla_{\boldsymbol{\theta}} \log p(\boldsymbol{e}_{i,T_i} \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{\theta}, \boldsymbol{\phi})$, we compute the gradient with respect to the five parameters in $\boldsymbol{\theta} = (\nu_1, \nu_2, \mu, \boldsymbol{\beta}_d, \sigma_d^2)$.

1.

$$\begin{aligned}
&\frac{\partial \log p(\boldsymbol{e}_{i,T_i} \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{\theta}, \boldsymbol{\phi})}{\partial \boldsymbol{\beta}_d} \\
= \;& \frac{\partial}{\partial \boldsymbol{\beta}_d} \frac{1}{2} \Big(- Jlog(2\pi) - Jlog(\sigma_d^2) - 1/\sigma_d^2 \sum_{j=1}^{J} (d_{i,j} - a_{i,j}\boldsymbol{\beta}_d)^2 \Big) \\
= \;& \frac{\partial}{\partial \boldsymbol{\beta}_d} \Big(- \frac{1}{2\sigma_d^2} \sum_{j=1}^{J} (-2 d_{i,j} a_{i,j} \boldsymbol{\beta}_d + (a_{i,j}\boldsymbol{\beta}_d)^2) \Big).
\end{aligned}$$

Let $\beta_{d,l}$ and $a_{i,j,l}$ refer to the l-th dimension of $\boldsymbol{\beta}_d$ and $a_{i,j} = (1, y_{i,j}, \boldsymbol{x}_i)$. Also, let $\beta_{d,-l}$ and $a_{i,j,-l}$ refer to the other dimensions, excluding the l-th dimension of $\boldsymbol{\beta}_d$ and $a_{i,j}$. Then, we have

$$\begin{aligned}
&\frac{\partial \log p(\boldsymbol{e}_{i,T_i} \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{\theta}, \boldsymbol{\phi})}{\partial \beta_{d,l}} = \frac{\partial}{\partial \beta_{d,l}} \Big(- \frac{1}{2\sigma_d^2} \sum_{j=1}^{J} (-2 d_{i,j} a_{i,j,l} \beta_{d,l} + (a_{i,j,l}\beta_{d,l})^2) \Big) \\
= \;& \sum_{j=1}^{J} \frac{d_{i,j} a_{i,j,l} - a_{i,j,l}^2 \beta_{d,l}}{\sigma_d^2}.
\end{aligned}$$

2.

$$\begin{aligned}
&\frac{\partial \log p(\boldsymbol{e}_{i,T_i} \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{\theta}, \boldsymbol{\phi})}{\partial \sigma_d} \\
= \;& \frac{\partial}{\partial \sigma_d} \frac{1}{2} \Big(- Jlog(2\pi) - Jlog(\sigma_d^2) - 1/\sigma_d^2 \sum_{j=1}^{J} (d_{i,j} - a_{i,j}\boldsymbol{\beta}_d)^2 \Big) \\
= \;& -\frac{J}{\sigma_d} + \frac{\sum_{j=1}^{J} (d_{i,j} - a_{i,j}\boldsymbol{\beta}_d)}{\sigma_d^3}.
\end{aligned}$$

3.

$$\begin{aligned}
&\frac{\partial \log p(\boldsymbol{e}_{i,T_i} \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{\theta}, \boldsymbol{\phi})}{\partial \mu} = - \int_0^{t_{i,J_i}} \frac{\partial \lambda_i(t)}{\partial \mu} dt + \sum_{j=1}^{J} \frac{\partial \log \lambda_i(t_{i,j})}{\partial \mu}. \\
= \;& - \int_0^{t_{i,J_i}} \frac{\partial \lambda_i(t)}{\partial \mu} dt + \sum_{j=1}^{J} \frac{1}{\lambda_i(t_{i,j})} \frac{\partial \lambda_i(t_{i,j})}{\partial \mu} \\
= \;& - \int_0^{t_{i,J_i}} \exp(\mu) dt + \sum_{j=1}^{J} \frac{\exp(\mu)}{\lambda_i(t_{i,j})} = -t_{i,J_i} \exp(\mu) + \sum_{j=1}^{J} \frac{\exp(\mu)}{\lambda_i(t_{i,j})}.
\end{aligned}$$

4.

$$\frac{\partial \log p(\boldsymbol{e}_{i,T_i} \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{\theta}, \boldsymbol{\phi})}{\partial \nu_2} = - \int_0^{t_{i,J_i}} \frac{\partial \lambda_i(t)}{\partial \nu_2} dt + \sum_{j=1}^{J} \frac{1}{\lambda_i(t_{i,j})} \frac{\partial \lambda_i(t_{i,j})}{\partial \nu_2}.$$

For $t \in (t_{i,j}, t_{i,j+1}]$ and $\boldsymbol{V}_{i,j} = (1, y_{i,j})$:

$$
\begin{aligned}
\frac{\partial \lambda_i(t)}{\partial \nu_2} &= \frac{\partial}{\partial \nu_2}\Big(\frac{2(t-t_{i,j})^{\kappa-1} e^{-(\kappa-1)(t-t_{i,j})/\exp(\nu_1)} \frac{((\kappa-1)/\exp(\nu_1))^{\kappa}}{\Gamma(\kappa)}}{1+\exp(\boldsymbol{V}_{i,j}\boldsymbol{\beta}_\alpha)}\Big) \\
= & \frac{\partial}{\partial \nu_2}\Big((t-t_{i,j})^{\exp(\nu_2)} \frac{e^{-\exp(\nu_2)(t-t_{i,j})/\exp(\nu_1)} (\exp(\nu_2)/\exp(\nu_1))^{\exp(\nu_2)+1}}{\Gamma(\exp(\nu_2)+1)}\Big) \\
\times & \Big(\frac{2}{1+\exp(\boldsymbol{V}_{i,j}\boldsymbol{\beta}_\alpha)}\Big) \\
= & \frac{\partial}{\partial \nu_2}\Big(\Big(\frac{\exp(\nu_2)}{\exp(\nu_1)}\Big)^{\exp(\nu_2)+1}(t-t_{i,j})^{\exp(\nu_2)} \frac{e^{-\exp(\nu_2)(t-t_{i,j})/\exp(\nu_1)}}{\Gamma(\exp(\nu_2)+1)}\Big) \times \Big(\frac{2}{1+\exp(\boldsymbol{V}_{i,j}\boldsymbol{\beta}_\alpha)}\Big) \\
= & \Bigg(\Bigg(\Big((\exp(\nu_2)(\nu_2 - \log(\exp(\nu_1))) + \exp(\nu_2) + 1)(t-t_{i,j})^{\exp(\nu_2)} e^{-\exp(\nu_2)(t-t_{i,j})/\exp(\nu_1)} \\
+ & e^{\nu_2 - \exp(\nu_2)(t-t_{i,j})/\exp(\nu_1)} \log(t-t_{i,j})(t-t_{i,j})^{\exp(\nu_2)} \\
- & e^{\nu_2 - \exp(\nu_2)(t-t_{i,j})/\exp(\nu_1)}(t-t_{i,j})^{\exp(\nu_2)+1}/\exp(\nu_1)\Big) \frac{\big(\frac{\exp(\nu_2)}{\exp(\nu_1)}\big)^{\exp(\nu_2)+1}}{\Gamma(\exp(\nu_2)+1)}\Bigg) \\
- & \Big(\frac{\exp(\nu_2)}{\exp(\nu_1)}\Big)^{\exp(\nu_2)+1}(t-t_{i,j})^{\exp(\nu_2)} \frac{e^{-\exp(\nu_2)(t-t_{i,j})/\exp(\nu_1)} \phi(\exp(\nu_2)+1))}{\Gamma(\exp(\nu_2)+1)}\Bigg) \\
\times & \Big(\frac{2}{1+\exp(\boldsymbol{V}_{i,j}\boldsymbol{\beta}_\alpha)}\Big).
\end{aligned}
$$

5.

$$
\frac{\partial \log p(\boldsymbol{e}_{i,T_i} \mid \boldsymbol{y}_i, \boldsymbol{x}_i, \boldsymbol{\theta}, \boldsymbol{\phi})}{\partial \nu_1} = -\int_0^{t_{i,J_i}} \frac{\partial \lambda_i(t)}{\partial \nu_1} dt + \sum_{j=1}^{J} \frac{1}{\lambda_i(t_{i,j})} \frac{\partial \lambda_i(t_{i,j})}{\partial \nu_1}.
$$

For $t \in (t_{i,j}, t_{i,j+1}]$ and $\boldsymbol{V}_{i,j} = (1, y_{i,j})$:

$$
\begin{aligned}
\frac{\partial \lambda_i(t)}{\partial \nu_1} &= \frac{\partial}{\partial \nu_1}\Big(\frac{2(t-t_{i,j})^{\kappa-1} e^{-(\kappa-1)(t-t_{i,j})/\exp(\nu_1)} \frac{((\kappa-1)/\exp(\nu_1))^{\kappa}}{\Gamma(\kappa)}}{1+\exp(\boldsymbol{V}_{i,j}\boldsymbol{\beta}_{\alpha})}\Big) \\
&= \frac{\partial}{\partial \nu_1}\Big(\frac{2(t-t_{i,j})^{\kappa-1} e^{-(\kappa-1)(t-t_{i,j})/\exp(\nu_1)} \frac{((\kappa-1)/\exp(\nu_1))^{\kappa}}{\Gamma(\kappa)}}{1+\exp(\boldsymbol{V}_{i,j}\boldsymbol{\beta}_{\alpha})}\Big) \\
&= \frac{\partial}{\partial \nu_1}\Big(e^{-(\kappa-1)(t-t_{i,j})/\exp(\nu_1)} \exp(\nu_1)^{-\kappa}\Big)\Big(\frac{2(t-t_{i,j})^{\kappa-1} \frac{((\kappa-1))^{\kappa}}{\Gamma(\kappa)}}{1+\exp(\boldsymbol{V}_{i,j}\boldsymbol{\beta}_{\alpha})}\Big) \\
&= \frac{\partial}{\partial \nu_1}\Big(e^{-(\kappa-1)(t-t_{i,j})/\exp(\nu_1)} e^{-\nu_1\kappa}\Big)\Big(\frac{2(t-t_{i,j})^{\kappa-1} \frac{((\kappa-1))^{\kappa}}{\Gamma(\kappa)}}{1+\exp(\boldsymbol{V}_{i,j}\boldsymbol{\beta}_{\alpha})}\Big) \\
&= \frac{\partial}{\partial \nu_1}\Big(e^{-(\kappa-1)(t-t_{i,j})/\exp(\nu_1)-\nu_1\kappa}\Big)\Big(\frac{2(t-t_{i,j})^{\kappa-1} \frac{((\kappa-1))^{\kappa}}{\Gamma(\kappa)}}{1+\exp(\boldsymbol{V}_{i,j}\boldsymbol{\beta}_{\alpha})}\Big) \\
&= \Big(\frac{(\kappa-1)(t-t_{i,j})}{\exp(\nu_1)} - \kappa\Big) e^{-(\kappa-1)(t-t_{i,j})/\exp(\nu_1)-\nu_1\kappa} \Big(\frac{2(t-t_{i,j})^{\kappa-1} \frac{((\kappa-1))^{\kappa}}{\Gamma(\kappa)}}{1+\exp(\boldsymbol{V}_{i,j}\boldsymbol{\beta}_{\alpha})}\Big).
\end{aligned}
$$

Chapter 4

A Bayesian semiparametric model for learning biomarker trajectories and change points in Alzheimer's disease

4.1 Introduction

For many diseases, the manifestation of symptoms can be difficult to identify due to its slow development. As a result, many subjects that receive a diagnosis for these gradually developing diseases have already been experiencing its negative consequences for a substantial time. These detrimental effects can be subtly observed through longitudinal biomarker data that tracks the subject's well-being over time. In contrast to these "susceptible" or "uncured" subjects, others may never develop symptoms nor receive a diagnosis, and we typically refer to such subjects as being

"immune" or "cured".

In this work, we propose a Bayesian framework that addresses several key biomedical questions in such settings. First, we want to understand how the longitudinal biomarker patterns vary based on subject baseline characteristics and inherent disease cure status in order to compute the personalized probability of an inherent cure. For the uncured subjects, we are interested in how the diagnosis time depends on baseline characteristics and whether we can identify a change point in the evolution of the biomarker before diagnosis.

There is substantial literature on change point modeling for early detection of slowly progressing diseases. Hall et al. (2000) introduces a parametric change point model that detects when the rate of cognitive decline accelerates in people who eventually are diagnosed with dementia. Tang et al. (2017) further develops a two-phase joint longitudinal and survival model where the change point is measured relative to the survival event of manifestation of Alzheimer's disease (AD). Within the context of Alzheimer's disease, we will refer to the susceptible/uncured subjects as non-progressors and the immune/cured subjects as progressors to reflect whether or not subjects will eventually progress to developing AD. Although these models are able to robustly detect change points, they do not account for the possibility of disease non-progression. Addressing this possibility is difficult because non-progression cannot be directly observed and we cannot determine non-progression using only baseline characteristics. However, we can narrow down the subjects eligible for being considered non-progressors to subjects with right-censored diagnosis times, where the study ends before diagnosis. By hypothesizing that the progression patterns of the biomarker vary depending on whether subjects will eventually be diagnosed, we can make use of the longitudinal biomarker data to help determine inherent cure status

among the censored subjects.

Embedding a binary cure submodel, which assumes some subjects will never experience the survival event, into a joint longitudinal survival framework allows us to group subjects by whether they are non-progressors or progressors. Yu et al. (2004) developed a joint longitudinal-survival-cure model for the prostate cancer, where the survival event was clinical recurrence and the cured fraction is modeled as a logistic function of baseline covariates.

In this work, we develop a flexible semiparametric Bayesian framework that jointly models longitudinal biomarkers, individual cure status, and the disease diagnosis time by inferring the change point for AD progression before diagnosis. We apply our framework to an Alzheimer's disease context, a prime example of a disease with slowly developing symptoms. Many subjects who are diagnosed with AD are already cognitively impaired but these detrimental effects are only observable through certain longitudinal biomarkers that are difficult to track due to cost and invasiveness. Our data is collected through the BIOCARD study, which tracks cognitive related biomarkers for subjects that began as cognitively normal.

4.2 Bayesian Joint Model

In this section, we describe our proposed semiparametric Bayesian framework that jointly models longitudinal biomarker observations with a change point, diagnosis event times, and cure status. Section 4.2.1 introduces the cure fraction model for disease progressor/non-progressor status, section 4.2.2 introduces the time-to-event model for diagnosis times, and section 4.2.3 introduces the longitudinal model for biomarker measurements conditional on latent variables. Section 4.2.4 discusses the

nonparametric Dirichlet Process Mixture model that allows subjects to share certain structures and parameters within clusters.

4.2.1 Cure Fraction Incidence Model

Cure models are applied in survival analysis when event time data comes from a mixed population composed of two groups: subjects who will eventually progress and others who will never progress to the event of interest (Sy and Taylor, 2000). The primary clinical event of interest in our context is the diagnosis event time, but the survival event of subject death is also relevant as a censoring mechanism. For subject i $(i = 1, \ldots, I)$, let U_i denote the diagnosis event time, C_i denote the administrative censoring time, and D_i denote the survival event time. We assume that these three values are independent, and define the survival event time as $T_i = \min(C_i, D_i)$ and $c_i = \mathbb{I}_{(D_i \leq C_i)}$ to be the survival censoring indicator. If we denote censoring for the diagnosis event as $\delta_i = \mathbb{I}_{(U_i \leq T_i)}$, then U_i is observable only when $\delta_i = 1$.

We define the cure indicator, s_i, such that $s_i = 1$ if subject i is a progressor, and $s_i = 0$ if the subject is a non-progressor and diagnosis time $U_i = \infty$. Then, there are three possible diagnosis time censoring/cure scenarios for each subject:

1. $s_i = 1,\ \delta_i = 1 \quad (U_i \leq T_i)$

2. $s_i = 1,\ \delta_i = 0 \quad (T_i < U_i < \infty)$

3. $s_i = 0,\ \delta_i = 0 \quad (U_i = \infty)$

Only subjects with a censored diagnosis event time are considered for being a non-progressor. We use a probit link to model the cure probability and define a latent auxiliary variable s^*:

$$s_i^* = \boldsymbol{\beta}_c^T \boldsymbol{X}_{i,0} + \zeta_i \ , \zeta_i \sim \mathcal{N}(0, 1),$$

where $\boldsymbol{X}_{i,0}$ is a vector of baseline covariates that are observed when the subject initially enters the study at age $t_{i,0}$. Then the cure status, s_i, is modeled as an indicator for whether s_i^* is positive:

$$s_i = \begin{cases} 1 & s_i^* > 0 \\ 0 & \text{otherwise} x \end{cases} .$$

With this cure submodel, the probability of being a disease progressor can be directly computed as $P(s_i = 1 \mid \boldsymbol{X}_{i,0}, \boldsymbol{\beta}_c) = \boldsymbol{\Phi}(\boldsymbol{\beta}_c^T \boldsymbol{X}_{i,0})$, where $\boldsymbol{\Phi}$ is the cumulative distribution function of the standard normal distribution.

4.2.2 Time-to-Event Model

For progressors, we denote the distribution of the diagnosis event times, U_i, as $f(U_i \mid \cdot)$ and assume a linear model:

$$U_i \mid \boldsymbol{X}_{i,0} = \boldsymbol{\beta}_{ui}^T \boldsymbol{X}_{i,0} + \epsilon_{i,u} \ , \ \epsilon_{i,u} \sim \mathcal{N}(0, \sigma_{ui}^2)$$

If a progressor has a censored diagnosis event ($s_i = 1, \ \delta_i = 0$), then the unobserved U_i can be imputed from $f(U_i \mid \cdot)\mathbb{I}_{(U_i \geq T_i)}$. In contrast, if subject i belongs to the non-progressor group we assume a point mass of $U_i = \infty$. Thus, the two cure status possibilities lead to two models for U_i with different dimensions. To resolve this change in dimension that complicated posterior simulations, we use the pseudo prior approach described in Carlin and Chib (1995). This approach involves defining a prior $\pi(U)$ for subjects in the non-progressor group to match model dimensions and simplify MCMC updates. The pseudo prior is chosen by matching first and second moments with the marginal posterior of U under a model without cure, where $s_i = 1$

for all subjects.

4.2.3 Conditional Longitudinal Model

Denote $Y_{i,j}$, $\boldsymbol{X}_{i,j}$, and $\boldsymbol{Z}_{i,j}$ to be the longitudinal biomarker measurements, fixed effects covariates, and random effects covariates, respectively, for subject i at the j-th observation at time $t_{i,j}$ (subject age in years), $i = 1, \ldots, I$, $j = 1, \ldots, J_i$. Our longitudinal model for each subject is a linear mixed effects model conditional on cure status and cluster membership for progressors. For the non-progressor group, we assume:

$$Y_{i,j} \mid s_i = 0, \boldsymbol{X}_{i,j}, \boldsymbol{Z}_{i,j} \ = \boldsymbol{\beta}_1^T \boldsymbol{X}_{i,j} + \boldsymbol{b}_{1i}^T \boldsymbol{Z}_{i,j} + \epsilon_{i,j},$$

where $\epsilon_{i,j} \sim \mathcal{N}(0, \sigma_1^2)$ and $\boldsymbol{b}_{1i} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{\Sigma}_1)$. On the other hand, the longitudinal biomarker data for the progressor group is modeled by a separate linear mixed effects model to avoid identifiability issues in the cure model:

$$Y_{i,j} \mid s_i = 1, \boldsymbol{X}_{i,j}, \boldsymbol{Z}_{i,j}, U_i = \boldsymbol{\beta}_{2i}^T \boldsymbol{X}_{i,j} + \boldsymbol{b}_{2i}^T \boldsymbol{Z}_{i,j} + \gamma_i (t_{i,j} - U_i + \tau_i)^+ + \epsilon_{i,j},$$

where $\epsilon_{i,j} \sim \mathcal{N}(0, \sigma_2^2)$, $\boldsymbol{b}_{2i} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{\Sigma}_2)$, and $(t_{i,j} - U_i + \tau_h)^+ = \max(0, t_{i,j} - U_i + \tau_h)$. Under this model, the longitudinal biomarkers and time to diagnosis event data depend on each other for progressors: the change point for disease progression allows the biomarker trend to change at a latent point τ_h years before the diagnosis event. We impose a uniform prior on τ_h to restrict it to a clinically reasonable range.

4.2.4 Dirichlet Process Mixture Model

In section 2.3, we introduced the Dirichlet Process prior and a nonparametric Dirichlet Process mixture (DPM) model for survival event times. For our current framework, we assume a DPM model for selected components of our model to cluster subjects and account for subject heterogeneity. The model parameters that are shared among clusters include the diagnosis event parameters and linear fixed effects coefficients and change point parameters for progressors: $\boldsymbol{\theta}_i = \{\boldsymbol{\beta}_{ui}, \sigma^2_{ui}, \boldsymbol{\beta}_{2i}, \gamma_i, \tau_i\}$. The DPM model can be summarized as:

$$\begin{aligned} Y_{i,j}, U_i \mid \boldsymbol{\theta}_i &\sim F(\boldsymbol{\theta}_i) \\ \boldsymbol{\theta}_i \mid G &\sim G \\ G &\sim DP(G_0, \alpha), \end{aligned}$$

where G is the mixing distribution for $\boldsymbol{\theta}$, G_0 represents the base distribution, and α refers to the concentration parameter. We can also represent our model using stick breaking representations for U_i and $Y_{i,j}$, both using normal distributions as the mixture distribution:

$$U_i \mid \boldsymbol{X}_{i,0} \sim \sum_{h=1}^{\infty} w_h \mathcal{N}(\boldsymbol{\beta}_{uh}^T \boldsymbol{X}_{i,0}, \sigma^2_{uh})$$

$$Y_{i,j} \mid s_i = 1, \boldsymbol{X}_{i,j}, \boldsymbol{Z}_{i,j}, U_i \sim \sum_{h=1}^{\infty} w_h \mathcal{N}(\boldsymbol{\beta}_{2h}^T \boldsymbol{X}_{i,j} + \boldsymbol{b}_{2i}^T \boldsymbol{Z}_{i,j} + \gamma_h (t_{i,j} - U_i + \tau_h)^+, \sigma_2^2)$$

$$w_h = v_h \prod_{k<h} (1 - v_k) \ , \ v_h \sim \text{Beta}(1, \alpha),$$

where h represents cluster membership. We implement a finite DPM model by setting an upper bound for the total number of clusters to a finite number H. This DPM

model allows us to account for heterogeneity among subjects by clustering them using their diagnosis times and longitudinal measurements.

4.2.5 Joint Likelihood of the Model

In summary, we propose a joint model consisting of a cure model for individual disease progressor status, time to event model for disease diagnosis time, a conditional longitudinal model that assumes different biomarker evolution patterns based on subject-specific latent cure status, and a DPM model that clusters subjects based on their longitudinal and diagnosis event processes. In our context, subjects are not enrolled if they have been diagnosed with the disease or have died before the initial recruitment time. Therefore we must modify the likelihood to account for bias $\min(U_i, D_i) \geq t_{i,0}$ and the full complete data likelihood is

$$
\begin{aligned}
L = \prod_{i=1}^{I} & p(s_i = 0 \mid \boldsymbol{\beta}_c, \boldsymbol{X}_{i,0}) \Bigg(\prod_{j=1}^{J_i} p(Y_{i,j} \mid s_i = 0, \boldsymbol{\beta}_1, \boldsymbol{b}_{1i}, \sigma_1^2, \boldsymbol{X}_{i,j}, \boldsymbol{Z}_{i,j}) \Bigg)^{1-s_i} \times \\
& \Bigg(\Bigg(f(U_i \mid \boldsymbol{X}_{i,0}, \boldsymbol{\beta}_{uh}, \sigma_{uh}^2) \prod_{j=1}^{J_i} p(Y_{i,j} \mid s_i = 1, \boldsymbol{\beta}_{2h}, \boldsymbol{b}_{2i}, \sigma_2^2, \gamma_h, \tau_h, \boldsymbol{X}_{i,j}, \boldsymbol{Z}_{i,j}, U_i) \Bigg)^{\delta_i} \times \\
& \Bigg(f(u \mid \boldsymbol{X}_{i,0}, \boldsymbol{\beta}_{uh}, \sigma_{uh}^2) \mathbb{I}_{(u \geq T_i)} \prod_{j=1}^{J_i} p(Y_{i,j} \mid s_i = 1, \boldsymbol{\beta}_{2h}, \boldsymbol{b}_{2i}, \sigma_2^2, \gamma_h, \tau_h, \boldsymbol{X}_{i,j}, \boldsymbol{Z}_{i,j}, U_i = u) \Bigg)^{1-\delta_i} \times \\
& p(s_i = 1 \mid \boldsymbol{\beta}_c, \boldsymbol{X}_{i,0}) \Bigg)^{s_i} \times \frac{1}{p(\min(U_i, D_i) \geq t_{i,0} \mid \cdot)},
\end{aligned}
$$

where

$$p(\min(U_i, D_i) \geq t_{i,0} \mid \cdot) = p(U_i \geq t_{i,0} \mid \cdot) \times p(D_i \geq t_{i,0})$$
$$= \left[p(s_i = 1 \mid \boldsymbol{\beta}_c, \boldsymbol{X}_{i,0}) p(U_i \geq t_{i,0} \mid s_i = 1, \boldsymbol{\beta}_{uh}, \sigma^2_{uh}, \boldsymbol{X}_{i,0}) + p(s_i = 0 \mid \boldsymbol{\beta}_c, \boldsymbol{X}_{i,0})\right] \times$$
$$p(D_i \geq t_{i,0}).$$

To compute this bias adjustment in the likelihood, we can model the death time using a normal model: $D_i \sim \mathcal{N}(\mu_d, \sigma^2_d)$. However, the modeling of death times is not necessary for posterior inference because we can treat $p(D_i \geq t_{i,0})$ as a constant. We complete the model by imposing the following priors: $\boldsymbol{\beta}_c \sim \mathcal{N}(\mathbf{m}_c, \boldsymbol{V}_c)$, $\sigma^2_{uh} \sim$InvGamma(a_0, b_0), $\boldsymbol{\beta}_{uh} \sim \mathcal{N}(\mathbf{m}_u, \boldsymbol{V}_u)$, $\boldsymbol{\beta}_1 \sim \mathcal{N}(\mathbf{m}_1, \boldsymbol{V}_1)$, $\boldsymbol{\Sigma}_1 \sim$ InvWishart$(\boldsymbol{S}_1^{-1}, \eta_1)$, $\sigma^2_1 \sim$ InvGamma(a_1, b_1), $\boldsymbol{\beta}_{2h} \sim \mathcal{N}(\mathbf{m}_2, \boldsymbol{V}_2)$, $\gamma_h \sim \mathcal{N}(m_\gamma, \sigma^2_\gamma)$, $\tau_h \sim$ Uniform(a_τ, b_τ), $\boldsymbol{\Sigma}_2 \sim$ InvWishart$(\boldsymbol{S}_2^{-1}, \eta_2)$, $\sigma^2_2 \sim$ InvGamma(a_2, b_2). We carry out posterior inference using the Markov chain Monte Carlo (MCMC) sampler. The details are included in Section 4.6.

4.3 Simulation Study

To evaluate our model's ability to identify clusters of similar subjects and individual cure status while being able to accurately model the underlying event time and longitudinal processes, we develop a simulated dataset and fit our model.

4.3.1 Simulation Setup

We simulated a dataset mimicking the AD dataset composed of longitudinal measurements, diagnosis events, and survival events for 552 subjects. 52 subjects are left

truncated (9.4% truncation rate) because their diagnosis or survival event occurred before the first longitudinal measurement time, leaving us with $I = 500$ subjects.

We considered four covariates in $\boldsymbol{X}_{i,j}$: an intercept, years of education (standardized), sex, and time (age of subject in years). The baseline covariates, $\boldsymbol{X}_{i,0}$, and random effects covariates, $\boldsymbol{Z}_{i,j}$, are composed of an intercept, education, and gender. For each subject, the years of education was generated from $\mathcal{N}(17, 2.4)$ and sex (female=1) was sampled from a Bernoulli(0.5). Initial age was generated from $\mathcal{N}(59, 7^2)$ and time (years) between longitudinal measurements was generated from Exponential(1). Survival time was sampled from a Lognormal$(4.4, 0.07)$ and 438 subjects have an observed survival event.

In the cure incidence model, $\boldsymbol{\beta}_c$ was set to $(0.9, 0.3, -0.3)$, which results in 377 out of 500 subjects identified as progressors. Out of these, 226 subjects have an observed diagnosis event (40% censoring rate). The progressors were generated from two clusters, with equal probability of being in either cluster: 182 subjects are in the first cluster, and 195 are in the second cluster. The diagnosis time model parameters for these clusters were set to $\boldsymbol{\beta}_{u1} = (67, 3, 3)$, $\boldsymbol{\beta}_{u2} = (70, 3, 3)$, and $\sigma^2_{u1} = \sigma^2_{u2} = 5^2$. The median diagnosis time was 70 years, with the earliest diagnosis event occurring at 52.1 years and the latest at 85.4 years.

For the longitudinal model, the parameters were $\boldsymbol{\beta}_1 = (-16, 0.15, 0.8, 1)$, $\boldsymbol{\beta}_{21} = (-16, 0.2, 1, 1.2)$, $\boldsymbol{\beta}_{22} = (-20, 0.35, 1.2, 1.4)$, $\tau_1 = 3$, $\tau_2 = 5$, $\gamma_1 = \gamma_2 = 0.1$, $\boldsymbol{V}_{b_1} = \boldsymbol{V}_{b_2} = \begin{bmatrix} 0.04 & 0 & 0 \\ 0 & 0.01 & 0 \\ 0 & 0 & 0.01 \end{bmatrix}$, and $\sigma_1^2 = 0.2^2, \sigma_2^2 = 0.3^2$. The simulated dataset had a total of 8236 longitudinal measurements.

4.3.2 Results: model fitting and posterior inference

We applied the proposed Bayesian joint model to the simulated dataset. The hyperparameters were set to be $\alpha = 1$, $\mathbf{m}_c = \mathbf{m}_u = \mathbf{m}_1 = \mathbf{m}_2 = \mathbf{0}$, $\boldsymbol{V}_c = \boldsymbol{V}_u = \boldsymbol{V}_1 = \boldsymbol{V}_2 = 100^2\mathbf{I}$, $\boldsymbol{S}_1 = \boldsymbol{S}_2 = 1000^2\mathbf{I}$, $\eta_1 = \eta_2 = 4$, $a_0 = b_0 = a_1 = b_1 = a_2 = b_2 = 0.01$, $m_\gamma = 0$, $\sigma_\gamma^2 = 1^2$, $a_\tau = 0$, and $b_\tau = 25$. The model fitting process under our framework begins with establishing a pseudo prior for the diagnosis event times. We follow the recommendations of (Carlin and Chib, 1995) and (Zhang et al., 2010) for constructing the pseudo prior for U, $\pi(U)$, using preliminary data analysis under a model without cure, where all subjects have $s_i = 1$. Under this simplified model with no dimensional change, we ran 10,000 MCMC iterations with an initial burn-in of 5,000 iterations, and a thinning factor of 10, using a finite DPM with five total clusters. Then, we assume a normal distribution for $\pi(U)$ and match the first two moments of this pseudo prior to the marginal posterior density of U.

After the initial pseudo prior specification, the we fit the data under our proposed model with five total clusters and 10,000 MCMC iterations with a burn-in of 5000 iterations and a thinning factor of 10. Although it's common to run into the label-switching problem with mixture models, our MCMC chain does not yield such issues because the cluster memberships do not change post burn-in. For our simulation study with five total clusters, this point estimate gives three empty clusters, two large clusters with 179 and 198 subjects that correspond to the two true clusters in the simulated data, and all 123 subjects correctly identified as non-progressors.

The convergence for the MCMC chain was assessed using R package *coda* and displayed no issues of non-convergence. The trace plots (Figure 4.1) and summary

table (Table 4.1) of the change point parameters, τ_h and γ_h, and the cure status regression parameter, $\boldsymbol{\beta}_c$, show good mixing of the chain and the estimates are close to the truth.

	Cluster 1	Cluster 2
True τ_h	3	5
τ_h Post. Mean (95% CI)	3.044(2.634,3.388)	5.081(4.682,5.533)
True γ_h	0.1	0.1
γ_h Post. Mean (95% CI)	0.103(0.1,0.106)	0.099(0.096,0.103)

Table 4.1: Change point parameter estimation results

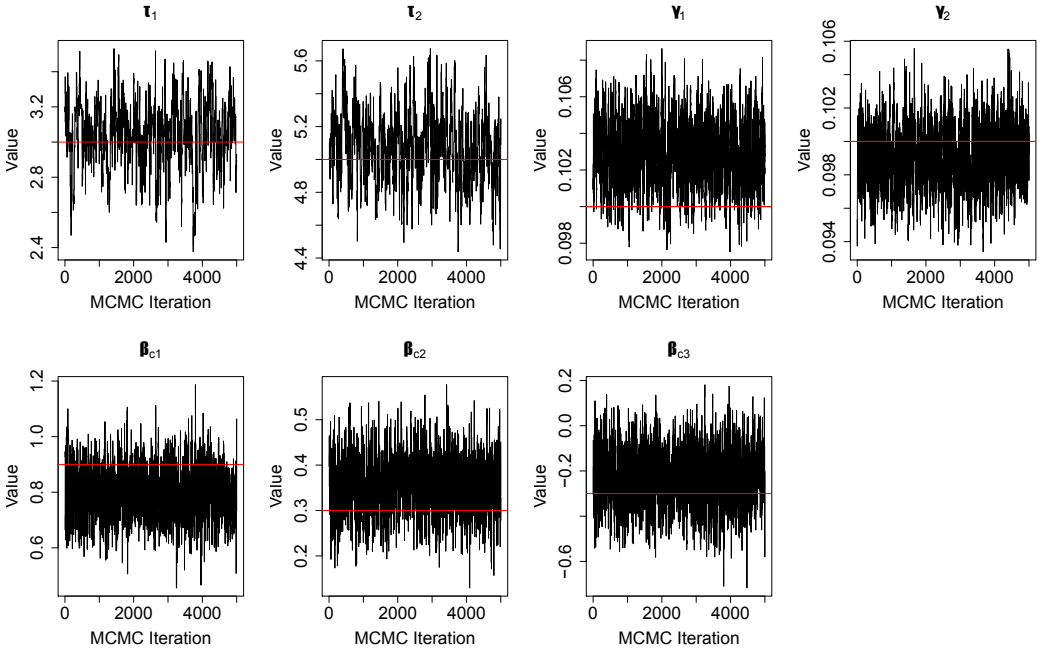

Figure 4.1: Post burn-in trace plots for change point and cure status parameters. The red line represents the true value. The last index in $\boldsymbol{\beta}_c$ refers to the dimension number.

To evaluate the longitudinal model performance, we computed the predictive mean and 95% credible intervals for selected subjects using the MCMC posterior samples. The longitudinal measurements and true longitudinal mean across time is plotted in

Figure 4.2 along with the predictive mean and 95% credible intervals for three randomly selected subjects with different cluster membership and cure status, demonstrating that our model is able to reliably recover the true longitudinal process with change points.

Our model fitting results also allows us to carry out predictive inference for the cure status and diagnosis times of current or future subjects. The details of predictive posterior inference is provided in Section 4.6.1. We randomly select a progressor from our dataset that is female with 20.5 years of education and had a censored diagnosis event. We sample the predictive probability of being a progressor and plot the density in Figure 4.3(a). The predictive density covers the true value probability (0.916) and gives a predictive mean probability of 0.912. The density of the predictive diagnosis event time conditional on the subject being a progressor is provided in Figure 4.3(b) and covers the true unobserved diagnosis time of 78.9 years. Thus, we can construct the predictive density of the diagnosis event time for this subject to be a mixture distribution with probability 0.088 assigned to the point mass of $U_i = \infty$ and with probability 0.912 assigned to the density in Figure 4.3(b).

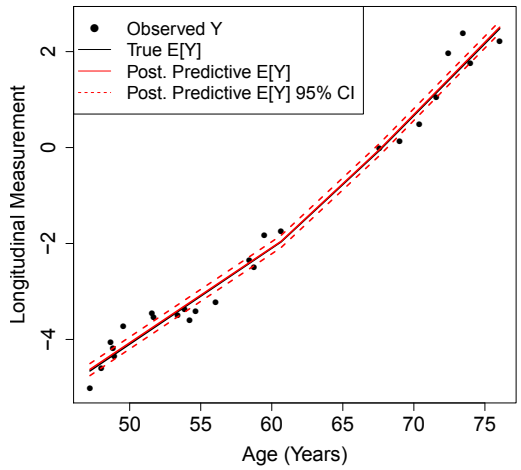

(a) Subject 1: Cluster 1, $s_i = 1$

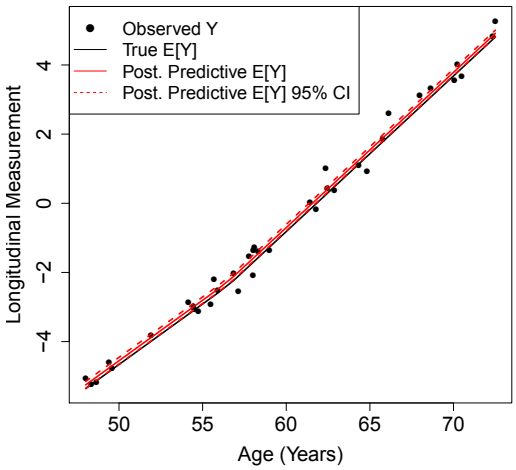

(b) Subject 2: Cluster 2, $s_i = 1$

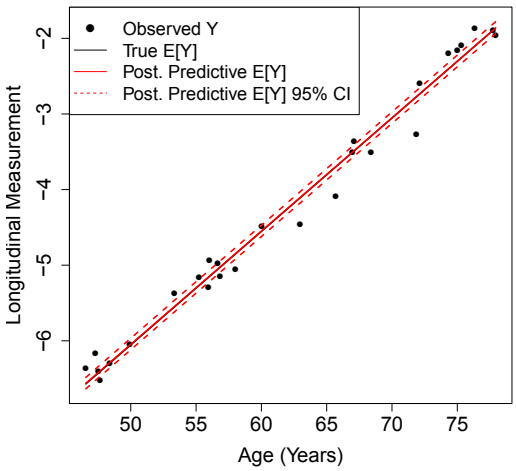

(c) Subject 3: $s_i = 0$

Figure 4.2: These panels plot the longitudinal observations (dots) and true longitudinal mean (black line) across time and the predictive mean (red line) and 95% credible intervals (dashed red lines) for three randomly selected subjects, with different cluster membership and cure status.

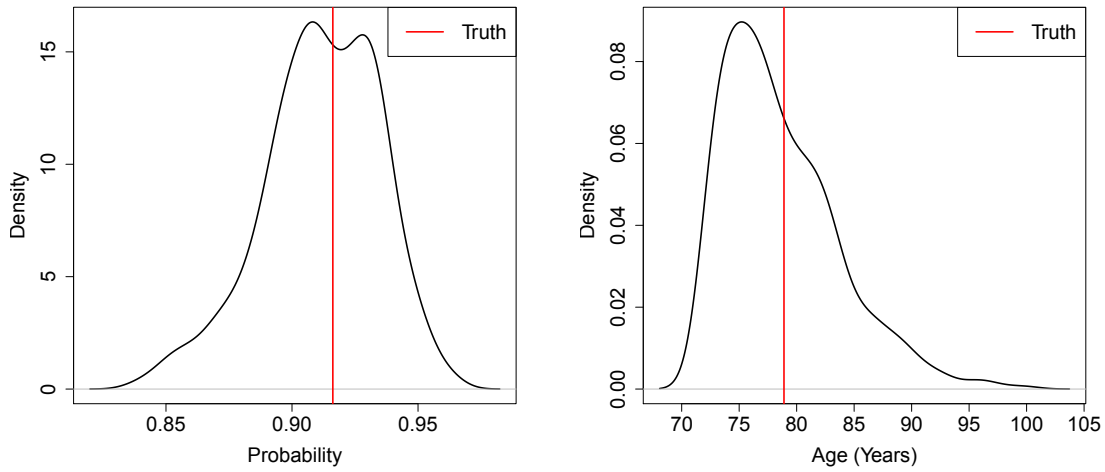

(a) Progressor Probability

(b) Diagnosis event time

Figure 4.3: These panels plot the posterior predictive densities for the diagnosis time and probability of being a progressor for a randomly selected subject. The red lines represent the true unobserved diagnosis time and probability of being a progressor.

4.4 Real Data Analysis

In this section, we analyze the data from the BIOCARD study for for $I = 313$ subjects that longitudinally tracked biomarkers associated with the development of cognitive impairment, particularly Alzheimer's disease. Our analysis uses the biomarker measurements for plasma levels of phosphorylated tau181 (ptau181) in the cerebrospinal fluid (CSF) of subjects. All subjects in this study started with a normal cognitive status and several baseline risk factors were initially collected when subjects joined the study: years of education (Educ), sex (female=1), and the presence of the $\epsilon 4$ allele in the Apolipoprotein E gene (APOE), defined as 1 indicating presence of the allele, and 0 indicating no presence. Table 4.2 summarizes subjects' baseline characteristics.

Years of Education (years)	
Mean ± SD	17 ± 2.4
Median	18
Sex, n(%)	
Male	131 (41.9%)
Female	182 (58.1%)
APOE, n(%)	
Yes	108 (34.5%)
No	205 (65.5%)

Table 4.2: Subject characteristics at baseline.

The subjects' cognitive abilities are repeatedly diagnosed and the diagnosis event of Alzheimer's disease is defined in our context as an onset of mild cognitive impairment (MCI) or dementia. Out of our 313 subjects, 101 have an observed diagnosis time (67.7 % censoring rate) and the largest age at diagnosis was 97.4 years. The Kaplan-Meier plot for the diagnosis times by subject age is provided in Figure 4.4

and suggests that roughly 20% of subjects will not have a diagnosis for AD by age 97.

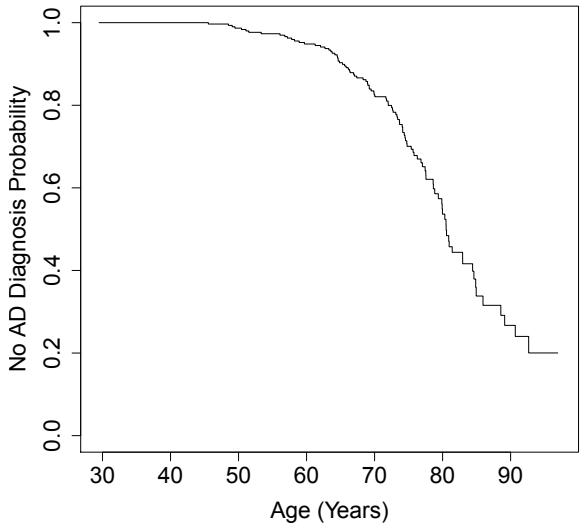

Figure 4.4: Kaplan-Meier plot for the Alzheimer's disease diagnosis event.

4.4.1 Results: model fitting and posterior inference

We applied the proposed Bayesian joint model to the BIOCARD data with $\boldsymbol{X}_{i,j} = (1, \text{Educ}, \text{Sex}, \text{APOE}, t_{i,j})$, $\boldsymbol{Z}_{i,j} = \boldsymbol{X}_{i,0} = (1, \text{Educ}, \text{Sex}, \text{APOE})$. The hyperparameters were set to the same as in the simulation study and we ran the same three-step MCMC inference. In each step, we used a finite DPM with five total clusters and had 10,000 MCMC iterations, an initial burn-in of 5,000 iterations, and a thinning factor of 10. The first step of assuming no cure to construct a pseudo prior for the diagnosis event times, U, yielded a marginal posterior density with a mean of 84.2 years and a standard deviation of 12.3 years.

After constructing a pseudo prior for U, we fit our proposed model to the data using another 10,000 MCMC iterations. Using these posterior samples, we sample from the predictive posterior distributions of diagnosis event times and probability of being a progressor for two randomly selected subjects and plot the densities in Figure 4.5. Subject R1 is male with 16 years of education and positive APOE status while subject R2 is female with 18 years of education and negative APOE status. As a result of their varying baseline characteristics, these two subjects have significant differences in the predictive densities for probability of being a progressor and subject R1 has a predictive mean probability of 0.98 while subject R2 has a probability of 0.54. Subject R1's predictive density for diagnosis time is a mixture distribution with a 0.02 probability for $U_i = \infty$ (non-progressor) and a 0.98 probability assigned to the predictive density in Figure 4.5(b), which has a mean of 80.3 years. In contrast, Subject R2's diagnosis time predictive density is a mixture distribution with a 0.46 probability of $U_i = \infty$ and a 0.54 probability to follow the density in Figure 4.5(d), which has a mean of 82.6 years.

Parameter estimation under a mixture model such as the DPM for diagnosis event times is complicated by the label switching problem, where the unidentifiability of clusters makes it difficult to interpret parameter estimation results. To resolve this issue, we construct point estimates of clustering based on the least-squares distances from the posterior probabilities of subject pairs sharing the same cluster (Vannucci et al., 2009). We then use these point estimates of clusters to fix the cluster membership of each subject, and continue to run another 10,000 MCMC iterations. The convergence for all three MCMC chains, each with 10,000 iterations, was assessed using R package *coda* and showed no issues of non-convergence. Fixing the cluster membership results in 75 subjects (24% of all subjects) being identified as non-progressors.

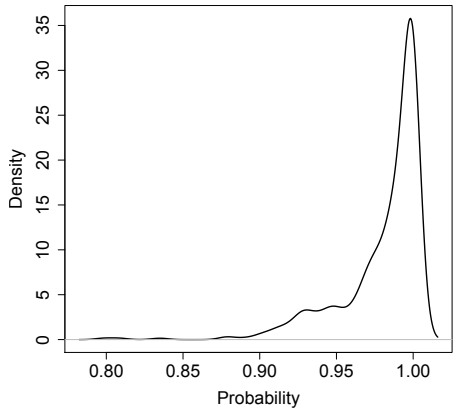

(a) Subject R1- Progressor Probability

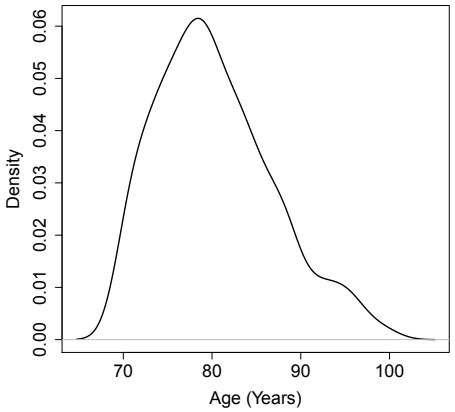

(b) Subject R1- Diagnosis event time

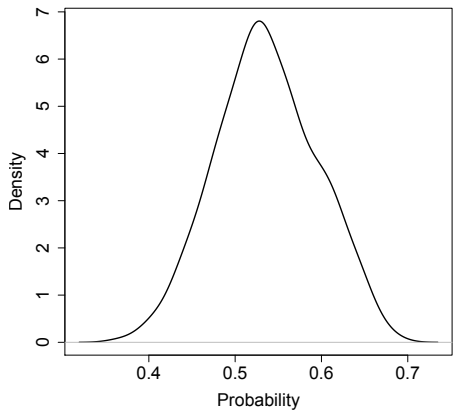

(c) Subject R2- Progressor Probability

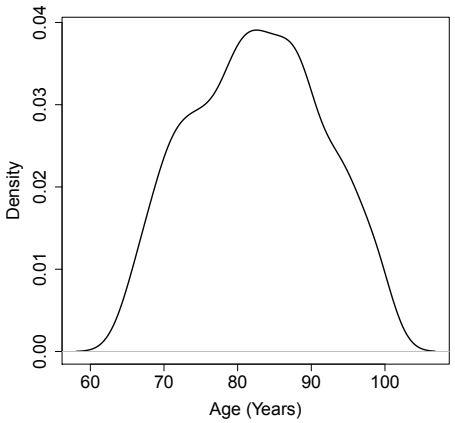

(d) Subject R2- Diagnosis event time

Figure 4.5: These panels plot the posterior predictive densities for the diagnosis time and probability of being a progressor for two randomly selected subjects.

The remaining 238 progressors were clustered into five groups and baseline statistics and change point estimation results are summarized in Table 4.3. Cluster 5 contains most of the progressors and the slope change at the change point, γ_h, is close to zero and its 95% credible interval covers zero. Thus, we believe these 153 subjects have no change point. The other four clusters have varying change points between 3 and 13 years before the diagnosis time. Subjects in clusters 1 and 2 experienced a decrease in the slope at the change point while clusters 3 and 4 had their slopes increase at

the change point. Compared to the five progressor clusters, the non-progressors had more education (mean of 17.35 years), a significantly higher percentage of females (74.67%), and a significantly lower percentage of APOE (2.67%).

	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5
N	39	16	26	4	153
Mean Educ.	17.14	16.94	16.96	17	16.78
% Female	56.41	37.5	57.69	50	52.94
% APOE	51.28	68.75	46.15	50	39.87
τ_h Post. Mean	12.83	10.18	7.66	3.54	17.72
τ_h Post.95% CI	(1.64,20.92)	(6.75,14.94)	(6.26,9.03)	(0.57,7.66)	(11.52,24.34)
γ_h Post. Mean	-1.17	-2.12	3.4	2.66	0.01
γ_h Post.95% CI	(-2.43,-0.65)	(-2.47,-1.82)	(2.97,3.81)	(1.31,4)	(-0.09,0.12)

Table 4.3: Summary of results with fixed clusters

We plot the longitudinal measurements of ptau181 grouped by cluster membership in Figure 4.6, where the zero on the x-axis is the posterior mean of the change point. The slopes of the longitudinal process significantly changes at the change point for clusters 1-4 but remains stable in cluster 5. Figure 4.7 plots the longitudinal measurements of the non-progressors and progressors. The non-progressors have a significantly flatter slope and lower variance in the ptau181 longitudinal trajectories compared to the progressors. These findings agree with studies in the literature that identify high levels of ptau181 as a diagnostic tool for Alzheimer's disease (Thijssen et al., 2020).

Figure 4.8 plots the trace plots for the cure status coefficient, $\boldsymbol{\beta}_c$, which has a posterior mean of (0.83,-0.28,-0.83,2.05). Years of education is negatively associated with being a progressor for AD, which aligns with findings in Sharp and Gatz (2011) that have shown that low education increases the risk of dementia. Furthermore, our analysis finds that the presence of the APOE $\epsilon 4$ allele significantly increases the risk of developing AD, agreeing with previous clinical studies that identify APOE $\epsilon 4$ as

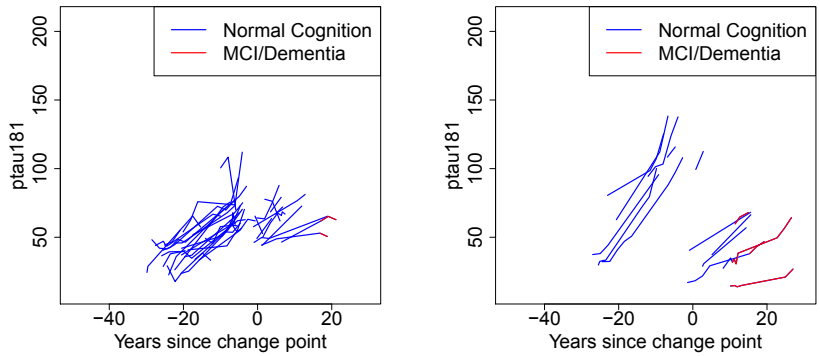

(a) Cluster 1

(b) Cluster 2

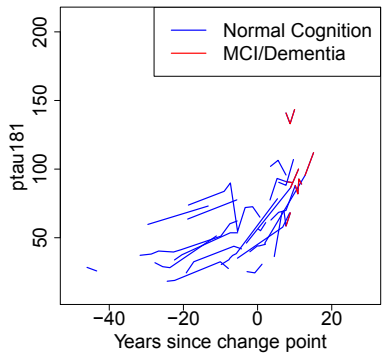

(c) Cluster 3

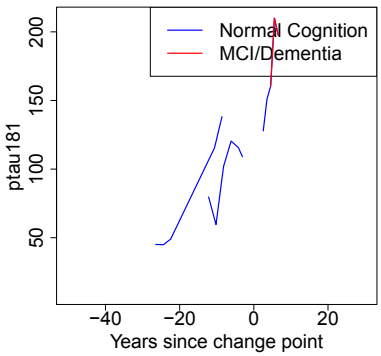

(d) Cluster 4

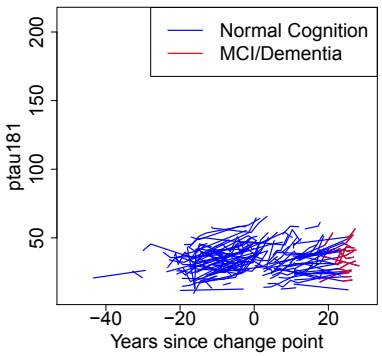

(e) Cluster 5

Figure 4.6: These panels plot the ptau181 biomarker for progressors grouped by cluster membership. Zero on the x-axis is the posterior mean of the change point. Each line represents the measurements for one subject and the red line segments are measurements after the AD diagnosis, when the subject is experiencing MCI/dementia.

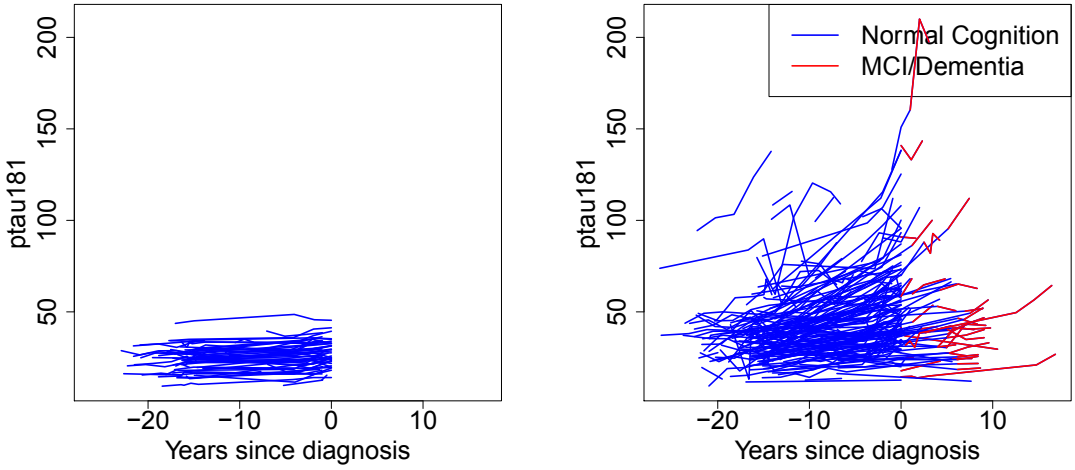

(a) Non-Progressors (b) Progressors

Figure 4.7: These panels plot the ptau181 biomarker for each subject, grouped by cure status. The zero on the x-axis is T_i.

the most prevalent genetic risk factor of AD (Safieh et al., 2019)

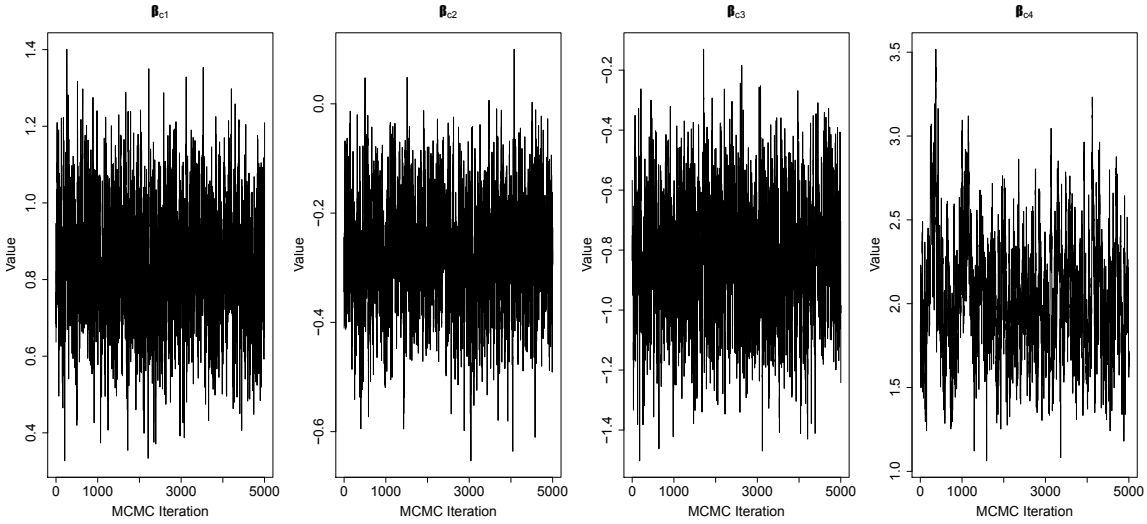

Figure 4.8: Post burn-in trace plots for cure status parameters, where the last index in $\boldsymbol{\beta}_c$ refers to the dimension number.

4.5 Conclusion

In this chapter, we developed a semiparametric Bayesian framework that jointly models longitudinal biomarker measurements, latent cure status, and diagnosis event times. The model simultaneously identifies a change point in the longitudinal evolution of a disease that precedes diagnosis time, accounts for the possibility of individual disease non-progression, and clusters subjects using longitudinal and diagnosis time data. Through simulation studies, we demonstrated that our model is able to reliably cluster subjects with similar longitudinal and diagnosis event time patterns, compute the personalized probability of being a disease progressor, and recover the true underlying longitudinal and diagnosis event time processes. We applied our proposed model to the BIOCARD dataset, yielding clinically relevant results and detecting longitudinal change points in the ptau181 biomarker for select subjects occurring between 3 and 13 years before an official diagnosis for Alzheimer's disease.

There are a number of possible extensions for our methods. Firstly, we consider one longitudinal biomarker measurement in the longitudinal process, which can be expanded to a multivariate analysis of multiple longitudinal biomarkers. Secondly, we can allow for individual heterogeneity in the change point years before diagnosis event, τ_h. In our current setup, this value is fixed and shared among all subjects in the same cluster.

4.6 MCMC Details

For our finite DPM model, we work with mixtures of a large but finite number (H) of clusters. We use $r_i = h$ to indicate that i-th subject is assigned to h-th cluster. Let the bias adjustment in the likelihood be denoted as $B_{i,h}(\boldsymbol{\beta}_c, \boldsymbol{\beta}_{uh}, \sigma^2_{uh})$ for

progressors and $B_i(\boldsymbol{\beta}_c)$ for non-progressors. Then for progressors, $B_{i,h}(\boldsymbol{\beta}_c, \boldsymbol{\beta}_{uh}, \sigma^2_{uh}) = \Phi(\boldsymbol{\beta}_c^T \boldsymbol{X}_{i,0}) p(\widetilde{U}_i \geq t_{i1} \mid s_i = 1, r_i = h, \boldsymbol{\beta}_{uh}, \sigma^2_{uh}, \boldsymbol{X}_{i,0}) + (1 - \Phi(\boldsymbol{\beta}_c^T \boldsymbol{X}_{i,0}))$, where $\widetilde{U}_i$ represents the diagnosis time: either the observed diagnosis time, U_i ($\delta_i = 1$), or the imputed diagnosis time for subjects with a censored diagnosis time ($\delta_i = 0$). For non-progressors, $B_i(\boldsymbol{\beta}_c) = \Phi(\boldsymbol{\beta}_c^T \boldsymbol{X}_{i,0}) \sum_{h=1}^{H} w_h p(\widetilde{U}_i \geq t_{i1} \mid s_i = 1, r_i = h, \boldsymbol{\beta}_{uh}, \sigma^2_{uh}, \boldsymbol{X}_{i,0}) + (1 - \Phi(\boldsymbol{\beta}_c^T \boldsymbol{X}_{i,0}))$. Furthermore, let $\boldsymbol{Y}_i = (Y_{i,1}, \ldots, Y_{i,J_i})^T$, $\boldsymbol{t}_i = (t_{i,1}, \ldots, t_{i,J_i})$, $\boldsymbol{X}_i = (\boldsymbol{X}_{i,1}, \ldots, \boldsymbol{X}_{i,J_i})$, and $\boldsymbol{Z}_i = (\boldsymbol{Z}_{i,1}, \ldots, \boldsymbol{Z}_{i,J_i})$.

1. Update r_i for $i \in \mathcal{J}_0 := \{i : s_i = 0\}$:

   If $s_i = 1$ and $\delta_i = 0$:

   $$p(r_i = h \mid \cdots) = \frac{w_h \Phi(\widetilde{U}_i \mid \boldsymbol{\beta}_{uh}^T \boldsymbol{X}_{i,0}, \sigma^2_{uh}) p(\boldsymbol{Y}_i \mid r_i = h, \cdots) B_{i,h}(\boldsymbol{\beta}_c, \boldsymbol{\beta}_{uh}, \sigma^2_{uh})}{\sum_{h=1}^{H} w_h \Phi(\widetilde{U}_i \mid \boldsymbol{\beta}_{uh}^T \boldsymbol{X}_{i,0}, \sigma^2_{uh}) p(\boldsymbol{Y}_i \mid r_i = h, \cdots) B_{i,h}(\boldsymbol{\beta}_c, \boldsymbol{\beta}_{uh}, \sigma^2_{uh})}$$

   where Φ refers to the normal distribution CDF.

   If $s_i = 1$ and $\delta_i = 1$:

   $$p(r_i = h \mid \cdots) = \frac{w_h p(\widetilde{U}_i \mid \boldsymbol{\beta}_{uh}^T \boldsymbol{X}_{i,0}, \sigma^2_{uh}) p(\boldsymbol{Y}_i \mid r_i = h, \cdots) B_{i,h}(\boldsymbol{\beta}_c, \boldsymbol{\beta}_{uh}, \sigma^2_{uh})}{\sum_{h=1}^{H} w_h p(\widetilde{U}_i \mid \boldsymbol{\beta}_{uh}^T \boldsymbol{X}_{i,0}, \sigma^2_{uh}) p(\boldsymbol{Y}_i \mid r_i = h, \cdots) B_{i,h}(\boldsymbol{\beta}_c, \boldsymbol{\beta}_{uh}, \sigma^2_{uh})}.$$

2. Update v_h for $h = 1, \cdots, H-1$.

   $$f(v_h \mid \cdots) \propto \text{Beta}(1 + \sum_{i=1}^{I} \mathbb{I}_{(r_i = h)}, \alpha + \sum_{k=h+1}^{H} \sum_{i=1}^{I} \mathbb{I}_{(r_i = k)}) \,.$$

3. Update s_i for $i \in \mathcal{J}_0 := \{i : \delta_i = 0\}$

   $$P(s_i = 1 \mid \cdots) = \frac{P_1}{P_1 + P_0},$$

   where $P_1 = \Phi(\boldsymbol{\beta}_c^T \boldsymbol{X}_{i,0}) \sum_{h=1}^{H} w_h p(Y_i \mid \widetilde{U}_i, s_i = 1, r_i = h, \cdots) p(\widetilde{U}_i \mid \boldsymbol{\beta}_{uh}, \sigma^2_{uh}, \boldsymbol{X}_{i,0})$,

$P_0 = (1 - \Phi(\boldsymbol{\beta}_c^T \boldsymbol{X}_{i,0}))p(Y_i \mid \widetilde{U}_i, s_i = 0, \cdots)\pi(\widetilde{U}_i)$, and $\pi(\cdot)$ is pseudo prior for non-progressors.

4. Update s_i^*, (truncated normal),

$$p(\widetilde{s}_i^* \mid r_i = h, \cdots) = \begin{cases} \mathcal{N}(\boldsymbol{\beta}_c^T \boldsymbol{X}_{i,0}, 1)\mathbb{I}_{(0,\infty)} & s_i = 1 \\ \mathcal{N}(\boldsymbol{\beta}_c^T \boldsymbol{X}_{i,0}, 1)\mathbb{I}_{(-\infty,0)} & s_i = 0 \end{cases}.$$

5. Update $\boldsymbol{\beta}_c$

$$\widetilde{\boldsymbol{V}}_c = (\boldsymbol{V}_c^{-1} + \sum_{i=1}^{I} \boldsymbol{X}_{i,0}^T \boldsymbol{X}_{i,0})^{-1}$$
$$\widetilde{\mathbf{m}}_c = \widetilde{\boldsymbol{V}}_c(\boldsymbol{V}_c^{-1}\mathbf{m}_c + \sum_{i=1}^{I} \boldsymbol{X}_{i,0}^T s_i^*)\,.$$

$$f(\boldsymbol{\beta}_c \mid \cdots) \propto \mathcal{N}(\widetilde{\mathbf{m}}_c, \widetilde{\boldsymbol{V}}_c) \prod_{i=1}^{I} p(\min(\widetilde{U}_i, D_i) \geq t_{i1} \mid \boldsymbol{\beta}_c, \cdots)$$
$$\propto \mathcal{N}(\widetilde{\mathbf{m}}_c, \widetilde{\boldsymbol{V}}_c) \prod_{i:s_i=1} B_{i,r_i}(\boldsymbol{\beta}_c, \boldsymbol{\beta}_{ur_i}, \sigma^2_{ur_i}) \prod_{i:s_i=0} B_i(\boldsymbol{\beta}_c)$$

We use the Metropolis Hastings algorithm with a proposal distribution of $\mathcal{N}(\widetilde{\mathbf{m}}_c, \widetilde{\boldsymbol{V}}_c)$. The acceptance probability of the proposed $\boldsymbol{\beta}_c^*$ is:

$$A(\boldsymbol{\beta}_c^*, \boldsymbol{\beta}_c) = \min\left\{1, \frac{\prod_{i:s_i=1} B_{i,r_i}(\boldsymbol{\beta}_c^*, \boldsymbol{\beta}_{ur_i}, \sigma^2_{ur_i}) \prod_{i:s_i=0} B_i(\boldsymbol{\beta}_c^*)}{\prod_{i:s_i=1} B_{i,r_i}(\boldsymbol{\beta}_c, \boldsymbol{\beta}_{ur_i}, \sigma^2_{ur_i}) \prod_{i:s_i=0} B_i(\boldsymbol{\beta}_c)}\right\}$$

6. Update $\widetilde{U}_i$, for $i \in \mathcal{J}_0 := \{i : \delta_i = 0\}$

Given $s_i = 0$, we simulate $\widetilde{U}_i$ from pseudo prior $\pi(\cdot)$. If $s_i = 1$, then we want to

sample from:

$$p(\boldsymbol{Y}_i \mid \widetilde{U}_i, s_i = 1, \cdots) p(\widetilde{U}_i \mid r_i = h, \boldsymbol{\beta}_{uh}, \sigma_{uh}) \mathbb{I}_{(\widetilde{U}_i > T_i)}$$

$$\propto \exp\Big(- \frac{(\widetilde{U}_i - \boldsymbol{\beta}_{uh}^T \boldsymbol{X}_{i,0})^2}{2\sigma_{uh}^2} - \sum_{j=1}^{J_i} \frac{\mu_{i,j} - \gamma_h (t_{i,j} - \widetilde{U}_i + \tau_h)^+\big)^2}{2\sigma_2^2} \Big) \mathbb{I}_{(\widetilde{U}_i > T_i)}$$

, where $\mu_{i,j} = Y_{i,j} - \boldsymbol{\beta}_{2h}^T \boldsymbol{X}_{i,j} - \boldsymbol{b}_{2i}^T Z_i$. If we let K_i denote the number of visits after $T_i - \gamma_h$, then we need to sample from a mixture of $K_i + 1$ truncated normal distributions. If we define $t_{i,0} = 0$ and $t_{i,J_i+1} = \infty$, the k-th truncated normal distribution is in the range $t_{i,J_i-k} + \tau_h < \widetilde{U}_i \leq t_{i,J_i-k+1} + \tau_h$:

$$\mathcal{N}(m_u, v_u) \mathbb{I}_{(\widetilde{U}_i > T_i)} \mathbb{I}_{t_{i,J_i-k} + \tau_h < \widetilde{U}_i \leq t_{i,J_i-k+1} + \tau_h}$$

$$v_u = \frac{\sigma_{uh}^2 \sigma_2^2}{\sigma_2^2 + k \sigma_{uh}^2 \gamma_h^2}$$

$$m_u = \frac{\sigma_2^2 \boldsymbol{\beta}_{uh}^T \boldsymbol{X}_{i,0} - \sigma_{uh}^2 \sum_{j=J_i-k+1}^{J_i} \mu_{i,j} \gamma_h}{\sigma_2^2 + k \sigma_{uh}^2 \gamma_h^2} \mathbb{I}_{t_{i,J_i-k} + \tau_h < \widetilde{U}_i \leq t_{i,J_i-k+1} + \tau_h}$$

The probability of sampling from each of the $K_i + 1$ truncated normal distributions is proportional to the integral of their respective normal distribution over its truncated support.

7. Update $\boldsymbol{\beta}_{uh}$

$$\widetilde{\boldsymbol{V}}_u = \left(\boldsymbol{V}_u^{-1} + \sigma_{uh}^{-2} \sum_{i: r_i = h, s_i = 1} \boldsymbol{X}_{i,0} \boldsymbol{X}_{i,0}^T \right)^{-1}$$

$$\widetilde{\mathbf{m}}_u = \widetilde{\boldsymbol{V}}_u \left(\boldsymbol{V}_u^{-1} \mathbf{m}_u + \sigma_k^{-2} \sum_{i: r_i = h, s_i = 1} \boldsymbol{X}_{i,0} \widetilde{U}_i \right) .$$

$$f(\boldsymbol{\beta}_{uh} \mid \cdots) \propto \mathcal{N}(\widetilde{\mathbf{m}}_u, \widetilde{\boldsymbol{V}}_u) \prod_{i:r_i=h,s_i=1} p(\min(\widetilde{U}_i, D_i) \geq t_{i1} \mid \boldsymbol{\beta}_{uh}, \cdots)$$
$$\propto \mathcal{N}(\widetilde{\mathbf{m}}_u, \widetilde{\boldsymbol{V}}_u) \prod_{i:r_i=h,s_i=1} B_{i,h}(\boldsymbol{\beta}_c, \boldsymbol{\beta}_{uh}, \sigma^2_{uh})$$

We use the Metropolis Hastings algorithm with a proposal distribution of $\mathcal{N}(\widetilde{\mathbf{m}}_u, \widetilde{\boldsymbol{V}}_u)$. The acceptance probability of the proposed $\boldsymbol{\beta}^*_{uh}$ is:

$$A(\boldsymbol{\beta}^*_{uh}, \boldsymbol{\beta}_{uh}) = \min\left\{1, \frac{\prod_{i:r_i=h,s_i=1} B_{i,h}(\boldsymbol{\beta}_c, \boldsymbol{\beta}^*_{uh}, \sigma^2_{uh})}{\prod_{i:r_i=h,s_i=1} B_{i,h}(\boldsymbol{\beta}_c, \boldsymbol{\beta}_{uh}, \sigma^2_{uh})}\right\}$$

8. Update σ^2_{uh}. The full conditional for σ^{-2}_{uh} is

$$\widetilde{a}_0 = a_0 + \frac{\sum_{i=1}^{I} \mathbb{I}_{(r_i=h,s_i=1)}}{2}$$
$$\widetilde{b}_0 = b_0 + \frac{1}{2} \sum_{i:r_i=h,s_i=1} (\widetilde{U}_i - \boldsymbol{\beta}^T_{uh} \boldsymbol{X}_{i,0})^2 .$$

$$f(\sigma^2_{uh} \mid \cdots) \propto \text{Inv-Gamma}(\widetilde{a}_0, \widetilde{b}_0) \prod_{i:r_i=h,s_i=1} p(\min(\widetilde{U}_i, D_i) \geq t_{i1} \mid \sigma^2_{uh}, \cdots)$$
$$\propto \text{Inv-Gamma}(\widetilde{a}_0, \widetilde{b}_0) \prod_{i:r_i=h,s_i=1} B_{i,h}(\boldsymbol{\beta}_c, \boldsymbol{\beta}_{uh}, \sigma^2_{uh})$$

We use the Metropolis Hastings algorithm with a proposal distribution of Inv-Gamma$(\widetilde{a}_0, \widetilde{b}_0)$. The acceptance probability of the proposed ${\sigma^2_{uh}}^*$ is:

$$A({\sigma^2_{uh}}^*, \sigma^2_{uh}) = \min\left\{1, \frac{\prod_{i:r_i=h,s_i=1} B_{i,h}(\boldsymbol{\beta}_c, \boldsymbol{\beta}_{uh}, {\sigma^2_{uh}}^*)}{\prod_{i:r_i=h,s_i=1} B_{i,h}(\boldsymbol{\beta}_c, \boldsymbol{\beta}_{uh}, \sigma^2_{uh})}\right\}$$

9. Update $\boldsymbol{\beta}_{1h}$, $\boldsymbol{\beta}_{2h}$, and γ_h. The full conditional distributions are in standard

form, given as

$$
\begin{aligned}
f(\boldsymbol{\beta}_{1h} \mid \cdots) &\propto \mathcal{N}(\widetilde{\mathbf{m}}_1, \widetilde{\boldsymbol{V}}_1) \\
\widetilde{\boldsymbol{V}}_1 &= \left(\boldsymbol{V}_1^{-1} + \sigma_1^{-2} \sum_{i:r_i=h, s_i=0} \boldsymbol{X}_i \boldsymbol{X}_i^T \right)^{-1} \\
\widetilde{\mathbf{m}}_1 &= \widetilde{\boldsymbol{V}}_1 \left(\boldsymbol{V}_1^{-1} \mathbf{m}_1 + \sigma_1^{-2} \sum_{i:r_i=h, s_i=0} \boldsymbol{X}_i (\boldsymbol{Y}_i - \boldsymbol{b}_{1i}^T \boldsymbol{Z}_i) \right) \\
f(\boldsymbol{\beta}_{2h} \mid \cdots) &\propto \mathcal{N}(\widetilde{\mathbf{m}}_2, \widetilde{\boldsymbol{V}}_2) \\
\widetilde{\boldsymbol{V}}_2 &= \left(\boldsymbol{V}_2^{-1} + \sigma_2^{-2} \sum_{i:r_i=h, s_i=1} \boldsymbol{X}_i \boldsymbol{X}_i^T \right)^{-1} \\
\widetilde{\mathbf{m}}_2 &= \widetilde{\boldsymbol{V}}_2 \left(\boldsymbol{V}_2^{-1} \mathbf{m}_2 + \sigma_2^{-2} \sum_{i:s_i=1} \boldsymbol{X}_i (\boldsymbol{Y}_i - \boldsymbol{b}_{2i}^T \boldsymbol{Z}_i - \gamma_h (\boldsymbol{t}_i - \widetilde{U}_i - \tau_h)^+) \right) \\
f(\gamma_h \mid \cdots) &\propto \mathcal{N}(\widetilde{m}_\gamma, \widetilde{\sigma}_\gamma^2) \\
\widetilde{\sigma}_\gamma^2 &= \left(\sigma_\gamma^{-2} + \sigma_2^{-2} \sum_{i:r_i=h, s_i=1} ((\boldsymbol{t}_i - \widetilde{U}_i + \tau_h)^+)^2 \right)^{-1} \\
\widetilde{m}_\gamma &= \widetilde{\sigma}_\gamma^2 \left(\sigma_\gamma^{-2} m_\gamma + \sigma_2^{-2} \sum_{i:r_i=h, s_i=1} (\boldsymbol{t}_i - \widetilde{U}_i + \tau_h)^+ (\boldsymbol{Y}_i - \boldsymbol{\beta}_{2h}^T \boldsymbol{X}_i - \boldsymbol{b}_{2i}^T \boldsymbol{Z}_i) \right) .
\end{aligned}
$$

10. Update $\boldsymbol{b}_{1i}$ for $i \in \mathcal{J}_1 := \{i : s_i = 0\}$ and $\boldsymbol{b}_{2i}$ for $i \in \mathcal{J}_2 := \{i : s_i = 1\}$,

$$
\begin{aligned}
f(\boldsymbol{b}_{1i} \mid r_i = h, \cdots) &\propto \mathcal{N}(\widetilde{\mathbf{m}}_{b_1}, \widetilde{\boldsymbol{V}}_{b_1}) \\
\widetilde{\boldsymbol{V}}_{b_1} &= \left(\sigma_1^{-2} \sum_{i: r_i = h, s_i = 0} \boldsymbol{Z}_i \boldsymbol{Z}_i^T + \boldsymbol{V}_{b_1}^{-1} \right)^{-1} \\
\widetilde{\mathbf{m}}_{b_1} &= \widetilde{\boldsymbol{V}}_{b_1} \left(\sigma_1^{-2} \sum_{i: r_i = h, s_i = 0} \boldsymbol{Z}_i (\boldsymbol{Y}_i - \boldsymbol{\beta}_{1h}^T \boldsymbol{X}_i) \right) \\
f(\boldsymbol{b}_{2i} \mid r_i = h, \cdots) &\propto \mathcal{N}(\widetilde{\mathbf{m}}_{b_2}, \widetilde{\boldsymbol{V}}_{b_2}) \\
\widetilde{\boldsymbol{V}}_{b_2} &= \left(\sigma_2^{-2} \sum_{i: r_i = h, s_i = 1} \boldsymbol{Z}_i \boldsymbol{Z}_i^T + \boldsymbol{V}_{b_2}^{-1} \right)^{-1} \\
\widetilde{\mathbf{m}}_{b_2} &= \widetilde{\boldsymbol{V}}_{b_2} \left(\sigma_2^{-2} \sum_{i: r_i = h, s_i = 1} \boldsymbol{Z}_i (\boldsymbol{Y}_i - \boldsymbol{\beta}_{2h}^T \boldsymbol{X}_i - \gamma_h (\boldsymbol{t}_i - \widetilde{U}_i + \tau_h)^+) \right) .
\end{aligned}
$$

11. Update $\boldsymbol{V}_{b_1}$, and $\boldsymbol{V}_{b_2}$,

$$
\begin{aligned}
f(\boldsymbol{V}_{b_1} \mid \cdots) &\propto \text{InvWishart} \left(\left(\boldsymbol{S}_1 + \sum_{i: r_i = h, s_i = 0} \boldsymbol{b}_{1i} \boldsymbol{b}_{1i}^T \right)^{-1}, \eta_1 + \sum_i \mathbb{I}_{(r_i = h, s_i = 0)} \right) \\
f(\boldsymbol{V}_{b_2} \mid \cdots) &\propto \text{InvWishart} \left(\left(\boldsymbol{S}_2 + \sum_{i: r_i = h, s_i = 1} \boldsymbol{b}_{2i} \boldsymbol{b}_{2i}^T \right)^{-1}, \eta_2 + \sum_i \mathbb{I}_{(r_i = h, s_i = 1)} \right) .
\end{aligned}
$$

12. Update σ_1, and σ_2,

$$
\begin{aligned}
f(\sigma_1 \mid \cdots) &\propto \text{InvGamma}(a_1 + \frac{\sum_{i: s_i = 0} J_i}{2}, \widetilde{b}_{1h}) \\
\widetilde{b}_{1h} &= b_1 + \frac{1}{2} \sum_{i: r_i = h, s_i = 0} (\boldsymbol{Y}_i - \boldsymbol{\beta}_{1h}^T \boldsymbol{X}_i - \boldsymbol{b}_{1i}^T \boldsymbol{Z}_i)^T (\boldsymbol{Y}_i - \boldsymbol{\beta}_{1h}^T \boldsymbol{X}_i - \boldsymbol{b}_{1i}^T \boldsymbol{Z}_i) \\
f(\sigma_2 \mid \cdots) &\propto \text{InvGamma}(a_2 + \frac{\sum_{i: s_i = 1} J_i}{2}, \widetilde{b}_{2h}) \\
\widetilde{b}_{2h} &= b_2 + \frac{1}{2} \sum_{i: r_i = h, s_i = 1} \mathbf{M}_{i,h}^T \mathbf{M}_{i,h},
\end{aligned}
$$

where $\mathbf{M}_{i,h} = \boldsymbol{Y}_i - \boldsymbol{\beta}_{2h}^T \boldsymbol{X}_i - \boldsymbol{b}_{1i}^T \boldsymbol{Z}_i - \gamma_h(\boldsymbol{t}_i - \widetilde{U}_i + \tau_h)^+$

13. To update change point variable τ_h, there is no closed form formula, and we can use Metropolis-Hastings algorithm to generate τ_h^* from Uniform(a_τ, b_τ), and accept it with probability

$$A(\tau_h^*, \tau_h) = \min\left\{1, \frac{\prod_{i:r_i=h, s_i=1} f(\boldsymbol{Y}_i \mid \tau_h^*, \cdots)}{\prod_{i:r_i=h, s_i=1} f(\boldsymbol{Y}_i \mid \tau_h, \cdots)}\right\}.$$

4.6.1 Predictive Inference

The posterior predictive distribution for the diagnosis time of a progressor i with a censored diagnosis time is given by:

$$p(\widetilde{U}_i \mid \boldsymbol{X}_{i,0}) = \sum_{h=1}^{H} \frac{I_h}{\alpha + I_s} \int \mathcal{N}(\boldsymbol{\beta}_{uh}^T \boldsymbol{X}_{i,0}, \sigma_{uh}^2) \mathbb{I}_{(\widetilde{U}_i > T_i)} p(\boldsymbol{\beta}_{uh}, \sigma_{uh}^2 \mid D_n) d\sigma_{uh}^2 d\boldsymbol{\beta}_{uh}$$

$$+ \frac{\alpha}{\alpha + I_s} \int \mathcal{N}(\boldsymbol{\beta}_{ui}^T \boldsymbol{X}_{i,0}, \sigma_{uh}^2) \mathbb{I}_{(\widetilde{U}_i > T_i)} p(\boldsymbol{\beta}_{ui}) p(\sigma_{ui}^2) d\sigma_{ui}^2 d\boldsymbol{\beta}_{ui}$$

$$= \sum_{h=1}^{H} \frac{I_h}{\alpha + I_s} \int \mathcal{N}(\boldsymbol{\beta}_{uh}^T \boldsymbol{X}_{i,0}, \sigma_{uh}^2) \mathbb{I}_{(\widetilde{U}_i > T_i)} p(\boldsymbol{\beta}_{uh}, \sigma_{uh}^2 \mid D_n) d\sigma_{uh}^2 d\boldsymbol{\beta}_{uh}$$

$$+ \frac{\alpha}{\alpha + I_s} t_{a_0}(\mathbf{m}_u^T \boldsymbol{X}_{i,0}, b_0 + \boldsymbol{X}_{i,0}^T \boldsymbol{V}_u \boldsymbol{X}_{i,0}) \mathbb{I}_{(\widetilde{U}_i > T_i)},$$

where t_{a_0} represents the t-distribution with a_0 degrees of freedom, the entire dataset is denoted $D_n = \{\boldsymbol{Y}_i, \boldsymbol{t}_i, \boldsymbol{X}_i, \boldsymbol{Z}_i, \boldsymbol{X}_{i,0}, \delta_i, U_i\}_{i=1}^n$, I_h denotes the number of subjects in the h-th cluster, and I_s denotes the total number of progressors.

The posterior predictive distribution for the probability of being a progressor, p_i,

is given by:

$$p(p_i \mid \boldsymbol{X}_{i,0}) = \int \Phi(\boldsymbol{\beta}_c^T \boldsymbol{X}_{i,0}) p(\boldsymbol{\beta}_c \mid D_n) d\boldsymbol{\beta}_c$$

Bibliography

Aalen, O., Borgan, O. and Gjessing, H. (2008) *Survival and event history analysis: a process point of view*. Springer Science & Business Media.

Andersson, B. S., Madden, T., Tran, H. T., Hu, W. W., Blume, K. G., Chow, D. S.-L., Champlin, R. E. and Vaughan, W. P. (2000) Acute safety and pharmacokinetics of intravenous busulfan when used with oral busulfan and cyclophosphamide as pretransplantation conditioning therapy: a phase i study. *Biology of Blood and Marrow Transplantation*, **6**, 548–554.

Andersson, B. S., Thall, P. F., Madden, T., Couriel, D., Wang, X., Tran, H. T., Anderlini, P., De Lima, M., Gajewski, J. and Champlin, R. E. (2002) Busulfan systemic exposure relative to regimen-related toxicity and acute graft-versus-host disease: defining a therapeutic window for iv bucy2 in chronic myelogenous leukemia. *Biology of Blood and Marrow Transplantation*, **8**, 477–485.

Andersson, B. S., Thall, P. F., Valdez, B. C., Milton, D. R., Al-Atrash, G., Chen, J., Gulbis, A., Chu, D., Martinez, C., Parmar, S., Popat, U., Nieto, Y., Kebriaei, P., Alousi, A., de Lima, M., Rondon, G., Meng, Q. H., Myers, A., Kawedia, J., Worth, L. L., Fernandez-Vina, M., Madden, T., Shpall, E. J., Jones, R. B. and Champlin, R. E. (2017) Fludarabine with pharmacokinetically guided IV busulfan is superior

to fixed-dose delivery in pretransplant conditioning of AML/MDS patients. *Bone Marrow Transplantation*, **52**, 580–587.

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**.

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**, 1453–1460. URL: `https://doi.org/10.1016/j.kint.2019.01.036`.

Bartelink, I. H., Bredius, R. G., Belitser, S. V., Suttorp, M. M., Bierings, M., Knibbe, C. A., Egeler, M., Lankester, A. C., Egberts, A. C., Zwaveling, J. and Boelens, J. J. (2009) Association between busulfan exposure and outcome in children receiving intravenous busulfan before hematologic stem cell transplantation. *Biology of Blood and Marrow Transplantation*, **15**, 231–241.

Bartelink, I. H., Lalmohamed, A., van Reij, E. M., Dvorak, C. C., Savic, R. M., Zwaveling, J., Bredius, R. G., Egberts, A. C., Bierings, M., Kletzel, M., Shaw, P. J., Nath, C. E., Hempel, G., Ansari, M., Krajinovic, M., Théorêt, Y., Duval, M., Keizer, R. J., Bittencourt, H., Hassan, M., Güngör, T., Wynn, R. F., Veys, P., Cuvelier, G. D., Marktel, S., Chiesa, R., Cowan, M. J., Slatter, M. A., Stricherz, M. K., Jennissen, C., Long-Boyle, J. R. and Boelens, J. J. (2016) Association of busulfan exposure with survival and toxicity after haemopoietic cell transplantation in children and young adults: a multicentre, retrospective cohort analysis. *The Lancet Haematology*, **3**, e526–e536.

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*, **2019**.

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**, 445–448.

Bredeson, C., LeRademacher, J., Kato, K., Dipersio, J. F., Agura, E., Devine, S. M., Appelbaum, F. R., Tomblyn, M. R., Laport, G. G., Zhu, X., McCarthy, P. L., Ho, V. T., Cooke, K. R., Armstrong, E., Smith, A., Rizzo, J. D., Burkart, J. M. and Pasquini, M. C. (2013) Prospective cohort study comparing intravenous busulfan to total body irradiation in hematopoietic cell transplantation. *Blood*, **122**, 3871–3878.

Butts, C. T. and Marcum, C. S. (2017) A Relational Event Approach to Modeling Behavioral Dynamics. In *Group Processes*, 51–92. Cham: Springer.

Carlin, B. P. and Chib, S. (1995) Bayesian model choice via markov chain monte carlo methods. *Journal of the Royal Statistical Society: Series B (Methodological)*, **57**, 473–484.

Chakraborty, B. (2013) *Statistical methods for dynamic treatment regimes*. Springer.

Chi, Y.-Y. and Ibrahim, J. G. (2006) Joint models for multivariate longitudinal and multivariate survival data. *Biometrics*, **62**, 432–445.

Clifton, J. and Laber, E. (2020) Q-learning: Theory and applications. *Annual Review of Statistics and its Application*, **7**, 279–301.

Copelan, E. A., Hamilton, B. K., Avalos, B., Ahn, K. W., Bolwell, B. J., Zhu, X., Aljurf, M., van Besien, K., Bredeson, C., Cahn, J. Y., Costa, L. J., de Lima, M., Gale, R. P., Hale, G. A., Halter, J., Hamadani, M., Inamoto, Y., Kamble, R. T., Litzow, M. R., Loren, A. W., Marks, D. I., Olavarria, E., Roy, V., Sabloff, M., Savani, B. N., Seftel, M., Schouten, H. C., Ustun, C., Waller, E. K., Weisdorf, D. J., Wirk, B., Horowitz, M. M., Arora, M., Szer, J., Cortes, J., Kalaycio, M. E., Maziarz, R. T. and Saber, W. (2013) Better leukemia-free and overall survival in aml in first remission following cyclophosphamide in combination with busulfan compared to tbi. *Blood*, **122**, 3863–3870.

De Iorio, M., Johnson, W. O., Müller, P. and Rosner, G. L. (2009) Bayesian nonparametric nonproportional hazards survival modeling. *Biometrics*, **65**, 762–771.

De Lima, M., Couriel, D., Thall, P. F., Wang, X., Madden, T., Jones, R., Shpall, E. J., Shahjahan, M., Pierre, B., Giralt, S., Korbling, M., Russell, J. A., Champlin, R. E. and Andersson, B. S. (2004) Once-daily intravenous busulfan and fludarabine: clinical and pharmacokinetic results of a myeloablative, reduced-toxicity conditioning regimen for allogeneic stem cell transplantation in aml and mds. *Blood*, **104**, 857–864.

Dix, S. P., Wingard, J. R., Mullins, R. E., Jerkunica, I., Davidson, T. G., Gilmore, C. E., York, R. C., Lin, L. S., Devine, S. M., Geller, R. B., Heffner, L. T., Hillyer, C. D., Holland, H. K., Winton, E. F. and Saral, R. (1996) Association of busulfan area under the curve with veno-occlusive disease following bmt. *Bone marrow transplantation*, **17**, 225–230.

Du, N., Dai, H., Trivedi, R., Upadhyay, U., Gomez-Rodriguez, M. and Song, L. (2016) Recurrent marked temporal point processes: Embedding event history to vector.

In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.

Escobar, M. D. and West, M. (1995) Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association*, **90**, 577–588.

Ferguson, T. S. (1973) A bayesian analysis of some nonparametric problems. *The annals of statistics*, 209–230.

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**, 469–479.

Geddes, M., Kangarloo, S. B., Naveed, F., Chaudhry, A. M., Jeje, O., Savoie, M. L., Bahlis, N. J., Stewart, D. A., Brown, C., Storek, J., Quinlan, D., Andersson, B. and Russell, J. A. (2008) High busulfan exposure is associated with worse outcomes in a daily iv busulfan and fludarabine allogeneic transplant regimen. *Biology of Blood and Marrow Transplantation*, **14**, 220–228.

Gelfand, A. E. and Kottas, A. (2003) Bayesian semiparametric regression for median residual life. *Scandinavian Journal of Statistics*, **30**, 651–665.

Gelman, A., Hwang, J. and Vehtari, A. (2014) Understanding predictive information criteria for bayesian models. *Statistics and computing*, **24**, 997–1016.

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**, 3781–3788.

Giesecke, K., Kakavand, H. and Mousavi, M. (2011) Exact simulation of point processes with stochastic intensities. *Operations Research*, **59**, 1233–1245.

Greensmith, E., Bartlett, P. L. and Baxter, J. (2004) Variance reduction techniques for gradient estimates in reinforcement learning. *Journal of Machine Learning Research*, **5**, 1471–1530.

Guan, Q., Reich, B. J., Laber, E. B. and Bandyopadhyay, D. (2019) Bayesian nonparametric policy search with application to periodontal recall intervals. *Journal of the American Statistical Association*, 1–13.

Hall, C. B., Lipton, R. B., Sliwinski, M. and Stewart, W. F. (2000) A change point model for estimating the onset of cognitive decline in preclinical Alzheimer's disease. *Statistics in Medicine*, **19**, 1555–1566.

Hanson, T. and Johnson, W. O. (2002) Modeling regression error with a mixture of polya trees. *Journal of the American Statistical Association*, **97**, 1020–1033.

Hawkes, A. G. (1971) Spectra of some self-exciting and mutually exciting point processes. *Biometrika*.

— (2018) Hawkes processes and their applications to finance: a review. *Quantitative Finance*, **18**, 193–198. URL: `http://doi.org/10.1080/14697688.2017.1403131`.

Hjelmeseth, 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**, 6.

Ishwaran, H. and Kogalur, U. (2007) Random survival forests for r. *R News*, **7**, 25–31. URL: `https://CRAN.R-project.org/doc/Rnews/`.

Ishwaran, H., Kogalur, U. B., Blackstone, E. H. and Lauer, M. S. (2008) Random survival forests. *Annals of Applied Statistics*, **2**, 841–860.

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**, 520.

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**, 283–288.

Kaelbling, L. P., Littman, M. L. and Moore, A. W. (1996) Reinforcement learning: A survey. *Journal of artificial intelligence research*, **4**, 237—-285.

Kanakry, C. G., O'Donnell, P. V., Furlong, T., de Lima, M. J., Wei, W., Medeot, M., Mielcarek, M., Champlin, R. E., Jones, R. J., Thall, P. F., Andersson, B. S. and Luznik, L. (2014) Multi-institutional study of post-transplantation cyclophosphamide as single-agent graft-versus-host disease prophylaxis after allogeneic bone marrow transplantation using myeloablative busulfan and fludarabine conditioning. *Journal of Clinical Oncology*, **32**, 3497–3505.

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**, 299–311. URL: http://dx.doi.org/10.1038/ki.2009.377.

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**, 237.

Kontoyiannis, D. P., Andersson, B. S., Lewis, R. E. and Raad, I. I. (2001) Progressive disseminated aspergillosis in a bone marrow transplant recipient: response with a high-dose lipid formulation of amphotericin b. *Clinical Infectious Diseases*, **32**, e94–e96.

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**, 1225.

Lamb, K. E. and Lodhi, S. (2011) Long-Term Renal Allograft Survival in the United States : A Critical Reappraisal. *American Journal ofTransplantation*, **11**, 450–462.

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**, 714–724.

Ljungman, P., Hassan, M., Bekassy, A., Ringden, O. and Öberg, G. (1997) High busulfan concentrations are associated with increased transplant-related mortality in allogeneic bone marrow transplant patients. *Bone Marrow Transplantation*, **20**, 909–913.

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.

MacEachern, S. N. (1999) Dependent nonparametric processes. In *ASA proceedings of the Section on Bayesian Statistical Science*, 50–55. American Statistical Association, pp. 50–55, Alexandria, VA.

MacKay, D. (1999) Introduction to Gaussian processes. *Tech. rep.*, Cambridge University, http://wol.ra.phy.cam.ac.uk /mackay/GP/.

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. *Nephrourol Mon.*, **8**, 1–5.

McCulloch, R., Sparapani, R., Gramacy, R., Spanbauer, C. and Pratola, M. (2018) *BART: Bayesian Additive Regression Trees*. URL: `https://CRAN.R-project.org/package=BART`. R package version 1.9.

McCune, J., Gooley, T., Gibbs, J., Sanders, J., Petersdorf, E., Appelbaum, F., Anasetti, C., Risler, L., Sultan, D. and Slattery, J. (2002) Busulfan concentration and graft rejection in pediatric patients undergoing hematopoietic stem cell transplantation. *Bone Marrow Transplantation*, **30**, 167–173.

Mei, H. and Eisner, J. (2017) The neural Hawkes process: A neurally self-modulating multivariate point process. *Advances in Neural Information Processing Systems*, **2017-Decem**, 6755–6765.

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**, 391–398.

Müller, P. and Mitra, R. (2013) Bayesian nonparametric inference – Why and how. *Bayesian Analysis*, **8**, 269–302.

Müller, P. and Rodriguez, A. (2013) Nonparametric Bayesian inference. *IMS-CBMS Lecture Notes. IMS*, **270**.

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**, 331–366.

Naesens, M., Kuypers, D. R. and Sarwal, M. (2009) Calcineurin inhibitor nephrotoxicity. *Clinical Journal of the American Society of Nephrology*, **4**, 481–508.

Nagler, A., Rocha, V., Labopin, M., Unal, A., Ben Othman, T., Campos, A., Volin, L., Poire, X., Aljurf, M., Masszi, T., Socie, G., Sengelov, H., Michallet, M., Passweg, J., Veelken, H., Yakoub-Agha, I., Shimoni, A. and Mohty, M. (2013) Allogeneic hematopoietic stem-cell transplantation for acute myeloid leukemia in remission: comparison of intravenous busulfan plus cyclophosphamide (Cy) versus total-body irradiation plus Cy as conditioning regimen – a report from the acute leukemia working party of the European group for blood and marrow transplantation. *Journal of Clinical Oncology*, **31**, 3549–3556.

Plummer, M., Best, N., Cowles, K. and Vines, K. (2006) Coda: Convergence diag-

nosis and output analysis for mcmc. *R News*, **6**, 7–11. URL: `https://journal.r-project.org/archive/`.

Randhawa, P. S., Starzl, T. E. and Demetris, A. J. (1997) Tacrolimus (FK506)-Associated Renal Pathology. *Advances in Anatomic Pathology*, **4**, 265.

Rasmussen, C. and Williams, C. (2006) *Gaussian Processes for Machine Learning*. ISBN 0-262-18253-X. MIT Press.

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**, 1385–1397.

Robbins, H. and Monro, S. (1951) A Stochastic Approximation Method. *The Annals of Mathematical Statistics*, **22**, 400–407.

Ruder, S. (2016) An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*, 1–14. URL: `http://arxiv.org/abs/1609.04747`.

Russell, J. A., Kangarloo, S. B., Williamson, T., Chaudhry, M. A., Savoie, M. L., Turner, A. R., Larratt, L., Storek, J., Bahlis, N. J., Shafey, M., Brown, C. B., Yang, M., Geddes, M., Zacarias, N., Yue, P., Duggan, P., Stewart, D. A. and Daly, A. (2013) Establishing a target exposure for once-daily intravenous busulfan given with fludarabine and thymoglobulin before allogeneic transplantation. *Biology of Blood and Marrow Transplantation*, **19**, 1381–1386.

Safieh, M., Korczyn, A. D. and Michaelson, D. M. (2019) ApoE4: an emerging therapeutic target for Alzheimer's disease. *BMC Medicine*, **17**, 1–17.

Santos, G. W., Tutschka, P. J., Brookmeyer, R., Saral, R., Beschorner, W. E., Bias, W. B., Braine, H. G., Burns, W. H., Elfenbein, G. J. and Kaizer, H. (1983) Marrow transplantation for acute nonlymphocytic leukemia after treatment with busulfan and cyclophosphamide. *New England Journal of Medicine*, **309**, 1347–1353.

Sethuraman, J. (1994) A constructive definition of Dirichlet priors. *Statistica Sinica*, **4**, 639–650.

Sharp, E. S. and Gatz, M. (2011) Relationship between education and dementia: An updated systematic review. *Alzheimer Disease and Associated Disorders*, **25**, 289–304.

Shibue, R. and Komaki, F. (2020) Deconvolution of calcium imaging data using marked point processes. *PLoS computational biology*, **16**, e1007650. URL: `http://dx.doi.org/10.1371/journal.pcbi.1007650`.

Slattery, J. T., Clift, R. A., Buckner, C. D., Radich, J., Storer, B., Bensinger, W. I., Soll, E., Anasetti, C., Bowden, R., Bryant, E., Chauncey, T., Deeg, H. J., Doney, K. C., Flowers, M., Gooley, T., Hansen, J. A., Martin, P. J., McDonald, G. B., Nash, R., Petersdorf, E. W., Sanders, J. E., Schoch, G., Stewart, P., Storb, R., Sullivan, K. M., Thomas, E. D., Witherspoon, R. P. and Appelbaum, F. R. (1997) Marrow transplantation for chronic myeloid leukemia: the influence of plasma busulfan levels on the outcome of transplantation. *Blood*, **89**, 3055–3060.

Sorror, M. L., Maris, M. B., Storb, R., Baron, F., Sandmaier, B. M., Maloney, D. G. and Storer, B. (2005) Hematopoietic cell transplantation (hct)-specific comorbidity index: a new tool for risk assessment before allogeneic hct. *Blood*, **106**, 2912–2919.

Sorror, M. L., Storb, R. F., Sandmaier, B. M., Maziarz, R. T., Pulsipher, M. A., Maris,

M. B., Bhatia, S., Ostronoff, F., Deeg, H. J., Syrjala, K. L., Estey, E., Maloney, D. G., Appelbaum, F. R., Martin, P. J. and Storer, B. E. (2014) Comorbidity-age index: a clinical measure of biologic age before allogeneic hematopoietic cell transplantation. *Journal of Clinical Oncology*, **32**, 3249–3256.

Sparapani, R. A., Logan, B. R., McCulloch, R. E. and Laud, P. W. (2016) Nonparametric survival analysis using bayesian additive regression trees (bart). *Statistics in Medicine*, **35**, 2741–2753.

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**, 1905–1909.

Sutton, R. S. and Barto, A. G. (2018) *Reinforcement learning: An introduction*.

Sy, J. P. and Taylor, J. M. (2000) Estimation in a Cox proportional hazards cure model. *Biometrics*, **56**, 227–236.

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**, 3988–3993.

Tang, X., Miller, M. I. and Younes, L. (2017) Biomarker change-point estimation with right censoring in longitudinal studies. *The annals of applied statistics*, **11**, 1738.

Thijssen, E. H., La Joie, R., Wolf, A., Strom, A., Wang, P., Iaccarino, L., Bourakova, V., Cobigo, Y., Heuer, H., Spina, S. et al. (2020) Diagnostic value of plasma phosphorylated tau181 in alzheimer's disease and frontotemporal lobar degeneration. *Nature medicine*, **26**, 387–397.

Tutschka, P. J., Copelan, E. A. and Klein, J. P. (1987) Bone marrow transplantation for leukemia following a new busulfan and cyclophosphamide regimen. *Blood*, **70**, 1382–1388.

van der Vaart, A. and van Zanten, H. (2011) Information rates of nonparametric Gaussian process methods. *Journal of Machine Learning Research*, **12**, 2095–2119.

Vannucci, M., Dahl, D. B. and Vannucci, M. (2009) Model-Based Clustering for Expression Data via a Dirichlet Process Mixture Model. *Bayesian Inference for Gene Expression and Proteomics*, 201–218.

Wachowiak, J., Sykora, K., Cornish, J., Chybicka, A., Kowalczyk, J., Gorczyńska, E., Choma, M., Grund, G. and Peters, C. (2011) Treosulfan-based preparative regimens for allo-hsct in childhood hematological malignancies: a retrospective study on behalf of the ebmt pediatric diseases working party. *Bone marrow transplantation*, **46**, 1510–1518.

Williams, R. J. (1992) Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, **8**, 229–256.

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**, 921–950.

Yu, M., Law, N. J., Taylor, J. M. and Sandler, H. M. (2004) Joint longitudinal-survival-cure models and their application to prostate cancer. *Statistica Sinica*, **14**, 835–862.

Zhang, S., Müller, P. and Do, K. A. (2010) A Bayesian semiparametric survival model with longitudinal markers. *Biometrics*, **66**, 435–443.

Zhou, H. and Hanson, T. (2018) A unified framework for fitting bayesian semiparametric models to arbitrarily censored survival data, including spatially-referenced data. *Journal of the American Statistical Association*, **113**, 571–581.

Zhou, H., Hanson, T. and Zhang, J. (2018) spBayesSurv: Fitting bayesian spatial survival models using R. *Journal of Statistical Software*, **accept minor**. URL: `http://arxiv.org/abs/1705.04584`.

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**.

Biographical sketch

William Hua was born in Toronto, Canada on June 22, 1998. He received a B.S. in Mathematics from California State University, Los Angeles in 2015, and then enrolled in the Applied Mathematics and Statistics Ph.D. program at Johns Hopkins University in the same year. At Johns Hopkins University, he is a recipient of the GAANN Fellowship, Rufus P. Isaacs Graduate Fellowship, and Professor Joel Dean Award for Excellence in Teaching.