ABSTRACT

RHODES, GRACE MARIE. Statistical Learning Methods for Dynamic Prediction of Residual Life and Estimation of Optimal Treatment Regimes with Electronic Medical Record Data. (Under the direction of Dr. Marie Davidian and Dr. Wenbin Lu).

Clinicians and patients must make treatment decisions at a series of key decision points throughout disease progression. When treating life-threatening diseases, prediction and optimization of remaining life, formally referred to as "residual life," can provide critical information at these decision points. Intuitively, it is desirable to leverage all available patient information at each decision point. Patient information is frequently stored in electronic medical records (EMRs), which often contain longitudinal biomarkers, repeatedly measured variables that reflect disease progression. Dynamic methods leverage accruing longitudinal biomarker measurements to provide individualized, updated results as new observations become available. In this thesis, we introduce statistical learning techniques to dynamically predict residual life and estimate optimal treatment regimes using EMR data. In Chapter 2, we introduce two methods for dynamically predicting mean residual life, the LSTM-GLM and the LSTM-NN, which use long short-term memory networks to construct encoded representations of the longitudinal biomarker trajectories, referred to as "context vectors." In Chapter 3, we introduce ReLiVE, a method that leverages context vectors to estimate the value of a fixed treatment regime, where the value is defined in terms of restricted residual life. Building on ReLiVE, we present ReLiVE-Q, a Q-learning method for estimating an optimal treatment regime that maximizes the expected value of restricted residual life. We demonstrate the utility of the proposed LSTM and ReLiVE methods through simulation studies, and we apply the methods to predict and optimize the residual life of patients diagnosed with sepsis, a complex medical condition that involves severe infections with life-threatening organ dysfunction. Specifically, we study the EMR data of septic patients in the intensive care unit (ICU) from the Multiparameter Intelligent Monitoring Intensive Care database (MIMIC-III).

Statistical Learning Methods for Dynamic Prediction of Residual Life and Estimation of Optimal Treatment Regimes with Electronic Medical Record Data

by
Grace Marie Rhodes

A dissertation submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the Degree of
Doctor of Philosophy

Statistics

Raleigh, North Carolina
2023

APPROVED BY:

Daowen Zhang	Kevin Flores
Marie Davidian Co-chair of Advisory Committee	Wenbin Lu Co-chair of Advisory Committee

DEDICATION

To Sherri and Tim Rhodes

BIOGRAPHY

Grace Rhodes grew up in Wadsworth, Ohio surrounded by a large, loud, and loving family. After graduating from Wadsworth High School, she moved to Columbus to attend The Ohio State University. In 2018, she graduated summa cum laude, earning a Bachelor of Science in Data Analytics with a minor in Mathematics. Inspired by an undergraduate experience shadowing at a clinical research organization, Grace developed a passion for applying statistics to improve patients' health and wellbeing. Consequently, she moved to Raleigh, North Carolina to pursue a PhD in Statistics at North Carolina State University. In 2021, Grace earned her Masters of Statistics from North Carolina State, with a specialization in Biostatistics. While completing her doctoral studies, Grace worked at the Duke Clinical Research Institute (DCRI) on a National Institutes of Health training grant. Grace's experience at DCRI deepened her interest in and passion for clinical research. Accordingly, Grace will begin working as a Research Advisor in clinical statistics at Eli Lilly and Company after completing her doctorate in 2023.

ACKNOWLEDGEMENTS

First and foremost, I am immensely grateful for my wonderful advisors, Dr. Marie Davidian and Dr. Wenbin Lu. Their consistent support, guidance, and expertise were instrumental to the success of my graduate school career. It has been an incredible honor to learn from each of them. Additionally, I am grateful for the time and insightful feedback shared by my advisory committee members, Dr. Daowen Zhang and Dr. Kevin Flores.

I would like to extend a special thanks to Dr. Marie Davidian and Dr. Ana-Maria Staicu for affording me the opportunity to participate in the NHLBI Integrated Biostatistical Training Program for CVD Research, which funded my doctoral research and my work at the Duke Clinical Research Institute (DCRI). I am especially thankful for Dr. Sean O'Brien and Hillary Mulder, whose kindness and support made my three years at DCRI such a pleasurable experience.

I would like to express my gratitude for the amazing faculty, staff, and students in the North Carolina State Department of Statistics. I owe a huge thanks to my wonderful office mate Cole Manschot, who infused our work days with light and laughter, and to Sarah Fairfax, whose camaraderie brightened every class and assignment. I'd also like to thank my incredible North Carolina friends, who treated me like family when my own was so far away. Thank you to Arielle Valint, Tim Valint, Karen Manschot, Kasia Dobrzycka, and Kade Young.

I am incredibly blessed to have been surrounded by individuals who supported and nurtured my love of mathematics from an early age. I owe a special thanks to Mrs. Leslie North, whose exceptional teaching at Wadsworth High School inspired me to pursue an education and career in the mathematical sciences.

Above all, thank you to my parents, Sherri and Tim Rhodes. Thank you for teaching me the importance of education and the value of hard work. Thank you for believing that I can accomplish whatever I set my mind to. Thank you for your unconditional love and support. This dissertation would not be possible without you and every phone call you answered.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER

1

INTRODUCTION

Clinicians and patients must make treatment decisions at a series of key decision points throughout disease progression. When treating life-threatening diseases, prediction and optimization of remaining life can provide critical information at these decision points. Formally, the remaining life of a patient at time t, given the patient has survived up to time t, is referred to as "residual life." At each decision point, it is desirable to predict and optimize residual life using all available patient information. Patient information is often stored in electronic medical records (EMRs), electronic charts that contain patient healthcare data collected, managed, and consulted by authorized clinicians and staff within a healthcare organization (Gogia (2019)). EMRs frequently contain data on biomarkers, variables that are "objectively measured and evaluated as an indicator of normal biological processes, pathogenic processes, or pharmacologic responses to a therapeutic intervention" (Strimbu and Tavel (2010)). Biomarkers that are measured repeatedly on patients over time, such as blood pressure, ventilator dependence, and white blood cell count, are referred to as "longitudinal." Dynamic estimation methods leverage accruing longitudinal measurements to provide individualized, updated results as new observations becomes available.

In this thesis, we first introduce two statistical learning methods to dynamically predict

mean residual life (MRL) using longitudinal EMR data. We then introduce a method to estimate the value of a dynamic treatment regime, a set of sequential decision rules that return treatment decisions based on accumulating patient information. We define the value of a treatment regime in terms of restricted residual life, the minimum of residual life and remaining time until a pre-specified constant. We then present a statistical learning method to estimate an optimal dynamic treatment regime that maximizes expected restricted residual life.

Our presented methods are particularly motivated by application to septic patient care. Sepsis is a complex medical condition that involves severe infections with life-threatening organ dysfunction, and it is a leading cause of death worldwide (Singer et al. (2016)). Treatment of septic patients remains highly challenging, as the heterogeneity of the septic patient population leads to differing responses to medical intervention (László et al. (2015)). Consequently, prediction and optimization techniques that provide personalized, real-time results based on accumulating patient information are highly desirable for sepsis treatment. Accordingly, we study a data set of septic patients constructed from the Multiparameter Intelligent Monitoring Intensive Care database (MIMIC-III). MIMIC-III is a freely-available database comprised of de-identified health records for over 40,000 patients who stayed in the critical care units at Beth Israel Deaconess Medical Center between 2001 and 2012 (Johnson et al. (2016)). MIMIC-III contains data on patients' demographics, vital signs, laboratory measurements, medications, imaging reports, chart notes, procedure codes, diagnostic codes, hospital stay, and survival. For a complete description of the MIMIC-III database, refer to Johnson et al. (2016).

In 2016, the definitions and clinical criteria for sepsis and septic shock were updated in the *Third International Consensus Definitions for Sepsis and Septic Shock (Sepsis-3)* (Singer et al. (2016)). Sepsis-3 defines sepsis as a "life-threatening organ dysfunction caused by a dysregulated host response to infection" and provides clinical criteria for diagnosing septic patients (Singer et al. (2016)). Komorowski (2019) developed code to identify patients in MIMIC-III fulfilling the Sepsis-3 criteria. Komorowski's code pulls relevant physiological parameters for each patient from up to 24 hours preceding their sepsis diagnosis until 48 hours after. The code aggregates the data into four-hour time-windows, recording an appropriate summary statistic when several measurements were taken in the same time-window. We use the code to construct a data set of 48 covariates measured on approximately 21,000 patients in the ICU diagnosed with sepsis, 15 percent of whom have an observed time of death.

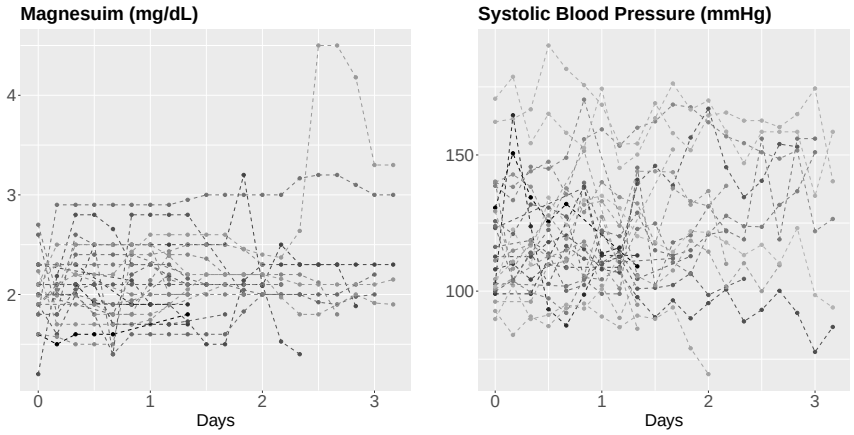

Figure 1.1: The longitudinal trajectories of magnesium and systolic blood pressure for 25 randomly selected septic patients in MIMIC-III.

The longitudinal trajectories of the biomarkers magnesium and systolic blood pressure are illustrated in Figure (1.1) for 25 randomly selected septic patients from MIMIC-III. The figures illustrate that longitudinal biomarker trajectories are often sophisticated functions of time that exhibit notable variation among patients. Moreover, EMR data often include a large number of longitudinal biomarkers, and the number of biomarker measurements often differs between patients. These complexities make it difficult to formulate parametric models that fully capture longitudinal biomarker processes and their relationship with residual life. In this thesis, we employ modern statistical learning techniques to predict and optimize residual life using EMR data with longitudinal covariates.

In Chapter 2, we introduce two methods to dynamically predict MRL using EMR data. In both methods, we begin by using long short-term memory networks (LSTMs) to construct encoded representations of the longitudinal biomarker trajectories, referred to as "context vectors." In our first method, the LSTM generalized linear model (LSTM-GLM), we dynamically predict MRL via a transformed MRL model that includes the context vectors as covariates. In our second method, the LSTM neural network (LSTM-NN), we dynamically predict MRL from the context vectors using a feed-forward neural network. We demonstrate the superior prediction performance of the proposed LSTM methods relative to competing methods in simulation studies, and we apply the proposed methods to predict the restricted MRL of septic patients in MIMIC-III. Ultimately, we demonstrate that the LSTM-GLM and the LSTM-NN leverage accruing patient information to produce individualized, real-time predictions of MRL that can inform patient care.

In Chapter 3, we present a method that leverages the context vectors introduced in Chapter 2 to estimate an optimal dynamic treatment regime from EMR data. An optimal treatment regime leads to the most favorable outcome on average when applied to a patient population, so identification of an optimal regime that maximizes residual life is highly desirable for septic patient care. Accordingly, we define the value of a fixed treatment regime to be the expected value of the sum of restricted residual life under the regime, where the sum is taken across all decision points at which a patient is alive. First, we introduce the residual life value estimator (ReLiVE), which estimates the value of a fixed regime. Building on ReLiVE, we present the ReLiVE Q-learning method (ReLiVE-Q), which uses backward induction to estimate an optimal treatment regime that maximizes the expected value of restricted residual life. We illustrate the utility of ReLiVE-Q in simulation studies, and we apply ReLiVE-Q to estimate an optimal treatment regime for septic patients in MIMIC-III. Ultimately, we demonstrate that ReLiVE-Q leverages accruing patient information to estimate personalized treatment regimes that maximize a clinically meaningful function of residual life.

CHAPTER

2

DYNAMIC PREDICTION OF RESIDUAL LIFE WITH LONGITUDINAL COVARIATES USING LONG SHORT-TERM MEMORY NETWORKS

2.1 Introduction

When making treatment decisions, clinicians and patients often desire accurate predictions of remaining life expectancy that leverage all available patient information, including longitudinal biomarker data. The National Institutes of Health (NIH) defines a biomarker as "a characteristic that is objectively measured and evaluated as an indicator of normal biological processes, pathogenic processes, or pharmacologic responses to a therapeutic intervention" (Strimbu and Tavel (2010)). Longitudinal biomarker measurements, such as blood pressure, ventilator dependence, and white blood cell count, are commonly available

in electronic medical records (EMRs). The recent proliferation of EMRs has led to a growing interest in using longitudinal biomarker data with "dynamic" prediction methods, which provide updated predictions as new biomarker measurements become available. Clinicians and patients are especially interested in using longitudinal biomarker data to dynamically predict mean residual life (MRL), the remaining life expectancy of a patient at time t, given the patient has survived up to time t.

MRL prediction is of particular interest for patients diagnosed with sepsis, a complex medical condition that involves severe infections with life-threatening organ dysfunction (Singer et al. (2016)). Sepsis is a leading cause of death worldwide (Singer et al. (2016)). Although international guidelines for sepsis treatment have been established, treating septic patients remains highly challenging (Evans et al. (2021)), as the heterogeneity of septic patient populations results in differing responses to medical intervention (László et al. (2015)). Dynamic predictions of mean residual life provide clinicians with individualized, real-time information that can help inform the treatment decisions of septic patients.

We dynamically predict the restricted mean residual life (RMRL) of septic patients in the intensive care unit (ICU) from EMR data. We conduct our study using a data set constructed from the Multiparameter Intelligent Monitoring Intensive Care database (MIMIC-III), described in Chapter 1. We study a single baseline covariate and 15 longitudinal biomarkers in the data set. The baseline covariate of interest is an indicator of whether the patient was previously admitted to the ICU during the given hospital stay. The longitudinal biomarkers of interest include two treatment variables: the median dosage of vasopressors provided to the patient in the given four-hour time-window and the cumulative amount of intravenous (IV) fluid administered to the patient since hospital admission. We also study a longitudinal indicator of mechanical ventilator dependence, as well as 12 vital signs and laboratory values: albumin, calcium, magnesium, hemoglobin, arterial lactate, arterial pH, fraction of inspired oxygen (FiO2), peripheral oxygen saturation (SpO2), Sequential Organ Failure Assessment (SOFA) score, respiratory rate, heart rate, and systolic blood pressure.

Biomarker trajectories of 25 randomly selected septic patients from MIMIC-III are illustrated in Figure (1.1). The figure demonstrates that longitudinal biomarker trajectories are sophisticated functions of time that exhibit notable variation among patients. Additionally, the number of biomarker measurements differs among patients. These complexities make it difficult to formulate a parametric model that fully captures the biomarker processes and their relationship with MRL.

The body of literature addressing how to model MRL with covariates measured only at

baseline is vast. Popular baseline MRL models include proportional MRL models (Maguluri and Zhang (1994)), additive MRL models (Chen (2007)), and transformed MRL models (Sun and Zhang (2009)). Although Sun et al. (2012) extended the family of transformed MRL models to accommodate time-dependent coefficients, none of the aforementioned models accommodate time-dependent covariates. Thus, they cannot be used to conduct dynamic prediction of MRL with longitudinal biomarker data.

A number of dynamic prediction models for survival risk have been developed via the landmarking approach (Zheng and Heagerty (2005), van Houwelingen (2007), van Houwelingen and Putter (2011), Rizopoulos et al. (2017), Zhu et al. (2019)). In a landmark analysis, a prediction model is fit at a series of time points, referred to as "landmark times," using only data collected prior to the landmark time on patients at risk at the landmark time. Lin et al. (2018) used the landmarking approach to incorporate longitudinal covariates into the transformed MRL model presented by Sun et al. (2012), effectively creating a dynamic prediction model for MRL.

To synthesize the longitudinal biomarker trajectories, Lin et al. (2018) proposed modeling the biomarkers using functional principal components analysis (FPCA). FPCA extracts dominant features from longitudinal trajectories as functional principal component (FPC) scores. Under the FPCA framework, Lin et al. (2018) introduced "window-specific FPC scores" that summarize a given longitudinal biomarker trajectory from baseline to the time of prediction. The authors proposed including the window-specific FPC scores as time-dependent covariates in their dynamic prediction model for MRL.

To calculate window-specific FPC scores at a given prediction time, Lin et al. (2018) presented a method that uses measurements collected from baseline to maximum follow-up time. Consequently, when the window-specific FPC scores are included as predictors in a dynamic MRL model, information collected after the time of prediction is used to predict MRL. This is undesirable in the dynamic prediction setting, where we wish to predict MRL using only the information available at the time of prediction.

Building on the work of Lin et al. (2018), we introduce two new methods to dynamically predict MRL from longitudinal biomarkers. The methods offer two potential advantages. First, the proposed methods uphold the dynamic nature of prediction. Second, the proposed methods may more effectively synthesize the complex longitudinal biomarker trajectories observed in MIMIC-III.

The proposed methods use long short-term memory networks (LSTMs) to construct "window-specific context vectors," which summarize the biomarker trajectories from base-

line to the time of prediction. LSTMs are especially suitable for constructing summaries of biomarker trajectories because they are capable of modeling complex, heterogeneous functions. Thus, LSTMs are an attractive alternative to FPCA for synthesizing the MIMIC-III biomarkers. To uphold the dynamic nature of prediction, the LSTMs construct the window-specific context vectors using only the information available at the time of prediction.

The first proposed method, the long short-term memory generalized linear model (LSTM-GLM), dynamically predicts MRL using a dynamic transformed MRL model that includes the baseline covariates and window-specific context vectors as predictors. The second proposed method, the long short-term memory neural network (LSTM-NN), dynamically predicts MRL from the baseline covariates and window-specific context vectors using a feed-forward neural network. We apply the LSTM-GLM and the LSTM-NN to dynamically predict the RMRL of septic patients in MIMIC-III. We demonstrate that the LSTM-GLM and the LSTM-NN produce accurate, individualized, dynamic predictions. Thus, the LSTM-GLM and the LSTM-NN can be used to inform the challenging treatment decisions of patients diagnosed with sepsis.

In Section 2.2, we introduce the dynamic transformed MRL model, and we present the LSTM-GLM and the LSTM-NN. In Section 2.3, we describe the procedure used to evaluate prediction performance. In Section 2.4, we present simulation studies, and in Section 2.5, we apply the proposed methods to predict the RMRL of septic patients in MIMIC-III. We conclude with a discussion of implications and open problems in Section 2.6.

2.2 Methods

2.2.1 Notation

Let there be $i = 1,\ldots,m$ patients, and let $t_{ij} \geq 0$ denote the time at which biomarker measurement j was collected on patient i, $j = 1,\ldots n_i$. We study $m = 20,954$ septic patients with $n_i \in [1,20]$ measurements. Let $T_i > 0$ and $C_i > 0$ denote the potential times to death and censoring, respectively, for patient i. We observe only $Y_i = \min(T_i, C_i)$ and $\Delta_i = I(T_i \leq C_i)$, the indicator of whether the death of patient i was observed ($\Delta_i = 1$) or censored ($\Delta_i = 0$). Let $\mathbf{X}_i$ denote the q-dimensional vector of baseline covariates measured on patient i. We study $q = 1$ indicator of whether patient i was previously admitted to the ICU during the given hospital stay. Let $\mathbf{Z}_i(t) = \{Z_{i1}(t),\ldots,Z_{ip}(t)\}$ denote the p-dimensional vector of longitudinal biomarkers measured on patient i at time $t \geq 0$. We study $p = 15$ longitudinal biomarkers,

including two treatment variables, an indicator of mechanical ventilator dependence, and 12 vital signs and laboratory values. Denote the covariate history of patient i at time $\tau \geq 0$ as $\mathscr{H}_i(\tau) = \{\mathbf{X}_i, \mathbf{Z}_i(t_{i1}), \ldots, \mathbf{Z}_i(t_{i\tau_i})\}$, where $\tau_i = \text{argmax}_j(t_{ij} < \tau)$. At time $\tau \geq 0$, the observed data for patient i are $\{Y_i, \Delta_i, \mathscr{H}_i(\tau)\}$. The mean residual life (MRL) of patient i at time $\tau \geq 0$ given the patient's covariate history is $E\{T_i - \tau \mid T_i > \tau, \mathscr{H}_i(\tau)\}$.

To avoid infinite remaining life expectancy and extreme propensity weights (see Section 2.2.2), we set a restricted lifetime of $L = 40$ days. Our potential survival time of interest is then $T_i^* = \min(T_i, L)$, and the restricted mean residual life (RMRL) for patient i at time $\tau \geq 0$ given the patient's covariate history is $E\{T_i^* - \tau \mid T_i > \tau, \mathscr{H}_i(\tau)\}$. We use the dynamic prediction methods presented in Section 2.2 to predict the RMRL of septic patients by redefining $Y_i = \min(T_i^*, C_i) = \min(T_i, C_i, L)$ and $\Delta_i = I(T_i^* \leq C_i) = I(T_i \leq C_i) + I\{L \leq \min(T_i, C_i)\}\{1 - I(T_i \leq C_i)\}$.

2.2.2 Dynamic Transformed MRL Model

Building on the work of Lin et al. (2018), we present a dynamic transformed MRL model that regresses residual life only on information collected prior to the time of prediction on patients at risk at the time of prediction. Let $f(\cdot)$ be a vector-valued function, and specify a prediction time $\tau \geq 0$. For patients with $Y_i > \tau$, define the υ-dimensional vector $\boldsymbol{\zeta}_i(\tau) = f\{\mathbf{Z}_i(t_{i1}), ..., \mathbf{Z}_i(t_{i\tau_i})\}$. Additionally, let $g(\cdot)$ be a pre-specified, non-negative link function that is twice continuously differentiable and strictly increasing. We specify the dynamic transformed MRL model as

$$E\{T_i - \tau | T_i > \tau, \mathscr{H}_i(\tau)\} = g\{\eta(\tau) + \boldsymbol{\gamma}(\tau)^T \mathbf{X}_i + \boldsymbol{\alpha}(\tau)^T \boldsymbol{\zeta}_i(\tau)\}, \tag{2.1}$$

where $\eta(\cdot)$ is a scalar, time-dependent parameter, $\boldsymbol{\gamma}(\cdot)$ is a q-dimensional, time-dependent parameter vector, and $\boldsymbol{\alpha}(\cdot)$ is a υ-dimensional, time-dependent parameter vector.

Define $w_i = \{\Delta_i I(Y_i > \tau)\}/\widehat{G}(Y_i)$, where $\widehat{G}(\cdot)$ is an estimate of the survival function of censoring time. We estimate the parameters in Equation (2.1) via a landmarking approach. Contrary to Lin et al. (2018), we do not adopt a supermodel approach for parameter estimation (van Houwelingen (2007); van Houwelingen and Putter (2011)). Instead, we pre-specify a set of positive prediction times $\mathscr{T}$. For each $\tau \in \mathscr{T}$, we use penalized maximum likelihood methods (Friedman et al. (2010)) to compute the values of $\eta(\cdot)$, $\boldsymbol{\gamma}(\cdot)$, and $\boldsymbol{\alpha}(\cdot)$ that minimize

the objective function

$$\frac{1}{2\sum_i w_i}\sum_{i=1}^{m} w_i\left[(Y_i-\tau)-g\{\eta(\tau)+\boldsymbol{\gamma}(\tau)^T\mathbf{X}_i+\boldsymbol{\alpha}(\tau)^T\boldsymbol{\zeta}_i(\tau)\}\right]^2+\lambda||\{\boldsymbol{\gamma}(\tau),\boldsymbol{\alpha}(\tau)\}||_1,$$

where λ is a scalar tuning parameter and $||\cdot||_1$ represents the L_1-norm. The inverse probability of censoring weights w_i account for censoring in the data. We assume censoring time is independent of the baseline covariates and longitudinal biomarkers, and we estimate $\widehat{G}(\cdot)$ using the Kaplan-Meier estimator. Alternatively, if censoring time is assumed to depend on only the baseline covariates, a Cox regression model can be used to estimate $\widehat{G}(\cdot)$.

We impose an L_1-penalty on $v+q$ regression parameters in the objective function to prevent overfitting. To ensure fair penalization, we apply proportion-of-maximum scaling (POMS) to the longitudinal biomarkers, such that

$$Z_{ik}(t)^{POMS}=\frac{Z_{ik}(t)-\min\limits_{i,u}\{Z_{ik}(u)\}}{\max\limits_{i,u}\{Z_{ik}(u)\}-\min\limits_{i,u}\{Z_{ik}(u)\}}.$$

Conducting dynamic prediction of MRL with $\boldsymbol{\zeta}_i(\tau)=\mathbf{Z}_i(\tau)$ is difficult in practice. Often, not all patients have longitudinal measurements $\mathbf{Z}_i(\tau)$ available at all desired prediction times $\tau\in\mathscr{T}$. In the MIMIC-III data set, longitudinal measurements are recorded at 4-hour time intervals. Over 80% of patients are missing at least one measurement, and over 25% of patients are missing at least 10 measurements, for all studied longitudinal biomarkers. Parametric models of $\mathbf{Z}_i(\cdot)$ that could be used to impute missing measurements are likely to be misspecified due to the complex, heterogeneous nature of biomarker processes. Moreover, regressing MRL only on the biomarker measurements taken at the time of prediction discards the information contained in the history of measurements.

Intuitively, it is desirable to select a function $f(\cdot)$ that summarizes the biomarker trajectories from baseline to prediction time τ. However, a simple summary function, such as average or slope, is unlikely to capture the complex trajectories of the longitudinal biomarkers. To address these complications, Lin et al. (2018) proposed summarizing the biomarker trajectories from baseline to prediction time using window-specific FPC scores. Alternatively, we propose summarizing the trajectories using window-specific context vectors constructed by LSTM autoencoders.

2.2.3 Context Vector Construction

The window-specific context vector $\boldsymbol{\psi}_{ik}(\tau)$ is an encoded representation of the trajectory of biomarker k from baseline to prediction time τ for patient i. At each prediction time τ, a distinct LSTM autoencoder is used to construct each of the $k = 1, \ldots, p$ sets of window-specific context vectors. An LSTM autoencoder is an unsupervised neural network that learns how to best encode temporal input into a context vector, so it can then reconstruct the original input from that context vector. To uphold the dynamic nature of prediction, the LSTM autoencoder used to construct $\boldsymbol{\psi}_{ik}(\tau)$ only accepts as input the biomarker measurements collected prior to time τ on patients at risk at time τ, where patient i is defined to be at risk at time τ if $Y_i > \tau$.

An LSTM autoencoder, which consists of an encoder and a decoder, is illustrated in Figure (2.1). Proportion-of-maximum scale the biomarker data, and let $\mathbf{Z}_{ik}^{\tau} = \{Z_{ik}(t_{i1}), \ldots, Z_{ik}(t_{i\tau_i})\}$ be the n_{ik}^{τ}-dimensional vector of scaled measurements of biomarker k collected on patient i prior to time τ. For each patient i with $Y_i > \tau$, input $\mathbf{Z}_{ik}^{\tau}$ into the encoder. The encoder compresses $\mathbf{Z}_{ik}^{\tau} \in \mathbb{R}^{n_{ik}^{\tau}}$ into the window-specific context vector $\boldsymbol{\psi}_{ik}(\tau) \in \mathbb{R}^{s}$. The decoder then constructs an estimate of the input scaled biomarker measurements, $\widehat{\mathbf{Z}}_{ik}^{\tau}$, from $\boldsymbol{\psi}_{ik}(\tau)$. The autoencoder is trained to minimize the reconstruction error

$$\sum_{i:Y_i>\tau} \sum_{j=1}^{\tau_i} \{Z_{ik}(t_{ij}) - \widehat{Z}_{ik}(t_{ij})\}^2.$$

After training the autoencoder, the decoder can be removed from the network, so the encoder outputs the context vector $\boldsymbol{\psi}_{ik}(\tau)$ directly to the user.

In an LSTM autoencoder, both the encoder and decoder are a type of recurrent neural network called a "long short-term memory network." Recurrent neural networks (RNNs) are a class of artificial neural networks designed to process sequential data. RNNs contain a feedback loop that enables information from previous time steps to be passed to future time steps. The information is passed in an s-dimensional vector $\mathbf{h}_i(\cdot)$, referred to as a "hidden vector." The parameters in an RNN are estimated via the back propagation through time (BPTT) algorithm (Werbos (1990)). In the BPTT algorithm, derivatives are multiplied across time steps. Consequently, in RNNs with a large number of time steps, if the derivatives are large, the gradients will increase exponentially and "explode." If the derivatives are small, the gradients will decrease exponentially and "vanish." This is referred to as the "vanishing and exploding gradient problem" (Aggarwal (2018)). The vanishing and exploding gradient

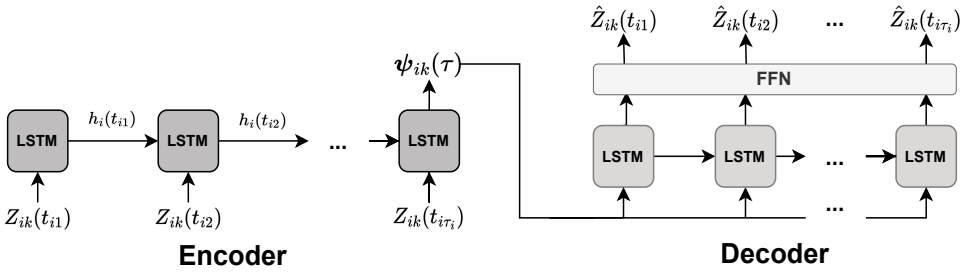

Figure 2.1: An LSTM autoencoder at time τ, comprised of an encoder and a decoder. Both the encoder and decoder consist of a series of LSTM temporal units, labelled "LSTM." The decoder also contains a feed-forward neural network layer, labelled "FFN." The encoder compresses the input biomarker measurements $Z_{ik}(\cdot)$ into the window-specific context vector $\psi_{ik}(\tau)$, and the decoder attempts to reconstruct the original biomarker measurements from $\psi_{ik}(\tau)$. Information is passed between the LSTM temporal units via the hidden vectors $h_i(\cdot)$.

problem makes it difficult for simple RNNs to capture long-term dependencies. LSTMs were designed especially to mitigate the vanishing and exploding gradient problem.

An LSTM can be conceptualized as a network of temporal units, with a single temporal unit corresponding to each time step in the data. Figure (2.2) depicts the LSTM temporal unit corresponding to time t_{ij}. LSTM networks mitigate the vanishing and exploding gradient problem using an s-dimensional vector $\mathbf{c}_i(\cdot)$, referred to as the "cell state." Conceptually, the cell state can be thought of as a pseudo long-term memory that retains information from previous time steps (Aggarwal (2018)). The cell state is controlled by three s-dimensional gate control signals, the input gate $\mathbf{i}_i(\cdot)$, the forget gate $\mathbf{f}_i(\cdot)$, and the output gate $\mathbf{o}_i(\cdot)$. These gate control signals determine which information in the cell state is updated, discarded, and output to the next time step, respectively. Let $\mathbf{U}_i(t_{ij})$ represent the r-dimensional input vector for patient i at time t_{ij}, and let $\mathbf{h}_i(t_{i,j-1})$ represent the s-dimensional hidden vector output for patient i at time $t_{i,j-1}$. Then the three gate control signals for patient i are characterized by the equations

$$\mathbf{\Omega}_i(t_{ij}) = \sigma\{\mathbf{W}_\Omega \mathbf{U}_i(t_{ij}) + \mathbf{Q}_\Omega \mathbf{h}_i(t_{i,j-1}) + \mathbf{v}_\Omega\}, \qquad \sigma(x) = (1+e^{-x})^{-1},$$

where $\mathbf{\Omega} = \mathbf{i}, \mathbf{f}, \mathbf{o}$, $\mathbf{W}_\Omega$ is an $s \times r$ parameter matrix, $\mathbf{Q}_\Omega$ is an $s \times s$ parameter matrix, and $\mathbf{v}_\Omega$ is an $s \times 1$ bias vector. Note, $\mathbf{W}_\Omega$, $\mathbf{Q}_\Omega$, and $\mathbf{v}_\Omega$ are gate-specific and temporally-shared.

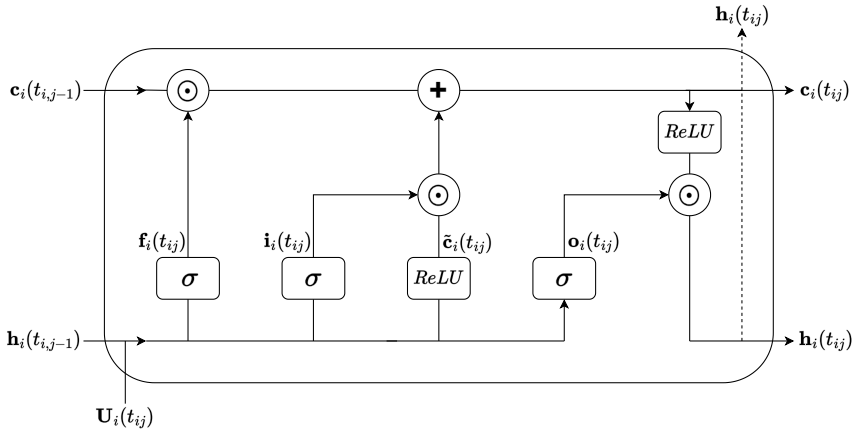

Figure 2.2: An LSTM temporal unit at time t_{ij}.

At time t_{ij}, an s-dimensional candidate cell state for patient i, $\tilde{\mathbf{c}}_i(t_{ij})$, is computed as

$$\tilde{\mathbf{c}}_i(t_{ij}) = \text{ReLU}\{\mathbf{W}_c\mathbf{U}_i(t_{ij}) + \mathbf{Q}_c\mathbf{h}_i(t_{i,j-1}) + \mathbf{v}_c\}, \qquad \text{ReLU}(x) = \max(0, x),$$

where $\mathbf{W}_c$ is an $s \times r$ parameter matrix, $\mathbf{Q}_c$ is an $s \times s$ parameter matrix, and $\mathbf{v}_c$ is an $s \times 1$ bias vector. Again, $\mathbf{W}_c$, $\mathbf{Q}_c$, and $\mathbf{v}_c$ are temporally-shared parameters.

At time t_{ij}, the cell state for patient i, $\mathbf{c}_i(t_{ij})$, is then computed as

$$\mathbf{c}_i(t_{ij}) = \mathbf{f}_i(t_{ij}) \odot \mathbf{c}_i(t_{i,j-1}) + \mathbf{i}_i(t_{ij}) \odot \tilde{\mathbf{c}}_i(t_{ij}),$$

where $\odot$ represents the Hadamard product.

Ultimately, each temporal unit outputs a hidden vector for patient i, $\mathbf{h}_i(t_{ij})$, computed as

$$\mathbf{h}_i(t_{ij}) = \mathbf{o}_i(t_{ij}) \odot \text{ReLU}\{\mathbf{c}_i(t_{ij})\}.$$

The hidden vector is then passed to the next temporal unit. Additionally, it may be output to the next layer in the LSTM autoencoder.

At each prediction time τ, for each of the $k = 1, \ldots, p$ biomarkers, we train a separate LSTM autoencoder to construct the window-specific context vectors $\boldsymbol{\psi}_{ik}(\tau)$ for all patients i such that $Y_i > \tau$. Let the superscript e signify elements of the encoder, and let the superscript d signify elements of the decoder. Each encoder accepts as input the scaled measurements of biomarker k collected at times $t_{ij} < \tau$ on patients with $Y_i > \tau$. For a given patient i, each biomarker measurement is input into a separate LSTM temporal unit, so $\mathbf{U}_i^e(t_{ij}) = Z_{ik}(t_{ij})$

and $r = 1$ for $j = 1, \ldots, \tau_i$. The hidden vector for patient i output by the last temporal unit is taken to be the context vector for patient i, so $\boldsymbol{\psi}_{ik}(\tau) = \mathbf{h}_i^e(t_{i\tau_i})$.

Similar to the encoder, the decoder contains an LSTM temporal unit corresponding to each biomarker measurement. In the decoder, each LSTM temporal unit accepts the context vector for patient i as input, so $\mathbf{U}_i^d(t_{ij}) = \boldsymbol{\psi}_{ik}(\tau)$ and $r = s$ for $j = 1, \ldots, \tau_i$. Each temporal unit outputs the hidden vector $\mathbf{h}_i^d(t_{ij})$, which is fed into a feed-forward neural network (FFN) layer. At each measurement time $j = 1, \ldots, \tau_i$, the FFN layer constructs an estimate of the scaled biomarker measurement for patient i, $\widehat{Z}_{ik}(t_{ij})$, from the hidden vector $\mathbf{h}_i^d(t_{ij})$ via linear regression. Specifically,

$$\widehat{Z}_{ik}(t_{ij}) = \mathbf{W}_n \mathbf{h}_i^d(t_{ij}) + \upsilon_n,$$

where $\mathbf{W}_n$ is an s-dimensional, temporally-shared parameter vector, and υ_n is a temporally-shared scalar bias.

As previously mentioned, the LSTM autoencoder is trained to minimize the reconstruction error $\sum_{i:Y_i>\tau} \sum_{j=1}^{\tau_i} \{Z_{ik}(t_{ij}) - \widehat{Z}_{ik}(t_{ij})\}^2$. After training the autoencoder for biomarker k at prediction time τ, the window-specific context vector $\boldsymbol{\psi}_{ik}(\tau)$ can be extracted for each patient i such that $Y_i > \tau$. Because each window-specific context vector is constructed using only measurements taken at times $t_{ij} < \tau$ on patients with $Y_i > \tau$, $\boldsymbol{\psi}_{ik}(\tau)$ can be used to conduct dynamic prediction. Moreover since the number and timing of biomarker measurements can differ between patients, imputation of missing or irregularly measured biomarkers is unnecessary.

Each LSTM autoencoder has several hyperparameters that influence how well the output window-specific context vectors summarize the input biomarker trajectories. Important hyperparameters include the dimension of the window-specific context vector, s, and the number of times the BPTT algorithm processes the entire data set, referred to as the number of "training epochs." Too many training epochs can lead to overfitting the data, and too few can lead to underfitting. These hyperparameters can be selected via traditional tuning methods such as hold-out validation or cross-validation. In Appendix A, Section A.1, we present an automated approach for selecting these hyperparameters for the LSTM-GLM.

2.2.4 LSTM-GLM

First, we dynamically predict MRL using a dynamic transformed MRL model that includes the baseline covariates and window-specific context vectors as predictors. Let $\boldsymbol{\psi}_i(\tau) = \{\boldsymbol{\psi}_{i1}(\tau), \ldots, \boldsymbol{\psi}_{ip}(\tau)\}$ be an sp-dimensional vector containing the p biomarker-specific, window-specific context vectors for patient i at prediction time τ. Additionally, let $g(\cdot)$ be a pre-specified, non-negative link function that is twice continuously differentiable and strictly increasing. The LSTM-GLM is specified as

$$E\{T_i - \tau | T_i > \tau, \mathscr{H}_i(\tau)\} = g\{\eta(\tau) + \boldsymbol{\gamma}(\tau)^T \mathbf{X}_i + \boldsymbol{\alpha}(\tau)^T \boldsymbol{\psi}_i(\tau)\}. \tag{2.2}$$

The LSTM-GLM is a special case of the dynamic transformed MRL model specified in Equation (2.1), where $\boldsymbol{\zeta}_i(\tau) = \boldsymbol{\psi}_i(\tau)$. Accordingly, we estimate the parameters in Equation (2.2) via penalized maximum likelihood by adopting the landmarking approach detailed in Section 2.2.2. Because a separate context vector is constructed for each biomarker, the parameter estimates of the LSTM-GLM can be used to gain insight into the relationship between the longitudinal biomarkers and mean residual life.

2.2.5 LSTM-NN

Next, we introduce an alternative method for dynamically predicting MRL from window-specific context vectors. The LSTM-NN dynamically predicts MRL using a feed-forward neural network that accepts the baseline covariates and window-specific context vectors as input. Compared to generalized linear models, feed-forward neural networks are more capable of modeling complex functional relationships. In fact, a neural network with a single non-linear hidden layer and a single linear output layer can compute almost any function (Aggarwal (2018)). This makes neural networks ideal for modeling the complex relationship between MRL and the biomarker processes.

The LSTM-NN is a feed-forward neural network comprised of one or more hidden layers and an output layer. The first hidden layer takes the baseline covariates and context vectors as input, and the output layer produces estimates of MRL. Additional hidden layers can be added to the LSTM-NN to tailor the network's flexibility to the complexity of the studied data set. For our simulation studies and the MIMIC-III data application, we consider an LSTM-NN with two hidden layers, as illustrated in Figure (2.3).

For a given prediction time τ, the first hidden layer, *FFN1*, accepts the baseline covariates

Figure 2.3: The LSTM-NN at prediction time τ.

$\mathbf{X}_i$ and the window-specific context vectors $\boldsymbol{\psi}_i(\tau)$ as input. *FFN1* then computes and outputs the u-dimensional hidden vector $\mathbf{O}_{i1}(\tau)$, which is calculated as

$$\mathbf{O}_{i1}(\tau) = \tanh[\ \mathbf{W}_1\{\mathbf{X}_i, \boldsymbol{\psi}_i(\tau)\} + \mathbf{v}_1\ ], \qquad \tanh(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}},$$

where $\mathbf{W}_1$ is a $u \times (q + sp)$ parameter matrix, and $\mathbf{v}_1$ is a u-dimensional bias vector.

The second hidden layer, *FFN2*, then accepts $\mathbf{O}_{i1}(\tau)$ as input. *FFN2* computes and outputs the u-dimensional hidden vector $\mathbf{O}_{i2}(\tau)$, which is calculated as

$$\mathbf{O}_{i2}(\tau) = \tanh[\ \mathbf{W}_2\mathbf{O}_{i1}(\tau) + \mathbf{v}_2\ ],$$

where $\mathbf{W}_2$ is a $u \times u$ parameter matrix, and $\mathbf{v}_2$ is a u-dimensional bias vector.

The output layer, *FFN3*, then accepts $\mathbf{O}_{i2}(\tau)$ as input. *FFN3* computes and outputs the estimate of MRL for patient i at time τ as

$$\widehat{R}_i(\tau) = \mathbf{W}_3\mathbf{O}_{i2}(\tau) + \upsilon_3,$$

where $\mathbf{W}_3$ is a u-dimensional parameter vector, and υ_3 is a scalar bias.

As with the LSTM-GLM, we estimate the parameters of the LSTM-NN via a landmarking approach. First, we specify a set of positive prediction times $\mathscr{T}$. Then for each $\tau \in \mathscr{T}$, we train the LSTM-NN to minimize the objective function

$$\left[\sum_{i=1}^{m} \frac{\Delta_i I(Y_i > \tau)}{\widehat{G}(Y_i)}\right]^{-1} \sum_{i=1}^{m} \frac{\Delta_i I(Y_i > \tau)}{\widehat{G}(Y_i)} \left[(Y_i - \tau) - \widehat{R}_i(\tau)\right]^2 + \lambda||\mathbf{W}_1||_2^2 + \lambda||\mathbf{W}_2||_2^2.$$

Again, we use inverse probability of censoring weights to account for censoring, where $\widehat{G}(\cdot)$ is the Kaplan-Meier estimate of the survival function of censoring time. Additionally, we apply an L_2-penalty to the parameter matrices $\mathbf{W}_1$ and $\mathbf{W}_2$ to prevent overfitting. Here, $||\cdot||_2$ represents the L_2-norm, and λ is a tuning parameter for the L_2-penalties.

The LSTM-NN provides more flexibility in modeling the relationship between MRL

and the longitudinal biomarkers than the LSTM-GLM. However, the complexity of the feed-forward neural network makes it difficult to interpret the relationship between the biomarkers and MRL. Moreover, the LSTM-NN has a number of hyperparameters that must be tuned, including the dimension of the parameter matrices u, the tuning parameter for the L_2-penalty λ, and the number of epochs used to train the LSTM-NN. These hyperparameters can be tuned using traditional processes such as hold-out validation or cross validation. However, these processes are computationally-intensive, and imperfect tuning can result in poor prediction performance.

2.3 Performance Evaluation

2.3.1 Comparative Methods

For the LSTM-GLM and the LSTM-NN to have utility in the clinical setting, the models must produce accurate dynamic predictions of MRL relative to competing dynamic prediction methods. Consequently, we evaluate the prediction performance of the LSTM-GLM and the LSTM-NN relative to six variations of the dynamic transformed MRL model specified in Equation (2.1). For each of the six dynamic transformed MRL models, we define a distinct function of the history of longitudinal biomarker measurements, $\boldsymbol{\zeta}_i(\tau) = f\{\mathbf{Z}_i(t_{i1}), ..., \mathbf{Z}_i(t_{i\tau_i})\}$. To maintain the dynamic nature of prediction, we construct $\boldsymbol{\zeta}_i(\tau)$ using only biomarker measurements taken at times $t_{ij} < \tau$ on patients with $Y_i > \tau$. First, we define $\boldsymbol{\zeta}_i^{(B)}(\tau)$ to be a vector of the baseline biomarker measurements. Second, we define $\boldsymbol{\zeta}_i^{(L)}(\tau)$ to be a vector of the biomarker measurements collected most recently before prediction time τ (i.e. the "last-value carried forward"). Third, we define $\boldsymbol{\zeta}_i^{(A)}(\tau)$ to be a vector of the average value of each biomarker prior to time τ. Next, we define two vectors containing the intercept and slope of each biomarker regressed against time. The first vector, $\boldsymbol{\zeta}_i^{(S)}(\tau)$, is formed by conducting an independent linear regression on each patient for each biomarker. The second vector, $\boldsymbol{\zeta}_i^{(M)}(\tau)$, is formed by fitting a single linear mixed effects model with a random intercept and slope to all patients for each biomarker. Lastly, we define $\boldsymbol{\zeta}_i^{(F)}(\tau)$ to be a vector of FPC scores computed independently on each biomarker. For each biomarker, $\boldsymbol{\zeta}_i^{(F)}(\tau)$ contains the minimum number of FPC scores required to explain 99 percent of the total variance of that biomarker. We provide technical specifications for each comparative method in Appendix A, Section A.2.

2.3.2 Performance Metrics

To evaluate prediction performance, we focus on measures of calibration and discrimination. In the MIMIC-III data set, the survival time of interest T_i^* was observed for only 15 percent of patients. Consequently, it is important for the measures of calibration and discrimination to account for censoring. Let $\widehat{RMRL}_i(\tau)$ represent a given model's estimate of RMRL for patient i at time τ. We assess the calibration of each model via the inverse probability of censoring weighted mean square error

$$\frac{1}{\sum_i \frac{\Delta_i I(Y_i>\tau)}{\widehat{G}(Y_i)}} \sum_i \frac{\Delta_i I(Y_i>\tau)}{\widehat{G}(Y_i)} \Big[(Y_i-\tau) - \widehat{RMRL}_i(\tau) \Big]^2.$$

We refer to this quantity as the "testing loss."

In addition to calibration, we assess each model's discrimination, i.e. its ability to accurately predict who among a given pair of patients will live longer. We compute the following discrimination metric based on Harrell's C-Index (Harrell et al. (1996)),

$$\frac{\sum_{i\neq j} I\{c_i(\tau)>c_j(\tau)\} I\{\widehat{c}_i(\tau)>\widehat{c}_j(\tau)\} \Delta_j}{\sum_{i\neq j} I\{c_i(\tau)>c_j(\tau)\} \Delta_j},$$

where $c_i(\tau) = Y_i - \tau$ and $\widehat{c}_i(\tau) = \widehat{RMRL}_i(\tau)$. We refer to this quantity as the "testing C-index."

2.3.3 Software

We conduct the simulation studies and MIMIC-III data application in Python and R. We build and train the LSTM autoencoders and LSTM-NNs in Python using the **Keras** library (Chollet (2015)). We fit the LSTM-GLMs and the six dynamic transformed MRL models in R using the **glmnet** package (Friedman et al. (2010)). We compute the Kaplan-Meier estimate of the survival function of censoring time, $\widehat{G}(\cdot)$, in R using the **survival** package (Therneau and Grambsch (2000)). We leverage the **fdapace** R package (Gajardo et al. (2021)) to construct the FPCA vectors $\boldsymbol{\zeta}_i^{(F)}$, and we use the **lme4** R package (Bates et al. (2015)) to construct the linear regression vectors $\boldsymbol{\zeta}_i^{(S)}$ and the mixed effects vectors $\boldsymbol{\zeta}_i^{(M)}$.

2.4 Simulations

We conduct simulation studies to assess the prediction performance of the LSTM-GLM and the LSTM-NN relative to the performance of the six dynamic transformed MRL models described in Section 2.3.1. We generate a single data set of covariates for $m = 5000$ patients. For each patient, we generate a single baseline covariate $X_i \overset{iid}{\sim} \mathbb{U}(0,1)$, where $\mathbb{U}$ denotes the uniform distribution. Let $(\tau_1, \tau_2, \ldots, \tau_{19}) = (0, 0.5, \ldots, 9)$, and let $\mathbb{N}(\mu, \Sigma)$ denote a normal distribution with mean μ and variance-covariance Σ. For each patient, we generate a single longitudinal biomarker at $j = 1, 2, \ldots, 19$ patient-specific measurement times $t_{ij} = \min(0, \tau_j + \varepsilon_{ij})$, where $\varepsilon_{ij} \overset{iid}{\sim} \mathbb{N}(0, 0.05^2)$. We generate the longitudinal biomarker measurements as $Z_i(t_{ij}) = B_i(t_{ij}) + \epsilon_{ij}$, where $\epsilon_{ij} \overset{iid}{\sim} \mathbb{N}(0, 0.5^2)$ represents the measurement error, and $B_i(\cdot)$ is a piecewise linear mixed effects model with 8 interior knots. Specifically,

$$B_i(t) = a + b_{i0} + (c_1 + b_{i1})t + \sum_{j=2}^{10} I\{t > (j-1)\}(c_j + b_{ij})\{t - (j-1)\},$$

where $a = -2$ and $(c_1, c_2, \ldots, c_{10}) = (4, -7, 5, -2.5, 3.5, -5, 1.5, 2, -2, 1)$. Let $\text{rep}(x, y)$ denote a vector containing x repeated y times. Then $(b_{i0}, b_{i1}, \ldots, b_{i10}) \sim \mathbb{N}(\mathbf{0}, \mathbf{D})$, where $\mathbf{D}$ is the diagonal matrix $\mathbf{D} = \text{diag}\{1, \text{rep}(0.05, 5), \text{rep}(0.01, 5)\}$. Figure (2.4) illustrates the longitudinal biomarker trajectories of 25 randomly selected patients in the generated data set of covariates.

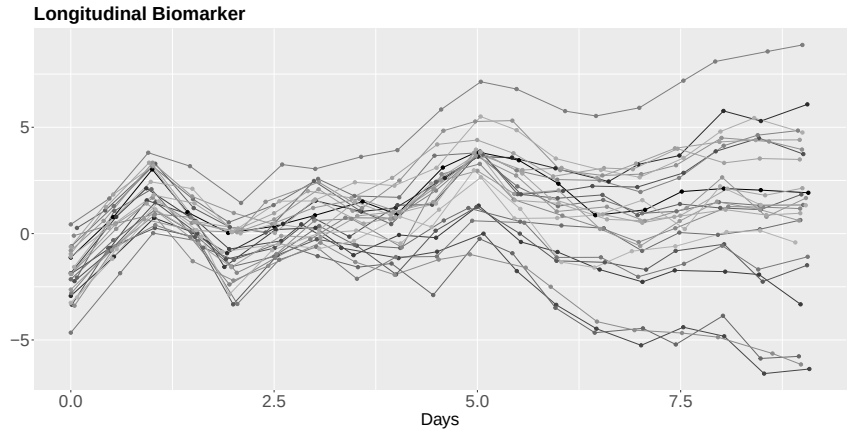

Figure 2.4: The biomarker trajectories of 25 randomly selected patients in the simulated covariate data set.

We conduct two simulation studies. In each study, we conduct 500 simulations using the aforementioned data set of covariates. For each of the 500 simulations, we generate a new data set of survival times T_i and censoring times C_i for all patients $i = 1, \ldots, 5000$. In both studies, we generate the censoring times as $C_i \overset{iid}{\sim} \mathbb{U}(0, 100)$. Conversely, we generate the survival times T_i using a different model for each study. In the first study, we generate T_i according to the accelerated failure time (AFT) model $\nu_i = \int_0^{T_i} \exp\{\beta_1 B_i(s) + \beta_2 X_i\} ds$, where $\beta_1 = \beta_2 = 1$, $\nu_i = \exp(\theta_i)$, and $\theta_i \overset{iid}{\sim} \mathbb{N}(3, 1)$. In the second study, we generate T_i according to a Cox proportional hazards model with an exponential baseline hazard function. Specifically, $h_i\{t \mid B_i(t), X_i\} = \lambda \exp\{\beta_1 B_i(t) + \beta_2 X_i\}$, where $\beta_1 = \beta_2 = 1$ and $\lambda = 0.05$. In both studies, we impose a restricted lifetime of $L = 50$. We describe the technical details of the survival time generation process in Appendix A, Section A.3.

We conduct prediction at time $\tau = 5$, so $n_i \in \{10, 11\}$. Across the 500 simulated AFT data sets, between 1824 and 1958 (mean=1886) patients are at risk at $\tau = 5$. Across the 500 simulated Cox data sets, between 1368 and 1502 (mean=1438) patients are at risk at $\tau = 5$. For both the AFT and Cox simulations, between 14% and 20% (mean=17%) of the patients at-risk at $\tau = 5$ are censored.

We log-transform the observed restricted residual lifetimes $Y_i - \tau$, and we define the link function for the LSTM-GLM and the six dynamic transformed MRL models to be the identity $g(x) = x$. Since we consider only a single longitudinal biomarker, we are not concerned with overfitting the data. Consequently, we set the tuning parameter to $\lambda = 0$ in all dynamic prediction models.

We specify the dimension of the window-specific context vectors to be $s = 5$, and we train the LSTM autoencoders for 500 epochs. We specify the dimension of the LSTM-NN parameter matrices to be $u = 3$, and we train the LSTM-NNs for 5000 epochs. We train all neural networks using the Adam optimization algorithm (Kingma and Ba (2017)). For the LSTM autoencoders, we use cross-validation to adaptively select the learning rate from the options $\{1e^{-3}, 1e^{-4}\}$ (O'Malley et al. (2019)). For the LSTM-NNs, we select a fixed learning rate of $1e^{-3}$.

For each simulation, we randomly divide the data into a training data set and a testing data set, stratifying on the censoring status Δ_i. We specify the training and testing data sets to each include 50 percent of the patients at risk at time $\tau = 5$. We estimate the survival function of censoring time $\widehat{G}(\cdot)$ independently on the training and testing data sets. We then fit each model on the training data set and compute the performance metrics specified in

Section 2.3.2 on the testing data set. We provide the Python and R code used to conduct the simulation studies at `https://github.com/gmrhodes/LSTM-GLM_LSTM-NN`.

We plot the distributions of the 500 testing losses and 500 testing C-indexes for each of the eight studied dynamic prediction models in Figure (2.5). For both the AFT and Cox simulations, the LSTM-NN results in the lowest median testing loss, followed by the LSTM-GLM. The LSTM-GLM consistently results in the highest median testing C-index. The LSTM-NN results in the second-highest median testing C-index for the AFT simulations and in the third-highest for the Cox simulations, where it is surpassed by the FPCA model. Overall, the simulation studies indicate that the LSTM-GLM and the LSTM-NN exhibit better calibration and discrimination than the baseline, last-value carried forward, average, linear regression, and mixed effects models, and that they exhibit at least comparable performance to the FPCA model. Thus, the simulation studies support that the LSTM-GLM and the LSTM-NN are useful tools for dynamically predicting MRL from longitudinal biomarker data.

In Appendix A, Section A.4, we repeat the simulation studies, reducing both the measurement error and the variation in measurement times. Compared to the simulations studies described above, the supplemental simulation studies demonstrate a more significant improvement in calibration and discrimination for the LSTM-GLM and the LSTM-NN relative to competing methods. Thus, the supplemental simulation studies indicate that the LSTM-GLM and the LSTM-NN are especially useful for producing accurate dynamic predictions of MRL in settings where the longitudinal biomarkers are measured using precise instruments.

2.5 Application to MIMIC-III

2.5.1 Data Analysis

We dynamically predict the RMRL of septic patients in the ICU from EMR data using the LSTM-GLM and the LSTM-NN. To evaluate the utility of the proposed methods in this clinical setting, we compare the prediction performance of the LSTM-GLM and the LSTM-NN to the performance of the six dynamic transformed MRL models described in Section 2.3.1. We conduct the study on the MIMIC-III data set described in Chapter 1. The distributions of the unrestricted and restricted survival times of the septic patients are depicted in Figure (2.6). Because the distribution of restricted survival times is notably

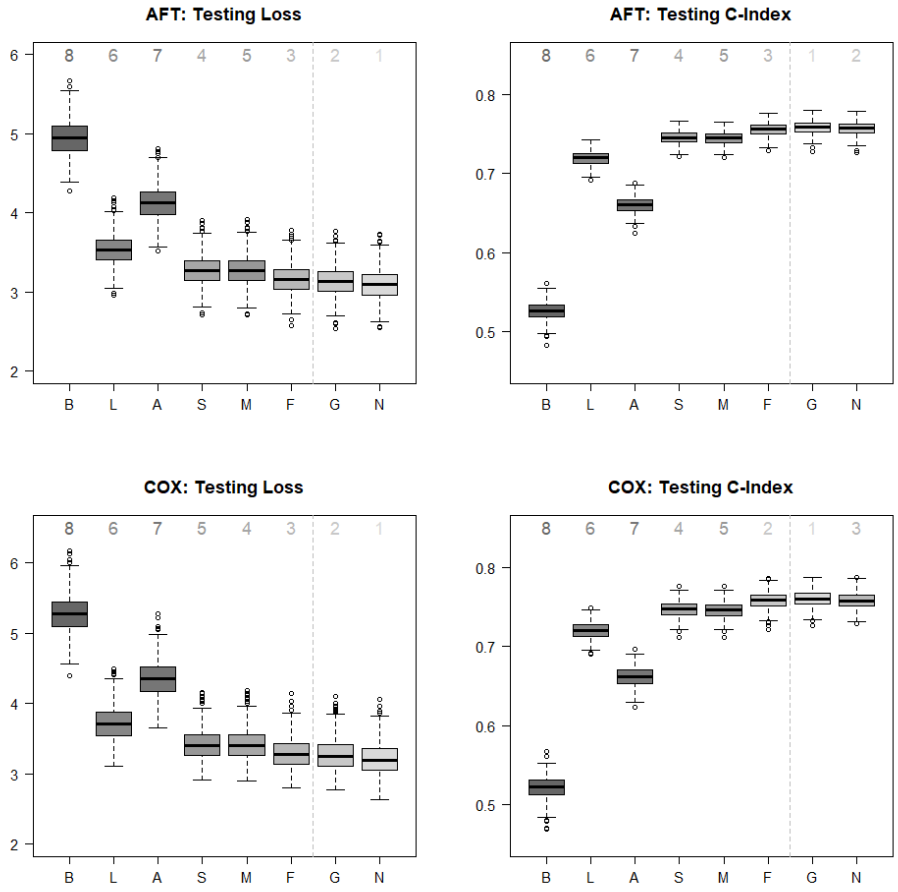

Figure 2.5: Distributions of the 500 testing losses and 500 testing C-indexes for each of the eight dynamic prediction models in the simulation studies. The models resulting in the lowest median testing loss and the highest median testing C-index are labelled 1. The models resulting in the highest median testing loss and the lowest median testing C-index are labelled 8. The six dynamic transformed MRL models are labelled according to their formulation of $\boldsymbol{\zeta}_i(\tau)$. "B" represents the baseline vector. "L" represents the last-value carried forward vector. "A" represents the average vector. "S" represents the linear regression vector. "M" represents the mixed effects vector. "F" represents the FPCA vector. Furthermore, "G" represents the LSTM-GLM, and "N" represents the LSTM-NN.

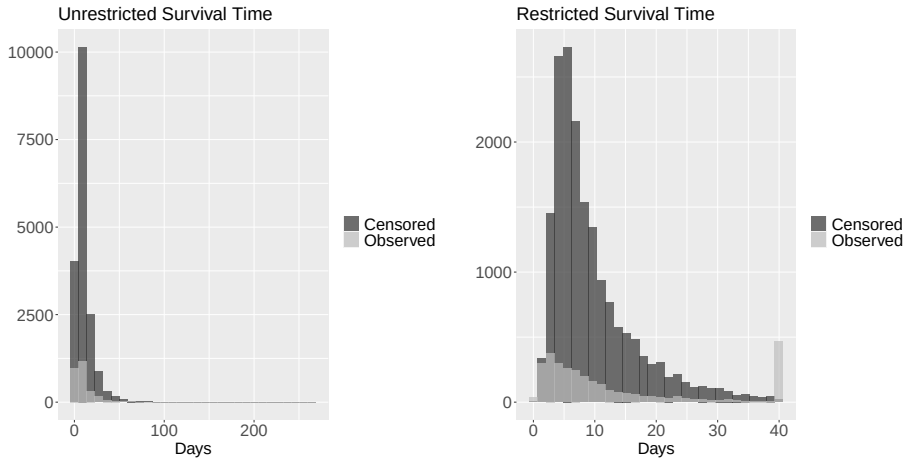

Figure 2.6: LEFT: The distribution of unrestricted survival time, stratified by censoring status. RIGHT: The distribution of survival time restricted to $L = 40$ days, stratified by censoring status.

right skewed and RMRL is non-negative by definition, we log-transform the observed restricted residual lifetimes $Y_i - \tau$. We then define the link function for the LSTM-GLM and the six dynamic transformed MRL models to be the identity $g(x) = x$. Because we conduct our study using $p = 15$ longitudinal biomarkers, we face the risk of overfitting the data. Consequently, we select the L_1-penalty tuning parameter λ via 5-fold cross validation for the LSTM-GLM and the six dynamic transformed MRL models.

We provide a detailed description of the hyperparameter selection process for the LSTM-GLM and the LSTM-NN in Appendix A, Section A.5. Let s denote the dimension of the window-specific context vectors, and let ep_a denote the number of epochs used to train the LSTM autoencoders. At each prediction time $\tau \in \mathscr{T}$, we construct four sets of window-specific context vectors using the hyperparameter settings $(s, ep_a) \in \{(3,150),(5,150),(5,300),(7,300)\}$, and we fit four LSTM-GLMs that each regress on one of the four sets of context vectors. For each $\tau \in \mathscr{T}$, we define the LSTM-GLM that results in the lowest median testing loss to be the "best" LSTM-GLM at time τ. The hyperparameter settings of the best LSTM-GLM at each $\tau \in \mathscr{T}$ can be seen in Table (2.1).

Additionally, we fit an "automated" LSTM-GLM, which selects the hyperparameter settings of the window-specific context vectors via cross-validation. We describe the technical details of the automated hyperparameter selection process in Appendix A, Section A.1.

Let λ denote the LSTM-NN L_2-penalty tuning parameter, let u denote the dimension of

Table 2.1: The hyperparameter settings of the best LSTM-GLM and the best LSTM-NN at each prediction time $\tau \in \mathscr{T}$, where the "best" model is defined to be the one resulting in the lowest median testing loss.

Prediction Time	LSTM-GLM		LSTM-NN				
Days	s	ep_a	s	ep_a	λ	u	ep_n
1	3	150	7	300	0.01	1	2000
1.5	7	300	7	300	0.005	2	3000
2	3	150	7	300	0.005	2	2000
2.5	5	300	7	300	0.01	2	2000
3	7	300	7	300	0.01	2	3000

the LSTM-NN parameter matrices, and let ep_n denote the number of LSTM-NN training epochs. We train eight LSTM-NNs on the window-specific context vectors constructed with hyperparameter settings $(7,300)$ using all eight possible combinations of $\lambda \in \{0.005, 0.01\}$, $u \in \{1,2\}$, and $ep_n \in \{2000, 3000\}$. For each $\tau \in \mathscr{T}$, we define the LSTM-NN that results in the lowest median testing loss to be the "best" LSTM-NN at time τ. The hyperparameter settings of the best LSTM-NN at each $\tau \in \mathscr{T}$ can be seen in Table (2.1).

We train all LSTM autoencoders and LSTM-NNs using the Adam optimization algorithm (Kingma and Ba (2017)). For each LSTM autoencoder, we use Bayesian optimization to adaptively select the learning rate from the options $\{1e^{-2}, 1e^{-3}, 1e^{-4}\}$ (O'Malley et al. (2019)). For the LSTM-NN, we select a fixed learning rate of $1e^{-4}$.

We conduct prediction at five prediction times, $\mathscr{T} = \{1, 1.5, 2, 2.5, 3\}$, where $\tau \in \mathscr{T}$ represents the number of days passed since the time of the given patient's first record in the data set. At prediction time $\tau \in \mathscr{T}$, we randomly divide the data into a training data set and a testing data set, stratifying on the censoring status Δ_i. The training data set is specified to include 70 percent of the patients at risk at time τ, and the testing data set is defined to include the other 30 percent of patients at risk. We estimate the survival function of censoring time $\widehat{G}(\cdot)$ independently on the training and testing data sets. We then fit each model on the training data set and compute the performance metrics specified in Section 2.3.2 on the testing data set. We repeat this process 100 times, using 100 unique divisions of the data.

2.5.2 Results

For each prediction time $\tau \in \mathscr{T}$, we plot the distribution of 100 testing losses for the best LSTM-GLM, the automated LSTM-GLM, the best LSTM-NN, and the six dynamic transformed MRL models in Figure (2.7), and we plot the distribution of 100 testing C-indexes in Figure (2.8).

First, we compare the calibration of the nine dynamic prediction models via the testing loss. Generally, the best LSTM-GLM, the automated LSTM-GLM, the best LSTM-NN, and the FPCA model result in the lowest median testing losses. The baseline model consistently results in the highest median testing loss.

At prediction times $\tau = 1, 1.5, 2.5, 3$, the best LSTM-GLM results in the lowest median testing loss. At $\tau = 2$, the best LSTM-NN results in the lowest median testing loss. In this application, the added flexibility of the LSTM-NN does not offset the cost of imperfect hyperparameter tuning.

In practice, we do not know which hyperparameter settings result in the "best" LSTM-GLM and the "best" LSTM-NN. Consequently, we use cross-validation to automatically select hyperparameters for the LSTM-GLM, as detailed in Appendix A, Section A.1. Generally, the automated hyperparameter selection process performs well. Disregarding the best LSTM-GLM and the best LSTM-NN, the automated LSTM-GLM results in the lowest median testing loss at prediction times $\tau = 1, 2.5, 3$ and in the second-lowest median testing loss at $\tau = 1.5, 2$, where it is beat only by the FPCA model.

Second, we compare the discrimination of the nine dynamic prediction models via the testing C-index. As with calibration, the best LSTM-GLM, the automated LSTM-GLM, the best LSTM-NN, and the FPCA model generally result in good discrimination, and the baseline model consistently results in the poorest discrimination. Interestingly, the last-value carried forward model results in the highest median testing C-index at prediction times $\tau = 1.5, 2$.

The best LSTM-NN results in a higher testing C-index than the best LSTM-GLM only at prediction times $\tau = 1.5, 3$. Again, this indicates that the added flexibility of the LSTM-NN is offset by imperfect hyperparameter tuning.

The automated hyperparameter selection process for the LSTM-GLM performs well with respect to discrimination. Disregarding the best LSTM-GLM and the best LSTM-NN, the automated LSTM-GLM results in the highest median testing C-index at prediction times $\tau = 1, 3$, the second-highest at $\tau = 2$, and the third-highest at $\tau = 1.5, 2.5$. In this case, the

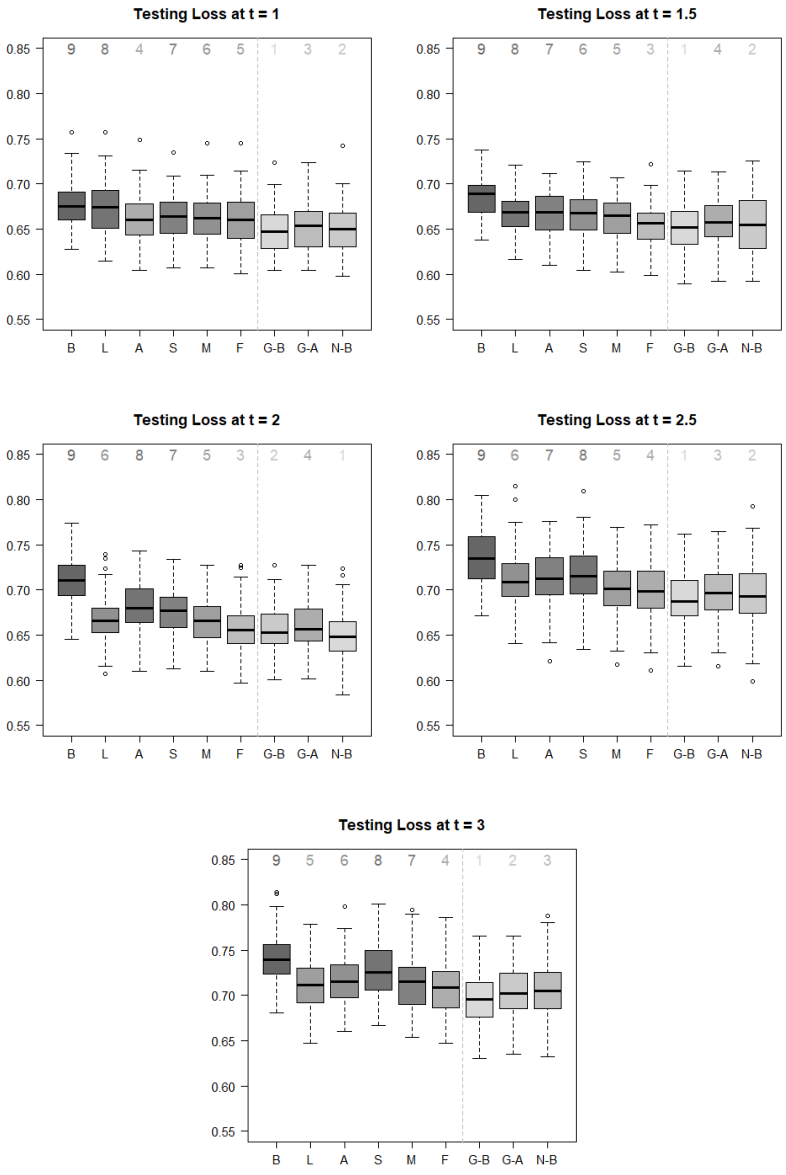

Figure 2.7: Distribution of the 100 testing losses for each of the nine studied dynamic prediction models at each prediction time $\tau \in \mathcal{T} = \{1, 1.5, 2, 2.5, 3\}$. The model resulting in the lowest median testing loss is labelled 1. The model resulting in the highest median testing loss is labelled 9. The six dynamic transformed MRL models are labelled according to their formulation of $\boldsymbol{\zeta}_i(\tau)$. "B" represents the baseline vector. "L" represents the last-value carried forward vector. "A" represents the average vector. "S" represents the linear regression vector. "M" represents the mixed effects vector. "F" represents the FPCA vector. The best LSTM-GLM is labelled "G-B." The automated LSTM-GLM is labelled "G-A." The best LSTM-NN is labelled "N-B."

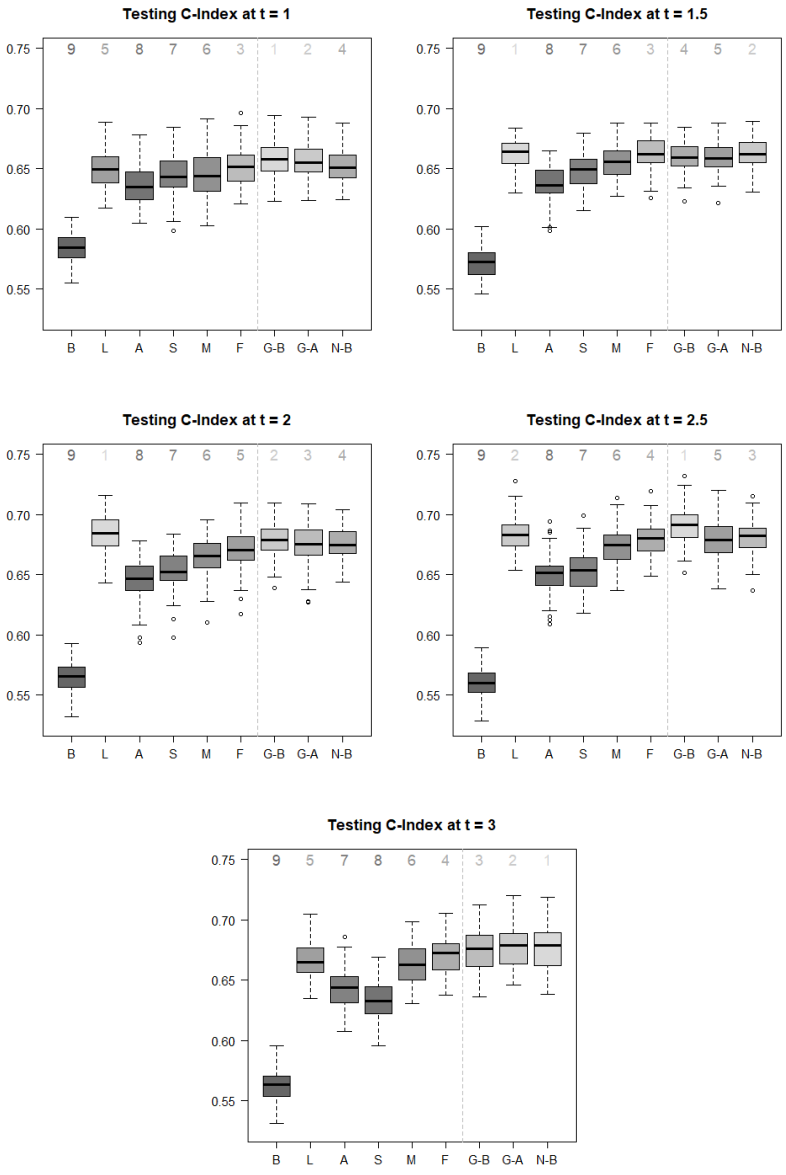

Figure 2.8: Distribution of the 100 testing C-indexes for each of the nine studied dynamic prediction models at each prediction time $\tau \in \mathscr{T} = \{1, 1.5, 2, 2.5, 3\}$. The model resulting in the highest median testing C-index is labelled 1. The model resulting in the lowest median testing C-index is labelled 9. The six dynamic transformed MRL models are labelled according to their formulation of $\boldsymbol{\zeta}_i(\tau)$. "B" represents the baseline vector. "L" represents the last-value carried forward vector. "A" represents the average vector. "S" represents the linear regression vector. "M" represents the mixed effects vector. "F" represents the FPCA vector. The best LSTM-GLM is labelled "G-B." The automated LSTM-GLM is labelled "G-A." The best LSTM-NN is labelled "N-B."

discrimination of the automated LSTM-GLM is beat only by the last-value carried forward model and the FPCA model.

In clinical application, calibration performance is often valued over discrimination performance. Typically, it is more important to accurately predict one septic patient's remaining life expectancy than to accurately predict which of two septic patients has a longer remaining life expectancy. Taking this into consideration, these results suggest that the LSTM-GLM and the LSTM-NN exhibit the best performance when dynamically predicting the RMRL of septic patients in the ICU from EMR data, as compared to the baseline, last-value carried forward, average, linear regression, and mixed effects models. Moreover, the LSTM-GLM and the LSTM-NN exhibit at least comparable performance to the FPCA model, and if their hyperparameters are properly tuned, the LSTM-GLM and the LSTM-NN exhibit better performance than the FPCA model. This study demonstrates that the LSTM-GLM and the LSTM-NN are useful methods for producing individualized, real-time predictions of the RMRL of septic patients in the ICU from EMR data.

2.6 Discussion

Sepsis is a leading cause of death worldwide and remains highly challenging to treat (Singer et al. (2016)). We introduce two methods, the LSTM-GLM and the LSTM-NN, to dynamically predict the RMRL of septic patients in the ICU from EMR data. Through simulation studies and application to the MIMIC-III data set, we demonstrate that the LSTM-GLM and the LSTM-NN exhibit superior prediction performance relative to six competing methods. Thus, the LSTM-GLM and the LSTM-NN offer an automatic method to synthesize complex longitudinal biomarker trajectories and produce accurate predictions of MRL, all while upholding the dynamic nature of prediction. By producing individualized, real-time predictions of the RMRL of septic patients, the LSTM-GLM and the LSTM-NN can help clinicians make informed treatment decisions, potentially improving septic patient care.

The LSTM-GLM and the LSTM-NN can process data containing a large number of patients and biomarkers, like that typically found in EMRs. However, the methods can be computationally expensive relative to the baseline, last-value carried forward, average, linear regression, mixed effects, and FPCA dynamic transformed MRL models. Because we propose training a distinct LSTM autoencoder to construct the context vector for each longitudinal biomarker, the relative computational cost of the LSTM-GLM and the LSTM-NN

is most notable in settings with a large number of longitudinal biomarkers. To improve the computational efficiency of the proposed methods, a context vector summarizing multiple biomarkers could be constructed by training a single LSTM autoencoder on multiple biomarkers. However, this approach further obfuscates the relationship between MRL and the biomarkers. Additional research needs to be conducted to evaluate the feasibility and utility of this approach.

Alternative to the presented methods, a joint model could be constructed to connect the longitudinal biomarkers with the hazard function of death (Tsiatis and Davidian (2004)), and MRL predictions could be derived from the hazard function. However, this procedure can make it difficult to interpret the relationship between the biomarkers and MRL. Additionally, constructing joint models with a large number of longitudinal biomarkers can be computationally challenging (Hickey et al. (2016)).

In this work, we assume censoring time is independent of the baseline covariates and longitudinal biomarkers, and we estimate the survival function of censoring time $\widehat{G}(\cdot)$ via the Kaplan-Meier method. If censoring time is dependent on both baseline and longitudinal covariates, a time-dependent Cox regression model can be used to estimate $\widehat{G}(\cdot)$. However, further research should be conducted to determine how best to incorporate the history of longitudinal biomarker measurements into the Cox regression model. Potentially, the window-specific context vectors could be used as predictors in the Cox model.

We present an automated hyperparameter selection process for the LSTM-GLM in Appendix A, Section A.1. Further research is required to formulate satisfactory methods for tuning the hyperparameters of the LSTM-NN.

In this chapter, we focus on obtaining accurate predictions of MRL. Consequently, we impose an L_1-penalty on the parameters of the LSTM-GLM to prevent overfitting. To conduct variable selection with the LSTM-GLM, a group LASSO penalty could instead be imposed to ensure that the entire context vector for a given longitudinal biomarker is either included in or removed from the model (Yuan and Lin (2006)). Further research is needed to determine the efficacy of using the LSTM-GLM to conduct variable selection.

CHAPTER

3

ESTIMATION OF OPTIMAL TREATMENT REGIMES WITH ELECTRONIC MEDICAL RECORD DATA USING THE RESIDUAL LIFE VALUE ESTIMATOR (RELIVE)

3.1 Introduction

Throughout disease progression, clinicians and patients must make treatment decisions at a series of key decision points. A dynamic treatment regime is a set of sequential decision rules that provides a formal process for making treatment decisions based on accumulating patient information. At each decision point, the decision rule accepts patient history as input and returns a recommended treatment from among the feasible options. When used to select treatments for a patient population, an optimal treatment regime leads to the most favorable outcome on average. Identifying an optimal regime is an integral component of

precision medicine, an approach to healthcare that focuses on tailoring treatment decisions to patient characteristics.

Across the fields of statistics and computer science, there exists a vast literature on methods for characterizing and estimating optimal treatment regimes, many of which developed from the seminal works of Murphy (2003) and Robins (2004). One class of methods seeks to directly optimize an estimator of the population mean outcome that would potentially be achieved under a given treatment regime, where optimization is conducted over all treatment regimes in a restricted class. Zhang et al. (2013) proposed direct optimization of a doubly robust inverse probability weighted estimator for the mean, while Orellana et al. (2010) proposed direct optimization of marginal structural mean models. Zhao et al. (2015a) introduced two outcome weighted learning methods, which reframe optimal regime estimation as a sequence of weighted classification problems.

In practice, the method of Q-learning (Murphy (2005)) is a popular approach for estimating an optimal treatment regime. Q-learning is a backward induction algorithm (Bellman (1957)) that first identifies an optimal decision rule at the final decision point, then sequentially identifies an optimal decision rule at each previous decision point. A model, referred to as the "Q-model," is posited for the expected outcome given the patient history at each decision point. At each decision point, the Q-model is optimized with respect to treatment, with all future treatment decisions taken to be optimal.

Identifying an optimal treatment regime that maximizes expected residual life is of particular interest for patients with potentially life-threatening conditions, such as sepsis. Sepsis is a complex medical condition that involves severe infections with life-threatening organ dysfunction and is a leading cause of death worldwide (Singer et al. (2016)). Although international guidelines for sepsis treatment have been established, treating septic patients remains highly challenging (Evans et al. (2021)). The heterogeneity of septic patient populations results in differing responses to medical intervention (László et al. (2015)). Consequently, determining optimal decision rules for selecting treatment based on patients' individualized characteristics is highly desirable for septic patient care.

We are especially interested in identifying optimal treatment regimes for septic patients in the intensive care unit (ICU) using their electronic medical record (EMR) data. In particular, we study a data set of septic patients constructed from the Multiparameter Intelligent Monitoring Intensive Care database (MIMIC-III), described in Chapter 1. Given a number of patient characteristics, including admission records, longitudinal vital signs, and longitudinal laboratory measurements, we seek to estimate optimal treatment regimes that

maximize the expected value of restricted residual life for the septic patients in MIMIC-III.

The dynamic treatment regime literature has been extended to accommodate censored time-to-event outcomes, such as residual life. A number of methods for direct optimization of an expected outcome estimator, including a number of classification-based methods, have been proposed for censored time-to-event outcomes (Zhao et al. (2015b); Jiang et al. (2017); Bai et al. (2017); Hager et al. (2018)). Additionally, Goldberg and Kosorok (2012) extended Q-learning to accommodate censored survival time data. See Chapter 8 of Tsiatis et al. (2020) for an overview of methods for estimating optimal regimes based on time-to-event outcomes.

Building on this work, we introduce the residual life value estimator (ReLiVE), an estimator for the expected value of the sum of restricted residual life under a fixed treatment regime, where the sum is taken across all decision points at which the patient is alive. Based on ReLiVE, we then present the residual life value estimator Q-learning method (ReLiVE-Q). ReLiVE-Q uses Q-learning to estimate an optimal treatment regime that maximizes the expected value of restricted residual life across all regimes.

We adopt a Q-learning approach because Q-learning naturally accommodates several challenges frequently incurred with EMR data like the studied MIMIC-III data set. First, EMR data often includes a large number of covariates, so directly optimizing an estimator of the population mean outcome over a restricted class of regimes may be infeasible. In contrast, Q-learning can be used to conduct estimation in settings with numerous covariates. Second, Q-learning allows us to specify flexible, non-parametric Q-models. This is critical given that the inherently complex relationship between patient history, treatment, and restricted residual life is unlikely to be well-represented by a linear model. At each decision point, we fit the Q-models in ReLiVE-Q using the random forest algorithm, an ensemble learning algorithm that combines the output of numerous decision trees (Breiman (2001)). Third, with Q-learning, we can incorporate the trajectories of longitudinal covariates into the Q-models using modern machine learning techniques. This is important because the trajectories of the longitudinal covariates in EMR data are often sophisticated functions of time, and the timing and number of measurements often differ between patients. Thus, simple summary statistics may not summarize the longitudinal covariates well. Consequently, we propose incorporating the longitudinal trajectories in the ReLiVE-Q Q-models using context vectors, as described in Chapter 2.

In Section 3.2, we present the statistical framework for ReLiVE and ReLiVE-Q, and in Section 3.3, we detail the ReLiVE and ReLiVE-Q methodology. In Section 3.4, we demon-

strate the utility of ReLiVE-Q in a simulation study, and in Section 3.5, we apply ReLiVE-Q to estimate optimal treatment regimes for septic patients in MIMIC-III. In both the simulation study and MIMIC-III data application, we demonstrate that ReLiVE-Q estimates personalized treatment regimes that optimize a clinically meaningful function of expected residual life. We conclude with a discussion of implications and open problems in Section 3.6.

3.2 Statistical Framework

3.2.1 Potential Outcomes & Treatment Regimes

We focus our study on settings where treatment decisions are made at fixed intervals, as is the scenario for the studied MIMIC-III data set. Consider a series of K decision points at which treatment decisions must be made, where the decision points occur at fixed times $\tau_1,\ldots,\tau_K$ with $\tau_1 = 0$. Let $\mathbf{x}_1$ denote the vector of covariates measured at τ_1. Then the patient history at decision point 1 is $\mathbf{h}_1 = (\tau_1, \mathbf{x}_1)$. For $k = 1,\ldots,K$, define $\mathscr{A}_k$ to be the finite set of all available treatment options at decision point k, and let $a_k \in \mathscr{A}_k$ be the treatment administered at decision point k. We assume all $a_k \in \mathscr{A}_k$ are feasible for all subjects at all decision points $k = 1,\ldots,K$, as this assumption holds for the treatments studied in MIMIC-III. For $k = 2,\ldots,K$, let $\mathbf{x}_k$ denote the vector of covariate information accumulated between decision points $k-1$ and k. Moreover, let $\overline{\tau}_k = (\tau_1,\ldots,\tau_k)$, $\overline{\mathbf{x}}_k = (\mathbf{x}_1,\ldots,\mathbf{x}_k)$, and $\overline{a}_k = (a_1,\ldots,a_k)$, where $k = 1,\ldots,K$, and let $\overline{a} = \overline{a}_K$. Then for $k = 2,\ldots,K$, a patient who does not experience the event of interest prior to decision point k will have patient history $\mathbf{h}_k = (\overline{\tau}_k, \overline{\mathbf{x}}_k, \overline{a}_{k-1})$ at decision point k. In contrast, a patient who experiences the event of interest at time t between decision points $k-1$ and k will have patient history $\mathbf{h}_j = (\overline{\tau}_{k-1}, \overline{\mathbf{x}}_{k-1}, \overline{a}_{k-1}, t)$ at each decision point $j = k,\ldots,K$.

For each decision point, $k = 1,\ldots,K$, let $d_k(\mathbf{h}_k)$ be a decision rule. If $\mathbf{h}_k$ indicates that the event of interest has not occurred prior to decision point k, then d_k maps $\mathbf{h}_k$ to $\mathscr{A}_k$. Else, for $k = 2,\ldots,K$, if $\mathbf{h}_k$ indicates that the event of interest has occurred prior to decision point k, then $d_k(\mathbf{h}_k)$ returns null. For notational simplicity, suppress the dependence of the decision rule on patient history, and let $d_k = d_k(\mathbf{h}_k)$. Define treatment regime $d = (d_1,\ldots,d_K)$, and denote the set of all such possible treatment regimes as $\mathscr{D}$. For convenience, let $\overline{d}_k = (d_1,\ldots,d_k)$, where $k = 1,\ldots,K$.

To formalize the definition of expected restricted residual life that would occur if an

individual were to follow a given regime $d \in \mathscr{D}$, we define the potential outcomes that would be achieved if a randomly selected patient were treated according to d. To this end, we first define the potential outcomes associated with a given sequence of treatments $\overline{a}$. Let $\kappa^*(\overline{a})$ be the potential number of decision points a randomly selected patient would reach if administered treatments $\overline{a}$. For $k = 2, \ldots, \kappa^*(\overline{a})$, let $\mathbf{X}_k^*(\overline{a}_{k-1})$ be the vector of potential covariate information that would occur between decision points $k-1$ and k if a patient was administered treatments $\overline{a}_{k-1}$. Further, let $T^*(\overline{a}_{\kappa^*(\overline{a})})$ be the potential time-to-event that would be observed if a patient was administered treatments $\overline{a}_{\kappa^*(\overline{a})}$. Define the set of all possible potential outcomes to be $W^* = \{$For all possible $\overline{a} \in \mathscr{A}_1 \times \cdots \times \mathscr{A}_K$: $\kappa^*(\overline{a}), \mathbf{X}_2^*(a_1), \mathbf{X}_3^*(\overline{a}_2), \ldots, \mathbf{X}_{\kappa^*(\overline{a})}^*(\overline{a}_{\kappa^*(\overline{a})-1}), T^*(\overline{a}_{\kappa^*(\overline{a})})\}$.

As discussed in Sections 6.2.3 and 8.3.2 of Tsiatis et al. (2020), it can be shown that the potential outcomes associated with a given regime $d \in \mathscr{D}$ can be defined in terms of W^*. Let $T^*(d)$ be the potential time-to-event under regime d. To avoid infinite event times, also specify a restricted event time $L > 0$. Then $\kappa_L^*(d) = \max[k : \tau_k < \min\{T^*(d), L\}]$ is the potential number of decision points reached under regime d. For $k = 2, \ldots, \kappa_L^*(d)$, let $\mathbf{X}_k^*(\overline{d}_{k-1})$ be the vector of potential covariate information that would occur between decision points $k-1$ and k if a randomly selected patient was treated according to regime $\overline{d}_{k-1}$. Denote the set of potential outcomes associated with regime d as $W_d^* = \{\kappa_L^*(d), \mathbf{X}_2^*(d_1), \mathbf{X}_3^*(\overline{d}_2), \ldots, \mathbf{X}_{\kappa_L^*(d)}^*(\overline{d}_{\kappa_L^*(d)-1}), T^*(d)\}$.

Because sepsis is a life-threatening medical condition, we specify $T^*(d)$ to be the potential time of death under regime d. Then, restricted residual life under regime d at a given time τ is defined as $\min\{T^*(d), L\} - \tau$. We define the value of regime $d \in \mathscr{D}$, $\mathscr{V}(d)$, to be the expected value of the sum of a patient's restricted residual life under regime d, where the sum is taken across all decision points at which the patient is alive. That is, we define

$$\mathscr{V}(d) = E\left\{\sum_{j=1}^{K} I[\min\{T^*(d), L\} > \tau_j] \cdot \left[\min\{T^*(d), L\} - \tau_j\right]\right\}. \tag{3.1}$$

An optimal treatment regime $d^{opt} \in \mathscr{D}$ satisfies the condition that $\mathscr{V}(d^{opt}) \geq \mathscr{V}(d)$ for all $d \in \mathscr{D}$. Based on observed EMR data, we first aim to estimate $\mathscr{V}(d)$ for any fixed regime $d \in \mathscr{D}$. We then aim to estimate an optimal regime d^{opt} and its value $\mathscr{V}(d^{opt})$ from the observed data.

3.2.2 Data & Assumptions

We now describe the observed data. As is conventional in the survival analysis literature, let $T > 0$ and $C > 0$ denote the potential times to death and censoring, respectively. We observe only $U = \min(T, C)$ and $\Delta = I(T < C)$, an indicator of whether death is observed ($\Delta = 1$) or censored ($\Delta = 0$). Accounting for the restricted lifetime $L > 0$, we observe the restricted outcome $U^L = \min(U, L)$ and associated indicator $\Delta^L = I\{\min(T, L) < C\} = \Delta + I(L < U)(1-\Delta)$. A patient is observed to reach $\kappa_L = \max\{k : \tau_k < U^L\}$ decision points, where $\kappa_L \leq K$. Let $\mathbf{X}_1$ denote the observed baseline information, A_k denote the treatment option administered at decision point k, $k = 1, \ldots, \kappa_L$, and $\mathbf{X}_k$ denote the observed covariate information accumulated between decision points $k-1$ and k, $k = 2, \ldots, \kappa_L$. Then the observed history at decision points $k = 1, \ldots, \kappa_L$ is $\mathbf{H}_k = (\overline{\tau}_k, \overline{\mathbf{X}}_k, \overline{A}_{k-1})$, and the observed history at decision points $k = \kappa_L + 1, \ldots, K$ is $\mathbf{H}_k = (\overline{\tau}_{\kappa_L}, \overline{\mathbf{X}}_{\kappa_L}, \overline{A}_{\kappa_L}, U^L, \Delta^L)$. Given m patients, the observed data are independent and identically distributed $(\kappa_{Li}, \mathbf{X}_{1i}, A_{1i}, \ldots, \mathbf{X}_{\kappa_{Li}i}, A_{\kappa_{Li}i}, U_i^L, \Delta_i^L)$, $i = 1, \ldots, m$.

We aim to estimate value (3.1) for a given regime $d \in \mathscr{D}$, and to estimate an optimal regime d^{opt} and its value, based on the observed data. Because $T^*(d)$ is defined in terms of W^*, we must be able to express relevant functionals of the distribution of the potential outcomes W^* in terms of the observed data. Such expression is possible under three key, standard assumptions: the Stable Unit Treatment Value Assumption (SUTVA), the Sequential Randomization Assumption (SRA), and the positivity assumption. SUTVA, commonly referred to as the consistency assumption, states that the observed data are the same as those that would potentially be achieved under the treatment decisions observed to be administered. Formally, we assume $\kappa_L = \kappa_L^*(\overline{A}_{\kappa_L})$ and $T = T^*(\overline{A}_{\kappa_L})$. We also assume $\mathbf{X}_k = \mathbf{X}_k^*(\overline{A}_{k-1})$ for $k = 2, \ldots, \kappa_L$. SRA states that the treatment administered at each decision point is independent of the set of potential outcomes conditional on the history. For $k = 1, \ldots, K$, we assume $W^* \perp A_k | \mathbf{H}_k, \kappa_L \geq k$, where $\perp$ denotes independence. The positivity assumption states that all treatment options at each decision point are represented in the data. This is necessary to specify the distributions of the potential outcomes for all $d \in \mathscr{D}$ in terms of the observed data. Formally, we assume $P(A_k = a_k \mid \mathbf{H}_k = \mathbf{h}_k, \kappa_L \geq k) > 0$ for all $a_k \in \mathscr{A}_k$ and for all possible $\mathbf{h}_k$ such that $P(\mathbf{H}_k = \mathbf{h}_k, \kappa_L \geq k) > 0$, $k = 1, \ldots, K$. For simplicity, we also assume censoring is independent of treatment assignment, patient characteristics, restricted lifetime, and potential outcomes. See Section 3.6 for a discussion on relaxing this assumption.

3.3 Methods

3.3.1 Methodological Foundation of ReLiVE

We first introduce the methodological foundation for ReLiVE, our proposed estimator for value (3.1) of a fixed regime. For a given regime $d \in \mathscr{D}$, we construct an estimator for $\mathscr{V}(d)$ by defining K backward recursive Q-functions. We incorporate inverse-probability weights in the Q-functions to account for censoring. For $x > \tau_k$, let $\mathscr{K}_k(x|\mathbf{h}_k, a_k) = P(C > x \mid U^L > \tau_k, \mathbf{H}_k = \mathbf{h}_k, A_k = a_k)$, where $k = 1, \ldots, K$. Under the assumptions specified in Section 3.2.2, $\mathscr{K}_k(x|\mathbf{h}_k, a_k) = P(C > x)/P(C > \tau_k)$. See Appendix B, Section B.1 for details.

For a fixed regime $d \in \mathscr{D}$, we define the Q-function at decision point K to be

$$Q_K^d(\mathbf{h}_K, a_K) = E\left\{ \frac{\Delta^L(U^L - \tau_K)}{\mathscr{K}_K(U^L|\mathbf{h}_K, a_K)} \middle| U^L > \tau_K, \mathbf{H}_K = \mathbf{h}_K, A_K = a_K \right\}.$$

Similarly, we define the Q-function at decision points $k = K-1, \ldots, 1$, to be

$$Q_k^d(\mathbf{h}_k, a_k) = E\left\{ \frac{\Delta^L(U^L - \tau_k)}{\mathscr{K}_k(U^L|\mathbf{h}_k, a_k)} + \frac{I(U^L > \tau_{k+1}) V_{k+1}^d(\mathbf{H}_{k+1})}{\mathscr{K}_k(\tau_{k+1}|\mathbf{h}_k, a_k)} \middle| U^L > \tau_k, \mathbf{H}_k = \mathbf{h}_k, A_k = a_k \right\},$$

where $V_k^d(\mathbf{h}_k) = Q_k^d\{\mathbf{h}_k, d_k(\mathbf{h}_k)\}$. Under the assumptions specified in Section 3.2.2, $V_1^d(\mathbf{h}_1) = Q_1^d\{\mathbf{h}_1, d_1(\mathbf{h}_1)\}$ is an unbiased estimator of $\mathscr{V}(d)$ if the true Q-functions are known. In Appendix B, Section B.2, we provide a proof demonstrating that $E\{V_1^d(\mathbf{h}_1)\} = \mathscr{V}(d)$.

In practice, the Q-functions for regime d are unknown and must be estimated from the data. In Section 3.3.5, we present the value estimator ReLiVE, which uses a flexible approach to model and estimate the Q-functions.

3.3.2 Methodological Foundation of ReLiVE-Q

Next, we introduce the foundation for ReLiVE-Q, our proposed Q-learning method for estimating an optimal treatment regime and its value. Building on the Q-functions presented in Section 3.3.1, we characterize an optimal treatment regime $d^{opt} \in \mathscr{D}$ such that $\mathscr{V}(d^{opt}) \geq \mathscr{V}(d)$ for all $d \in \mathscr{D}$. At decision point K, define the Q-function for d^{opt} to be

$$Q_K^{d^{opt}}(\mathbf{h}_K, a_K) = E\left\{ \frac{\Delta^L(U^L - \tau_K)}{\mathscr{K}_K(U^L|\mathbf{h}_K, a_K)} \middle| U^L > \tau_K, \mathbf{H}_K = \mathbf{h}_K, A_K = a_K \right\},$$

and let $d_K^{opt}(\mathbf{h}_K) = \underset{a_K \in \mathscr{A}_K}{\text{argmax}}\, Q_K^{d^{opt}}(\mathbf{h}_K, a_K)$. Similarly, at decision points $k = K-1, \ldots, 1$, define

$$Q_k^{d^{opt}}(\mathbf{h}_k, a_k) = E\left\{ \frac{\Delta^L(U^L - \tau_k)}{\mathscr{K}_k(U^L|\mathbf{h}_k, a_k)} + \frac{I(U^L > \tau_{k+1}) V_{k+1}^{d^{opt}}(\mathbf{H}_{k+1})}{\mathscr{K}_k(\tau_{k+1}|\mathbf{h}_k, a_k)} \middle| U^L > \tau_k, \mathbf{H}_k = \mathbf{h}_k, A_k = a_k \right\},$$

where $V_k^{d^{opt}}(\mathbf{h}_k) = Q_k^{d^{opt}}\{\mathbf{h}_k, d_k^{opt}(\mathbf{h}_k)\}$ and $d_k^{opt}(\mathbf{h}_k) = \underset{a_k \in \mathscr{A}_k}{\text{argmax}}\, Q_k^{d^{opt}}(\mathbf{h}_k, a_k)$. Then it can be shown that $d^{opt} = (d_1^{opt}, \ldots, d_K^{opt})$ is an optimal treatment regime satisfying $\mathscr{V}(d^{opt}) \geq \mathscr{V}(d)$ for all $d \in \mathscr{D}$, and the value of an optimal regime is $\mathscr{V}(d^{opt}) = E\{V_1^{d^{opt}}(\mathbf{h}_1)\}$. See Sections 7.2.3 - 7.2.4 of Tsiatis et al. (2020) for details.

In practice, the Q-functions for d^{opt} are unknown and must be estimated from the data. In Section 3.3.6, we present ReLiVE-Q, a Q-learning method that uses a flexible approach to estimate the Q-functions for d^{opt}.

3.3.3 Representation of Longitudinal Covariates

In practice, certain covariates, such as sex and birth date, are collected on patients only at baseline. Denote the vector of covariates measured only at baseline as $\mathbf{S}$. Other covariates, such as vital signs and laboratory values, may be measured repeatedly during a patient's follow up. Let $\mathscr{M}$ be the set of measurement times, and denote the vector of longitudinal covariates measured at time $t \in \mathscr{M}$ as $\mathbf{Z}(t) = \{Z_1(t), \ldots, Z_p(t)\}$. For longitudinal covariates measured at time $t = \tau_k$, $k = 1, \ldots, \kappa_L$, we follow the convention that longitudinal covariates are measured immediately prior to making treatment decisions. Then $\mathbf{X}_1 = \{\mathbf{S}, \mathbf{Z}(0)\}$, and $\mathbf{X}_k = \{\mathbf{Z}(t) : \tau_{k-1} < t \leq \tau_k\}$ for $k = 2, \ldots, \kappa_L$.

When estimating the Q-functions for ReLiVE and ReLiVE-Q, it is desirable to leverage all available patient information. To incorporate the trajectories of the longitudinal covariates into the Q-models specified in Sections 3.3.5 and 3.3.6, we must provide the Q-models with summaries of the longitudinal covariates. Let $f(\cdot)$ be a vector-valued function, and define vector $\boldsymbol{\zeta}(k) = f[\{\mathbf{Z}(t) : t \leq \tau_k\}]$ to be a function of the longitudinal covariate measurements accumulated by decision point k. Then $\overline{\mathbf{X}}_1 = \{\mathbf{S}, \mathbf{Z}(0)\}$, and for $k = 2, \ldots, \kappa_L$, we define $\overline{\mathbf{X}}_k = \{\mathbf{S}, \boldsymbol{\zeta}(k)\}$. The observed patient history is then given by $\mathbf{H}_1 = (\tau_1, \mathbf{X}_1) = \{\tau_1, \mathbf{S}, \mathbf{Z}(0)\}$ at decision point 1, and $\mathbf{H}_k = (\overline{\tau}_k, \overline{\mathbf{X}}_k, \overline{A}_{k-1}) = \{\overline{\tau}_k, \mathbf{S}, \boldsymbol{\zeta}(k), \overline{A}_{k-1}\}$ at decision points $k = 2, \ldots, \kappa_L$. At decision points $k = 1, \ldots, K$, we train the Q-models for ReLiVE and ReLiVE-Q over the predictor space $\overline{\mathbf{X}}_k$, as described in Section 3.3.4.

Intuitively, it is desirable to select a function $f(\cdot)$ that summarizes the longitudinal trajectories from baseline to decision point time. Several simple summary functions are commonly used in practice. The baseline vector $\boldsymbol{\zeta}^B(k) = \mathbf{Z}(0)$ involves only the baseline measures of the longitudinal covariates, while the last-value carried forward vector $\boldsymbol{\zeta}^L(k) = \mathbf{Z}(t_k)$, where $t_k = \max\{t \in \mathscr{M} : t \leq \tau_k\}$, summarizes each longitudinal trajectory through only the most recently observed value. The average vector $\boldsymbol{\zeta}^A(k) = \{Z_1^A(\tau_k), \ldots, Z_p^A(\tau_k)\}$, where $Z_l^A(\tau_k) = \left\{\sum_{t\in\mathscr{M}} I(t \leq \tau_k)\right\}^{-1} \sum_{t\in\mathscr{M}} I(t \leq \tau_k) Z_l(t)$, $l = 1, \ldots, p$, summarizes each trajectory using the average of the observed longitudinal measurements through decision point k.

The trajectories of two longitudinal covariates, magnesium and systolic blood pressure, are depicted in Figure (1.1) for 25 randomly selected patients from the MIMIC-III data set. Figure (1.1) illustrates that the aforementioned summary functions are unlikely to capture the complex nature of the longitudinal trajectories, as the trajectories are sophisticated functions of time that exhibit notable variation among patients. Consequently, we propose synthesizing the longitudinal covariates using window-specific context vectors.

The window-specific context vector $\boldsymbol{\psi}_l(\tau)$ is an encoded representation of the trajectory of longitudinal covariate $Z_l(\cdot)$ from baseline to time τ, where $l = 1, \ldots, p$. At decision point times $\tau_2, \ldots, \tau_{\kappa_L}$, we use a distinct long short-term memory (LSTM) autoencoder to construct each of the p window-specific context vectors for the p longitudinal covariates. An LSTM autoencoder is an unsupervised neural network that learns how to best encode temporal input into a context vector, so it can then reconstruct the original input from that context vector. The LSTM autoencoder used to construct $\boldsymbol{\psi}_l(\tau)$ only accepts measurements collected by time τ on patients with $U^L > \tau$, making context vectors ideal for dynamic estimation. For a detailed review of window-specific context vectors, refer to Chapter 2. At decision points $k = 2, \ldots, \kappa_L$, we define the context vector $\boldsymbol{\zeta}^C(k) = \{\boldsymbol{\psi}_1(\tau_k), \ldots, \boldsymbol{\psi}_p(\tau_k)\}$. Because only baseline measurements are collected by $\tau_1 = 0$, we define $\boldsymbol{\zeta}^C(1) = \mathbf{Z}(0)$ for decision point 1.

3.3.4 Q-Function Estimation Strategy

We propose methods to estimate the Q-functions for ReLiVE and ReLiVE-Q in Sections 3.3.5 and 3.3.6, respectively. In both sections, we model the Q-functions using a non-parametric approach to account for the inherently complex relationship between patient history and functions of residual life. Although extension to more complex treatment settings is possible,

we focus on settings with two treatment options such that $\mathscr{A}_k = \{0,1\}$ for $k = 1,\ldots,K$. Under this condition, our general strategy exploits the fact that the Q-function for regime $d \in \mathscr{D}$ can be written at each decision point, $k = 1,\ldots,K$, as

$$Q_k^d(\mathbf{h}_k, a_k) = \mu_{0k}^d(\mathbf{h}_k) + a_k\{\mu_{1k}^d(\mathbf{h}_k) - \mu_{0k}^d(\mathbf{h}_k)\},$$

where $\mu_{0k}^d(\mathbf{h}_k) = Q_k^d(\mathbf{h}_k, 0)$ and $\mu_{1k}^d(\mathbf{h}_k) = Q_k^d(\mathbf{h}_k, 1)$.

For a given regime $d \in \mathscr{D}$, we estimate $\mu_{0k}^d(\mathbf{h}_k)$ and $\mu_{1k}^d(\mathbf{h}_k)$ using the random forest algorithm, an ensemble learning algorithm that improves prediction by combining the output of numerous decision trees. A decision tree is a non-parametric prediction method that recursively partitions the predictor space such that observations in each partitioned subset are successively more homogeneous (Breiman (1984)). A decision tree predicts the outcome of a given observation to be the mean outcome of all observations in the partitioned subset to which that observation belongs. The random forest algorithm constructs a forest of unique, decorrelated decision trees using bootstrap aggregation (or bagging) and feature randomization. Bagging reduces the variance of the forest by creating a unique, bootstrapped sample to train each decision tree. Feature randomization decorrelates the decision trees by randomly selecting a subset of the predictors at each step in the recursive algorithm, so only predictors in the subset are considered as candidates for partitioning. In a random forest, the predicted outcome for a given observation is taken to be the mean of all predicted outcomes for that observation across the trees in the forest. For a detailed review of the random forest algorithm, see Breiman (2001).

In our proposed implementations of ReLiVE and ReLiVE-Q, we estimate $\mu_{0k}^d(\mathbf{h}_k)$ and $\mu_{1k}^d(\mathbf{h}_k)$ at each decision point using distinct random forests, where the outcomes used to train the forests are presented in the methods of Sections 3.3.5 and 3.3.6. At decision point 1, we train the random forests for $\mu_{0k}^d(\mathbf{h}_k)$ and $\mu_{1k}^d(\mathbf{h}_k)$ over the predictor space $\overline{\mathbf{X}}_1 = \{\mathbf{S}, \mathbf{Z}(0)\}$. Given a function $\boldsymbol{\zeta}(\cdot)$, at decision points $k = 2,\ldots,K$, we train the random forests for $\mu_{0k}^d(\mathbf{h}_k)$ and $\mu_{1k}^d(\mathbf{h}_k)$ over the predictor space $\overline{\mathbf{X}}_k = \{\mathbf{S}, \boldsymbol{\zeta}(k)\}$. We train the random forest for $\mu_{0k}^d(\mathbf{h}_k)$ on patients having $U^L > \tau_k$ and $A_k = 0$, and we train the random forest for $\mu_{1k}^d(\mathbf{h}_k)$ on patients having $U^L > \tau_k$ and $A_k = 1$.

3.3.5 ReLiVE: Value Estimation for a Fixed Regime

We now introduce the residual life value estimator (ReLiVE), which estimates value (3.1) for a fixed regime $d \in \mathscr{D}$. Beginning with decision point K, given patient history $\mathbf{H}_K = \mathbf{h}_K$, we train the random forests for $\mu^d_{0K}(\mathbf{h}_K)$ and $\mu^d_{1K}(\mathbf{h}_K)$ to estimate the outcome

$$\frac{\Delta^L(U^L - \tau_K)}{\widehat{\mathscr{K}}_K(U^L|\mathbf{h}_K, a_K)},$$

where $\widehat{\mathscr{K}}_k(x|\mathbf{h}_k, a_k) = \widehat{G}(x)/\widehat{G}(\tau_k)$ for $k = 1, \ldots, K$, and $\widehat{G}(t)$ is the Kaplan-Meier estimator for the survival function of censoring time C at time t (Kaplan and Meier (1958)). For $k = 1, \ldots, K$, denote the random forest estimators of $\mu^d_{0k}(\mathbf{h}_k)$ and $\mu^d_{1k}(\mathbf{h}_k)$ as $\widehat{\mu}^d_{0k}(\mathbf{h}_k)$ and $\widehat{\mu}^d_{1k}(\mathbf{h}_k)$, respectively. We estimate the Q-function at decision point K as

$$\widehat{Q}^d_K(\mathbf{h}_K, a_K) = \widehat{\mu}^d_{0K}(\mathbf{h}_K) + a_K\{\widehat{\mu}^d_{1K}(\mathbf{h}_K) - \widehat{\mu}^d_{0K}(\mathbf{h}_K)\}.$$

At decision points $k = K-1, \ldots, 1$, given patient history $\mathbf{H}_k = \mathbf{h}_k$, we then train the random forests for $\mu^d_{0k}(\mathbf{h}_k)$ and $\mu^d_{1k}(\mathbf{h}_k)$ to estimate the outcome

$$\frac{\Delta^L(U^L - \tau_k)}{\widehat{\mathscr{K}}_k(U^L|\mathbf{h}_k, a_k)} + \frac{I(U^L > \tau_{k+1})\widehat{V}^d_{k+1}(\mathbf{h}_{k+1})}{\widehat{\mathscr{K}}_k(\tau_{k+1}|\mathbf{h}_k, a_k)},$$

where $\widehat{V}^d_k(\mathbf{h}_k) = \widehat{Q}^d_k\{\mathbf{h}_k, d_k(\mathbf{h}_k)\}$. We take

$$\widehat{Q}^d_k(\mathbf{h}_k, a_k) = \widehat{\mu}^d_{0k}(\mathbf{h}_k) + a_k\{\widehat{\mu}^d_{1k}(\mathbf{h}_k) - \widehat{\mu}^d_{0k}(\mathbf{h}_k)\}.$$

Given m patients, we define ReLiVE to be $\widehat{\mathscr{V}}(d) = m^{-1}\sum_{i=1}^m \widehat{V}^d_1(\mathbf{h}_{1i})$, where $\widehat{V}^d_1(\mathbf{h}_{1i}) = \widehat{Q}^d_1\{\mathbf{h}_{1i}, d_1(\mathbf{h}_{1i})\}$ for patient i with history $\mathbf{h}_{1i}$.

3.3.6 ReLiVE-Q: Estimation of an Optimal Regime & Its Value

Next, we introduce the residual life value estimator Q-learning method (ReLiVE-Q), which estimates an optimal treatment regime $d^{opt} \in \mathscr{D}$ and its value $\mathscr{V}(d^{opt})$. Beginning with decision point K, given patient history $\mathbf{H}_K = \mathbf{h}_K$, we train the random forests for $\mu^{d^{opt}}_{0K}(\mathbf{h}_K)$

and $\mu_{1K}^{d^{opt}}(\mathbf{h}_K)$ to estimate the outcome

$$\frac{\Delta^L(U^L-\tau_K)}{\widehat{\mathscr{K}}_K(U^L|\mathbf{h}_K,a_K)}.$$

We estimate the Q-function for d^{opt} at decision point K as

$$\widehat{Q}_K^{d^{opt}}(\mathbf{h}_K,a_K)=\widehat{\mu}_{0K}^{d^{opt}}(\mathbf{h}_K)+a_K\{\widehat{\mu}_{1K}^{d^{opt}}(\mathbf{h}_K)-\widehat{\mu}_{0K}^{d^{opt}}(\mathbf{h}_K)\},$$

and we let $\widehat{d}_K^{opt}(\mathbf{h}_K)=\underset{a_K\in\mathscr{A}_K}{\operatorname{argmax}}\,\widehat{Q}_K^{d^{opt}}(\mathbf{h}_K,a_K)=I\{\widehat{\mu}_{1K}^{d^{opt}}(\mathbf{h}_K)>\widehat{\mu}_{0K}^{d^{opt}}(\mathbf{h}_K)\}$, where $a_K=0$ is the standard treatment.

At decision points $k=K-1,\ldots,1$, given patient history $\mathbf{H}_k=\mathbf{h}_k$, we then train the random forests for $\mu_{0k}^{d^{opt}}(\mathbf{h}_k)$ and $\mu_{1k}^{d^{opt}}(\mathbf{h}_k)$ to estimate the outcome

$$\frac{\Delta^L(U^L-\tau_k)}{\widehat{\mathscr{K}}_k(U^L|\mathbf{h}_k,a_k)}+\frac{I(U^L>\tau_{k+1})\widehat{V}_{k+1}^{d^{opt}}(\mathbf{h}_{k+1})}{\widehat{\mathscr{K}}_k(\tau_{k+1}|\mathbf{h}_k,a_k)},$$

where $\widehat{V}_k^{d^{opt}}(\mathbf{h}_k)=\widehat{Q}_k^{d^{opt}}\{\mathbf{h}_k,\widehat{d}_k^{opt}(\mathbf{h}_k)\}$. We estimate the Q-function for d^{opt} at decision points $k=K-1,\ldots,1$ as

$$\widehat{Q}_k^{d^{opt}}(\mathbf{h}_k,a_k)=\widehat{\mu}_{0k}^{d^{opt}}(\mathbf{h}_k)+a_k\{\widehat{\mu}_{1k}^{d^{opt}}(\mathbf{h}_k)-\widehat{\mu}_{0k}^{d^{opt}}(\mathbf{h}_k)\},$$

and we let $\widehat{d}_k^{opt}(\mathbf{h}_k)=\underset{a_k\in\mathscr{A}_k}{\operatorname{argmax}}\,\widehat{Q}_k^{d^{opt}}(\mathbf{h}_k,a_k)=I\{\widehat{\mu}_{1k}^{d^{opt}}(\mathbf{h}_k)>\widehat{\mu}_{0k}^{d^{opt}}(\mathbf{h}_k)\}$, where $a_k=0$ is the standard treatment.

We estimate an optimal treatment regime to be $\widehat{d}^{opt}=(\widehat{d}_1^{opt},\ldots,\widehat{d}_K^{opt})$. Given m patients, we estimate the value of an optimal treatment regime to be $\widehat{\mathscr{V}}(d^{opt})=m^{-1}\sum_{i=1}^m\widehat{V}_1^{d^{opt}}(\mathbf{h}_{1i})$, where $\widehat{V}_1^{d^{opt}}(\mathbf{h}_{1i})=\widehat{Q}_1^{d^{opt}}\{\mathbf{h}_{1i},\widehat{d}_1^{opt}(\mathbf{h}_{1i})\}$ for patient i with history $\mathbf{h}_{1i}$.

3.3.7 Testing & Validation Procedures

We evaluate the utility of ReLiVE-Q through testing and validation procedures that compare value estimates for an optimal treatment regime to value estimates for two fixed treatment regimes. First, the procedures estimate value (3.1) for two fixed regimes, the observed treatment regime and the no treatment regime, using ReLiVE. In the observed treatment regime d^{obs}, the decision rule at each decision point, $k=1,\ldots,K$, sets $d_k(\mathbf{h}_k)=a_k$, where a_k is the treatment actually administered to the patient at decision point k. In the no

treatment regime d^{no}, the decision rule at each decision point, $k = 1, \ldots, K$, sets $d_k(\mathbf{h}_k) = 0$. Additionally, the testing and validation procedures estimate an optimal treatment regime d^{opt} and its value $\mathcal{V}(d^{opt})$ using ReLiVE-Q.

We implement the testing procedure in both the simulation study and the MIMIC-III data application. The testing procedure obtains N value estimates $\widehat{\mathcal{V}}^{test}(d)$ for each regime $d \in \{d^{opt}, d^{obs}, d^{no}\}$ using N unique testing data sets. For N iterations, the procedure randomly assigns ω% of the m patients in the data to a training data set and $(100-\omega)$% of the patients to a testing data set, where $\omega \in (0, 100)$. For each iteration, the procedure computes $\widehat{\mathcal{V}}^{test}(d)$ for the fixed regimes by conducting ReLiVE separately on the training and testing data sets. $\widehat{\mathcal{V}}^{test}(d)$ for $d \in \{d^{obs}, d^{no}\}$ is taken to be the ReLiVE estimate $m_{test}^{-1} \sum_{i=1}^{m_{test}} \widehat{V}_1^d(\mathbf{h}_{1i})$ for the m_{test} patients in the testing data set.

For each iteration, the testing procedure also estimates d^{opt} and its value $\mathcal{V}(d^{opt})$ via a cross-validated adaptation of ReLiVE-Q. At each decision point, $k = 1, \ldots, K$, the procedure fits separate Q-models for d^{opt} on the training and testing data sets. Let the Q-model fit on the training data set at decision point k be denoted by $\widehat{Q}_{k,train}^{d^{opt}}(\mathbf{h}_k, a_k)$, and let the Q-model fit on the testing data set be denoted by $\widehat{Q}_{k,test}^{d^{opt}}(\mathbf{h}_k, a_k)$. At each backward iterative step in ReLiVE-Q, the procedure estimates the optimal treatment decision for patients in both the training and testing data sets using the training Q-model, so $\widehat{d}_k^{opt}(\mathbf{h}_k) = \underset{a_k \in \mathcal{A}_k}{\text{argmax}}\, \widehat{Q}_{k,train}^{d^{opt}}(\mathbf{h}_k, a_k)$. To obtain a cross-validated value estimate for $\widehat{d}^{opt}$, the procedure then sets $\widehat{V}_k^{d^{opt}}(\mathbf{h}_k) = \widehat{Q}_{k,train}^{d^{opt}}\{\mathbf{h}_k, \widehat{d}_k^{opt}(\mathbf{h}_k)\}$ for patients in the training data set, and sets $\widehat{V}_k^{d^{opt}}(\mathbf{h}_k) = \widehat{Q}_{k,test}^{d^{opt}}\{\mathbf{h}_k, \widehat{d}_k^{opt}(\mathbf{h}_k)\}$ for patients in the testing data set. $\widehat{\mathcal{V}}^{test}(d^{opt})$ is taken to be the ReLiVE-Q estimate $m_{test}^{-1} \sum_{i=1}^{m_{test}} \widehat{V}_1^{d^{opt}}(\mathbf{h}_{1i})$ for the m_{test} patients in the testing data set.

In the simulation study, we also implement a validation procedure that builds on the testing procedure. The validation procedure obtains N value estimates $\widehat{\mathcal{V}}^{val}(d)$ for each regime $d \in \{d^{opt}, d^{obs}, d^{no}\}$ by randomly generating N potential event times $T^*(d)$ for each patient in a validation data set. The validation procedure begins by generating baseline covariate data $\mathbf{X}_1 = \{\mathbf{S}, \mathbf{Z}(0)\}$ for 100,000 patients. All variables in the validation procedure are generated according to the simulation strategy described in Section 3.4.1.

The validation procedure first computes N estimates of $\widehat{\mathcal{V}}^{val}(d)$ for the fixed regimes d^{obs} and d^{no}. At decision points $k = 2, \ldots, K$, the procedure iteratively generates each patient's new covariate information $\mathbf{X}_k = \{\mathbf{Z}(t) : \tau_{k-1} < t \leq \tau_k\}$ as a function of the prior treatment decisions. For N iterations, given the fixed regime d, the validation procedure then ran-

domly generates $T^*(d)$ for each patient and computes $g\{T^*(d)\} = \sum_{j=1}^{K} I[\min\{T^*(d), L\} > \tau_j] \cdot [\min\{T^*(d), L\} - \tau_j]$. For each iteration, $\widehat{\mathcal{V}}^{val}(d)$ for $d \in \{d^{obs}, d^{no}\}$ is taken to be the mean of $g\{T^*(d)\}$ over the 100,000 validation patients.

The validation procedure also computes the value estimate $\widehat{\mathcal{V}}^{val}(d^{opt})$ for the N optimal treatment regimes estimated in the N iterations of the testing procedure. For each estimated regime, the validation procedure iteratively generates each patient's new covariate information $\mathbf{X}_k$ at decision points $k = 2, \ldots, K$ as a function of the prior estimated optimal treatment decisions. For each $\widehat{d}^{opt}$, the validation procedure then randomly generates $T^*(\widehat{d}^{opt})$ for each patient and computes $g\{T^*(\widehat{d}^{opt})\} = \sum_{j=1}^{K} I[\min\{T^*(\widehat{d}^{opt}), L\} > \tau_j] \cdot [\min\{T^*(\widehat{d}^{opt}), L\} - \tau_j]$. For each of the N estimated optimal regimes, $\widehat{\mathcal{V}}^{val}(d^{opt})$ is taken to be the mean of $g\{T^*(\widehat{d}^{opt})\}$ over the 100,000 validation patients.

When implementing the testing and validation procedures for the simulation study and the MIMIC-III data application, we specify that the random forests described in Section 3.3.4 be comprised of 250 decision trees, and we grow each decision tree until it reaches the minimum terminal node size of 50. For each tree, we specify that ceiling($r_k/3$) features from the predictor space be randomly selected as partition candidates in each step of the decision tree algorithm, where r_k is the dimension of the predictor space at decision point k. We repeat the testing and validation procedures four times, defining $\boldsymbol{\zeta}(\cdot)$ in terms of the baseline vector, the last-value carried forward vector, the average vector, and the context vector, as described in Section 3.3.3. Note, $\widehat{\mathcal{V}}^{test}(d^{opt})$, $\widehat{\mathcal{V}}^{test}(d^{obs})$, and $\widehat{\mathcal{V}}^{test}(d^{no})$ are dependent on the fitted Q-models, as is $\widehat{\mathcal{V}}^{val}(d^{opt})$. Consequently, these value estimates are dependent on the functional form of $\boldsymbol{\zeta}(\cdot)$. In contrast, $\widehat{\mathcal{V}}^{val}(d^{obs})$ and $\widehat{\mathcal{V}}^{val}(d^{no})$ are independent of the fitted Q-models, so these estimates do not depend on $\boldsymbol{\zeta}(\cdot)$.

3.3.8 Software

We conduct the simulation study and MIMIC-III data application in R and Python. We use the **randomForest** R package to construct the random forest Q-models (Liaw and Wiener (2002)), and we use the **survival** package to compute the Kaplan-Meier estimates for the censoring weights (Therneau and Grambsch (2000)). We construct the context vectors $\boldsymbol{\zeta}^C(\cdot)$ in Python using the **Keras** library (Chollet (2015)). We provide the code used to conduct the simulation study at `https://github.com/gmrhodes/ReLiVE-Q`.

3.4 Simulation Study

3.4.1 Simulation Strategy

We conduct a simulation study to evaluate the utility of ReLiVE-Q. We consider $K = 4$ decision points that occur at times $(\tau_1, \ldots, \tau_4) = (0, 3, 6, 9)$. We generate a single data set of $m = 10,000$ patients. For each patient, we generate a single treatment variable $A_k \in \mathscr{A}_k = \{0, 1\}$ at decision points $k = 1, \ldots, 4$ as $A_k \sim$Bernoulli(0.5). For each patient, we also generate a single covariate measured only at baseline, $S \sim \mathbb{U}(0, 1)$, where $\mathbb{U}$ denotes the uniform distribution. We consider ten measurement times $\mathscr{M} = (0, 1, \ldots, 9)$, and we generate three longitudinal covariates $Z_l(t) = B_l(t) + \epsilon_l(t)$ at each $t \in \mathscr{M}$, where $l = 1, 2, 3$. We specify the measurement error to be $\epsilon_l(t) \sim \mathbb{N}(0, 0.5^2)$, where $\mathbb{N}(\mu, \Sigma)$ denotes a normal distribution with mean μ and variance-covariance Σ. Define $a_{(j)} = a_1 I\{j \in (1,2,3)\} + a_2 I\{j \in (4,5,6)\} + a_3 I\{j \in (7,8,9)\} + a_4 I(j = 10)$, $j = 1, \ldots, 10$. For $l = 1, 2, 3$, we generate $B_l(t)$ according to the piecewise linear model

$$B_l(t) = c_0^l + \mathfrak{b}_0^l + \sum_{j=1}^{10} I\{t > (j-1)\} \cdot \{c_j^l + \mathfrak{b}_j^l + a_{(j)}(d_j^l + \mathfrak{f}_j^l)\} \cdot \{t - (j-1)\},$$

where $(c_0^l, \ldots, c_{10}^l)$ and $(d_1^l, \ldots, d_{10}^l)$ are fixed effects specified in Table (3.1), and $\mathfrak{b}^l = (\mathfrak{b}_0^l, \ldots, \mathfrak{b}_{10}^l)$ and $\mathfrak{f}^l = (\mathfrak{f}_1^l, \ldots, \mathfrak{f}_{10}^l)$ are random effects. Let rep(x, y) denote a vector containing x repeated y times, and let diag represent a diagonal matrix. We specify $\mathfrak{b}^l \sim \mathbb{N}(\mathbf{0}, \mathbf{D_1})$ and $\mathfrak{f}^l \sim \mathbb{N}(\mathbf{0}, \mathbf{D_2})$, where $\mathbf{D_1} = \text{diag}\{1, \text{rep}(0.025, 5), \text{rep}(0.01, 5)\}$ and $\mathbf{D_2} = \text{diag}\{1, \text{rep}(0.025, 5), \text{rep}(0.01, 4)\}$.

In two distinct analyses, we conduct the testing and validation procedures described in

Table 3.1: Fixed effects in the piecewise constant function $B_l(t)$.

l	c_0	c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
1	-2	4	-7	5	-2.5	3.5	-5	1.5	2	-2	1
2	3	-2	3	1	-3	-1	3.5	-2	-0.5	0.5	0.75
3	5	-0.5	-2	2	-2	2	3	-2	-1	2	-0.5

l	d_1	d_2	d_3	d_4	d_5	d_6	d_7	d_8	d_9	d_{10}
1	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
2	-0.25	-0.25	-0.25	-0.25	-0.25	-0.25	-0.25	-0.25	-0.25	-0.25
3	0.25	0.25	0.25	0.25	0.25	-0.25	-0.25	-0.25	-0.25	-0.25

Section 3.3.7 for $N = 500$ iterations, assigning $\omega = 50\%$ of patients to the training data set and 50% to the testing data set. In both analyses, we generate censoring times as $C \sim \mathbb{U}(0, 55)$, and we impose a restricted lifetime of $L = 50$. In each analysis, we use a distinct survival time generation process. For $t > 0$, define $A(t) = A_1 I(0 < t \leq 3) + A_2 I(3 < t \leq 6) + A_3 I(6 < t \leq 9) + A_4 I(t > 9)$. In the first analysis, we generate T according to the accelerated failure time (AFT) model

$$\nu = \int_0^T \exp\Big\{\beta_1 B_1(u) + \beta_2 B_2(u) + \beta_3 B_3(u) + \beta_4 S + \beta_5 A(u) + \beta_6 B_1(u)A(u) + \beta_7 B_2(u)A(u) + \beta_8 B_3(u)A(u) + \beta_9 SA(u)\Big\} d\,u,$$

where $(\beta_1, \ldots, \beta_9) = (0.25, -0.25, 0.25, 0.5, -1.5, -0.25, 0.25, 0.25, 0.25)$, $\nu = \exp(\theta)$, and $\theta \sim \mathbb{N}(3, 1)$. The resulting U^L have median (med)=6.4 with Quartile 1 (Q1)=3.2 and Quartile 3 (Q3)=10.2. Approximately 20.4% of the restricted survival times are censored, i.e. $\Delta^L = 0$.

In the second analysis, we generate T according to a Cox proportional hazards model with an exponential baseline hazard, specified as

$$h(t) = \lambda \exp\Big\{\beta_1 B_1(t) + \beta_2 B_2(t) + \beta_3 B_3(t) + \beta_4 S + \beta_5 A(t) + \beta_6 B_1(t)A(t) + \beta_7 B_2(t)A(t) + \beta_8 B_3(t)A(t) + \beta_9 SA(t)\Big\},$$

where $(\beta_1, \ldots, \beta_9) = (0.25, -0.5, 0.25, 0.5, -1.5, -0.5, -0.5, 0.25, -0.25)$ and $\lambda = 0.1$. The resulting U^L have med=5.4 with Q1=1.6 and Q3=8.7. Approximately 16.4% of the restricted survival times are censored, i.e. $\Delta^L = 0$. For technical details of the survival time generation process, refer to Appendix B, Section B.3.

We conduct both the AFT and Cox analyses four times, representing the longitudinal covariates using the baseline vector $\boldsymbol{\zeta}^B(\cdot)$, the last-value carried forward vector $\boldsymbol{\zeta}^L(\cdot)$, the average vector $\boldsymbol{\zeta}^A(\cdot)$, and the context vector $\boldsymbol{\zeta}^C(\cdot)$, as described in Section 3.3.3. In the testing procedures, we construct $\boldsymbol{\zeta}^C(\cdot)$ using the simulated data. We train the LSTM autoencoders for 500 epochs using the Adam optimization algorithm (Kingma and Ba (2017)), and we use cross-validation to adaptively select the learning rate from $\{1e^{-3}, 1e^{-4}\}$. We specify $\boldsymbol{\psi}_l(\tau_k)$ to have a constant dimension of 5 for $l = 1, 2, 3$ and $k = 2, 3, 4$. In the validation procedure for each analysis, we construct $\boldsymbol{\zeta}^C(\cdot)$ using the trained LSTM autoencoders from the corresponding testing procedure; we do not re-train the networks on the validation data.

3.4.2 Simulation Results

The 500 testing value estimates $\widehat{\mathscr{V}}^{test}(d)$ and validation value estimates $\widehat{\mathscr{V}}^{val}(d)$ are plotted in Figures (3.1) and (3.2) for $d \in \{d^{opt}, d^{obs}, d^{no}\}$. Additionally, the mean and standard deviation of the value estimates are presented in Tables (3.2) and (3.3). In both the AFT and Cox analyses, the distributions of $\widehat{\mathscr{V}}^{test}(d)$ and $\widehat{\mathscr{V}}^{val}(d)$ are consistently higher for the optimal treatment regime than for the observed treatment regime or the no treatment regime. Thus, the simulation study supports that $\mathscr{V}(\widehat{d}^{opt}) > \mathscr{V}(d^{obs})$ and $\mathscr{V}(\widehat{d}^{opt}) > \mathscr{V}(d^{no})$, which allows us to conclude that ReLiVE-Q successfully produces reasonable estimates of an optimal treatment regime.

Moreover, in both the AFT and Cox analyses, Q-models fit with the baseline vector $\boldsymbol{\zeta}^{B}(\cdot)$ result in lower value estimates for the optimal treatment regime than those fit with the average vector $\boldsymbol{\zeta}^{A}(\cdot)$, the last-value carried forward vector $\boldsymbol{\zeta}^{L}(\cdot)$, or the context vector $\boldsymbol{\zeta}^{C}(\cdot)$. In the AFT analysis, Q-models fit with $\boldsymbol{\zeta}^{L}(\cdot)$ or $\boldsymbol{\zeta}^{C}(\cdot)$ result in the highest value estimates for the optimal treatment regime, while in the Cox analysis, Q-models fit with $\boldsymbol{\zeta}^{C}(\cdot)$ result in the highest. Thus, the simulation study supports that sophisticated functions are necessary to synthesize the complex trajectories of the longitudinal covariates. Furthermore, the study demonstrates that representing longitudinal covariates with context vectors in ReLiVE-Q can lead to improved estimation of an optimal treatment regime.

3.5 Application to MIMIC-III

3.5.1 Data Analysis

We estimate an optimal treatment regime for septic patients in the ICU using EMR data from the MIMIC-III database, described in Chapter 1. We consider $K = 10$ decision points that occur every four hours at $(\tau_1, \tau_2, \ldots, \tau_{10}) = (0, 4, \ldots, 36)$. We study a single treatment $A_k \in \{0, 1\}$ at each decision point, $k = 1, \ldots, 10$, where $A_k = 1$ if the patient is provided vasopressors in the four-hour time-window $[\tau_k, \tau_{k+1})$ and $A_k = 0$ otherwise. We specify a restricted lifetime of $L = 40$ and conduct the testing procedure described in Section 3.3.7 for $N = 100$ iterations, assigning $\omega = 70\%$ of patients to the training data set and 30% to the testing data set.

We study a predictor space containing a single baseline covariate S and 14 longitudinal covariates $\mathbf{Z}(\cdot)$. The baseline covariate of interest is an indicator of whether the patient was

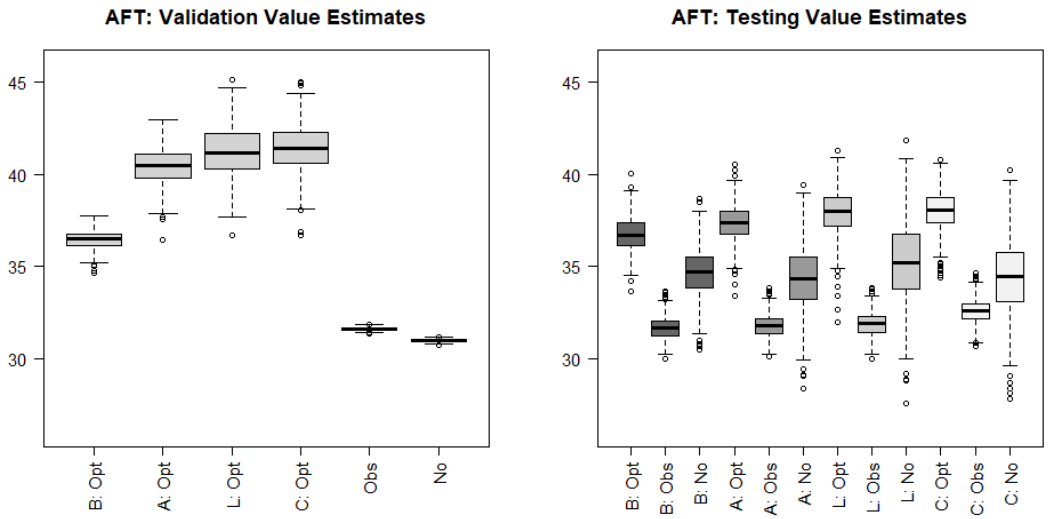

Figure 3.1: AFT Simulation Study: Boxplots of the value estimates from the validation procedure (left) and the testing procedure (right) for an optimal treatment regime (Opt), the observed treatment regime (Obs), and the no treatment regime (No). For scenarios dependent on the Q-models, value estimates are presented using the baseline vector (B), the average vector (A), the last-value carried forward vector (L), and the context vector (C).

Table 3.2: AFT Simulation Study: The mean (standard deviation) of the value estimates from the testing procedure (Test) and the validation procedure (Valid) for an optimal treatment regime (Opt), the observed treatment regime (Obs), and the no treatment regime (No). For scenarios dependent on the Q-models, value estimates are presented using the baseline vector (B), the average vector (A), the last-value carried forward vector (L), and the context vector (C).

AFT	Opt		Obs		No	
	Test	Valid	Test	Valid	Test	Valid
B	36.8 (0.9)	36.4 (0.5)	31.6 (0.6)	-	34.7 (1.3)	-
A	37.4 (1.0)	40.5 (1.0)	31.8 (0.6)	-	34.4 (1.7)	-
L	37.9 (1.3)	41.3 (1.4)	31.9 (0.6)	-	35.2 (2.2)	-
C	38.0 (1.1)	41.4 (1.3)	32.6 (0.7)	-	34.4 (2.0)	-
*	-	-	-	31.6 (0.1)	-	31.0 (0.1)

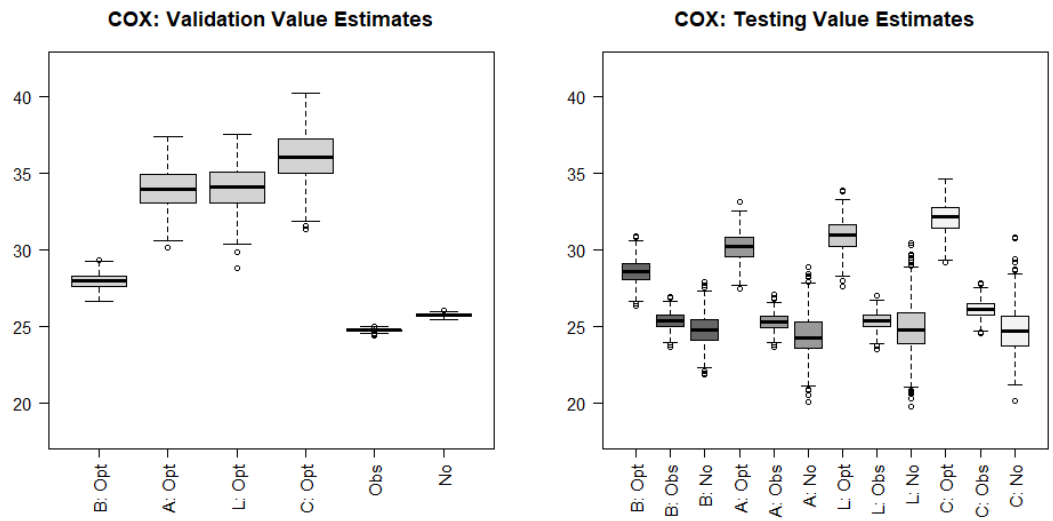

Figure 3.2: Cox Simulation Study: Boxplots of the value estimates from the validation procedure (left) and the testing procedure (right) for an optimal treatment regime (Opt), the observed treatment regime (Obs), and the no treatment regime (No). For scenarios dependent on the Q-models, value estimates are presented using the baseline vector (B), the average vector (A), the last-value carried forward vector (L), and the context vector (C).

Table 3.3: Cox Simulation Study: The mean (standard deviation) of the value estimates from the testing procedure (Test) and the validation procedure (Valid) for an optimal treatment regime (Opt), the observed treatment regime (Obs), and the no treatment regime (No). For scenarios dependent on the Q-models, value estimates are presented using the baseline vector (B), the average vector (A), the last-value carried forward vector (L), and the context vector (C).

Cox	Opt		Obs		No	
	Test	Valid	Test	Valid	Test	Valid
B	28.6 (0.8)	28.0 (0.5)	25.4 (0.6)	-	24.8 (1.0)	-
A	30.3 (0.9)	34.0 (1.3)	25.3 (0.6)	-	24.4 (1.4)	-
L	31.0 (1.0)	34.1 (1.5)	25.4 (0.6)	-	24.9 (1.7)	-
C	32.1 (1.0)	36.1 (1.6)	26.2 (0.6)	-	24.8 (1.5)	-
*	-	-	-	24.8 (0.1)	-	25.8 (0.1)

previously admitted to the ICU during the given hospital stay. The longitudinal covariates of interest include the cumulative amount of intravenous (IV) fluid administered to the patient since hospital admission, as well as a longitudinal indicator of mechanical ventilator dependence. We also study 12 longitudinal vital signs and laboratory values: albumin, calcium, magnesium, hemoglobin, arterial lactate, arterial pH, fraction of inspired oxygen (FiO2), peripheral oxygen saturation (SpO2), Sequential Organ Failure Assessment (SOFA) score, respiratory rate, heart rate, and systolic blood pressure.

In Chapter 2, we constructed four distinct context vectors for the studied longitudinal covariates by training LSTM autoencoders under four different hyperparameter settings. We construct the context vector $\boldsymbol{\zeta}^C(\cdot)$ using the most parsimonious settings presented in Chapter 2, as these settings were shown to result in competitively accurate predictions of restricted residual life. Specifically, we train the LSTM autoencoders for 150 epochs using the Adam optimization algorithm (Kingma and Ba (2017)). We use Bayesian optimization to adaptively select the learning rate from $\{1e^{-2}, 1e^{-3}, 1e^{-4}\}$ (O'Malley et al. (2019)). We specify $\boldsymbol{\psi}_l(\tau_k)$ to have a constant dimension of 3 for $l = 1, \ldots, 14$ and $k = 2, \ldots, 10$.

3.5.2 MIMIC-III Results

The 100 testing value estimates $\widehat{\mathscr{V}}^{test}(d)$ are plotted in Figure (3.3) for $d \in \{d^{opt}, d^{obs}, d^{no}\}$, and the mean and standard deviation of the value estimates are presented in Table (3.4). The distribution of $\widehat{\mathscr{V}}^{test}(d)$ is consistently higher for the optimal treatment regime than for the observed treatment regime or the no treatment regime. As in the simulation study, these results support that $\mathscr{V}(\widehat{d}^{opt}) > \mathscr{V}(d^{obs})$ and $\mathscr{V}(\widehat{d}^{opt}) > \mathscr{V}(d^{no})$. Thus, we can conclude that ReLiVE-Q successfully uses EMR data to produce reasonable estimates of an optimal treatment regime for septic patients in the ICU.

Again, Q-models fit with the baseline vector $\boldsymbol{\zeta}^B(\cdot)$ result in lower value estimates for the optimal treatment regime than those fit with the average vector $\boldsymbol{\zeta}^A(\cdot)$, the last-value carried forward vector $\boldsymbol{\zeta}^L(\cdot)$, or the context vector $\boldsymbol{\zeta}^C(\cdot)$. Moreover, Q-models fit with $\boldsymbol{\zeta}^C(\cdot)$ result in the highest value estimates for the optimal treatment regime. Similar to the simulation study, this application suggests that representing longitudinal covariates with context vectors in ReLiVE-Q can lead to improved estimation of an optimal treatment regime.

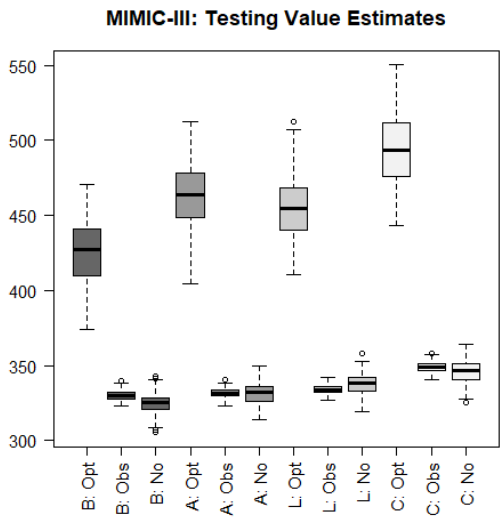

Figure 3.3: MIMIC-III Application: Boxplots of the value estimates from the testing procedure for an optimal treatment regime (Opt), the observed treatment regime (Obs), and the no treatment regime (No). Value estimates are presented from Q-models using the baseline vector (B), the average vector (A), the last-value carried forward vector (L), and the context vector (C).

Table 3.4: MIMIC-III Application: The mean (standard deviation) of the value estimates from the testing procedure for an optimal treatment regime (Opt), the observed treatment regime (Obs), and the no treatment regime (No). Value estimates are presented from Q-models using the baseline vector (B), the average vector (A), the last-value carried forward vector (L), and the context vector (C).

	MIMIC-III Testing Value Estimates		
	Opt	Obs	No
B	425.4 (22.4)	330.3 (3.4)	324.8 (7.7)
A	463.8 (22.0)	331.6 (3.3)	331.7 (7.4)
L	454.5 (22.5)	334.2 (3.0)	337.9 (7.6)
C	494.4 (23.3)	349.1 (3.4)	345.9 (8.1)

3.5.3 Covariate Importance

Having demonstrated the utility of ReLiVE-Q via the testing procedure, we estimate an optimal treatment regime for septic patients in the ICU by conducting ReLiVE-Q on $m = 20,952$ septic patients in MIMIC-III. We conduct ReLiVE-Q under the settings described in Section 3.5.1. In particular, we use context vectors to summarize the trajectories of the longitudinal covariates, so the predictor space is comprised of the baseline covariate S and the context vector $\boldsymbol{\zeta}^C(k)$, where $\boldsymbol{\zeta}^C(1) = \mathbf{Z}(0)$ and $\boldsymbol{\zeta}^C(k) = \{\boldsymbol{\psi}_1(\tau_k), \ldots, \boldsymbol{\psi}_{14}(\tau_k)\}$ for $k = 2, \ldots, 10$. Because $\boldsymbol{\psi}_l(\tau_k)$ is a 3-dimensional vector for $l = 1, \ldots, 14$, each longitudinal covariate is represented by three variables in the predictor space at decision points $k = 2, \ldots, 10$.

At each decision point, we compute two variable importance measures for each variable in the predictor space: one from the random forest for $\mu_{0k}^{d^{opt}}(\mathbf{h}_k) = Q_k^{d^{opt}}(\mathbf{h}_k, 0)$, and one from the random forest for $\mu_{1k}^{d^{opt}}(\mathbf{h}_k) = Q_k^{d^{opt}}(\mathbf{h}_k, 1)$. The variable importance measures quantify the variables' impact on the accuracy of the random forests' Q-function estimates. To compute the variable importance measures for a given forest, we first compute the mean squared error (MSE) of each tree in the forest using its out-of-bag (OOB) data, where the OOB data for a given tree is the sample of patients who were excluded from the tree's bootstrap training sample. We then permute the variable of interest and re-calculate the OOB MSE for each tree. The variable importance measure is taken to be the difference in the two OOB MSEs, averaged over all trees in the forest and normalized by the standard deviation of the differences (Liaw and Wiener (2002)). For each forest, we construct a single variable importance measure for each longitudinal covariate $l = 1, \ldots, 14$ at decision points $k = 2, \ldots, 10$ by summing the variable importance measures for the elements of $\boldsymbol{\psi}_l(\tau_k)$.

The variable importance measures from the random forests for $\mu_{0k}^{d^{opt}}(\mathbf{h}_k)$ and $\mu_{1k}^{d^{opt}}(\mathbf{h}_k)$ are depicted in Figures (3.4) and (3.5), respectively. The figures indicate that the cumulative amount of IV fluid administered to a patient, as well as a patient's albumin level and SOFA score, are influential predictors in both random forests. Thus, these covariates are important for accurate estimation of the Q-functions for septic patients in the ICU.

3.6 Discussion

In both the simulation study and the application to MIMIC-III, we demonstrate that the optimal treatment regime estimated via ReLiVE-Q results in higher estimates of value (3.1)

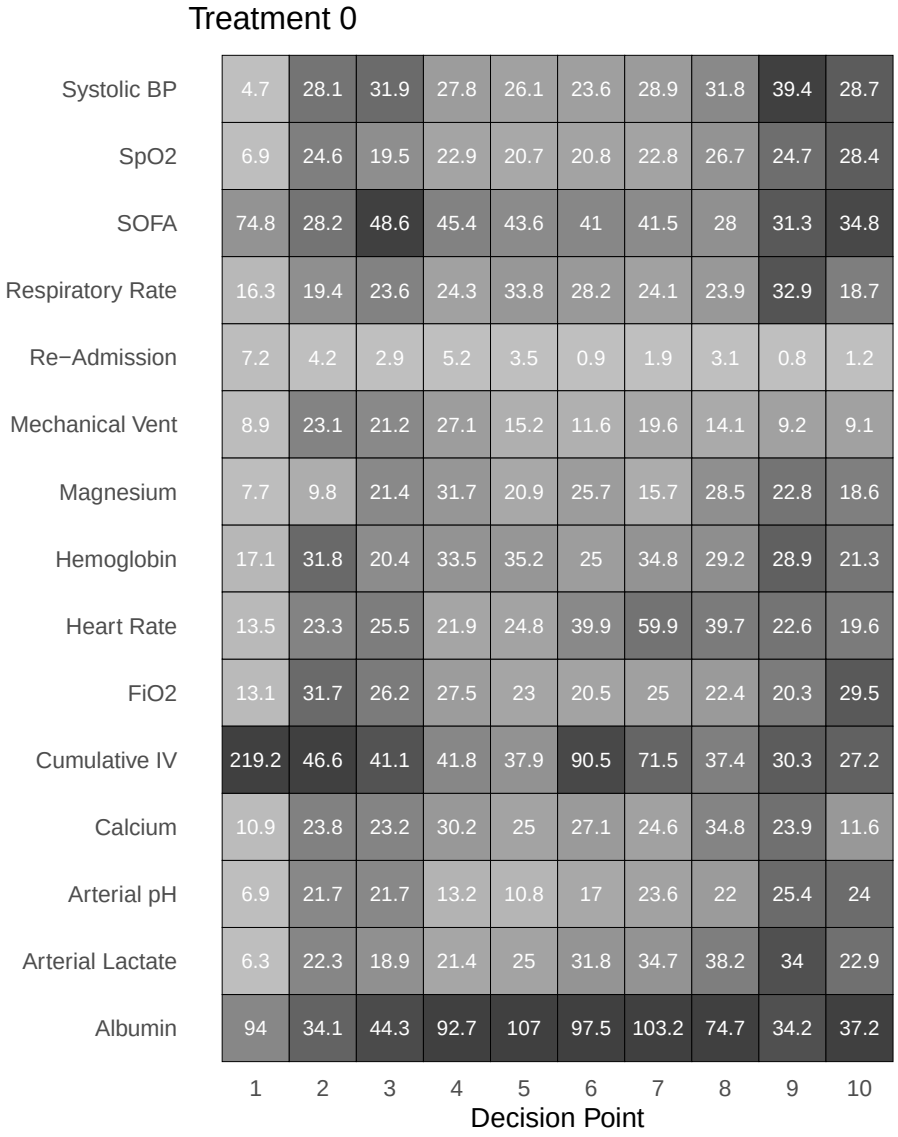

Figure 3.4: Variable importance measures from the random forest for $\mu_{0k}^{d^{opt}}(\mathbf{h}_k)$ at decision points $k = 1,\ldots,10$ for the baseline and longitudinal covariates studied in the MIMIC-III application. At each decision point, a color gradient is used to illustrate the relative magnitude of the variable importance measures.

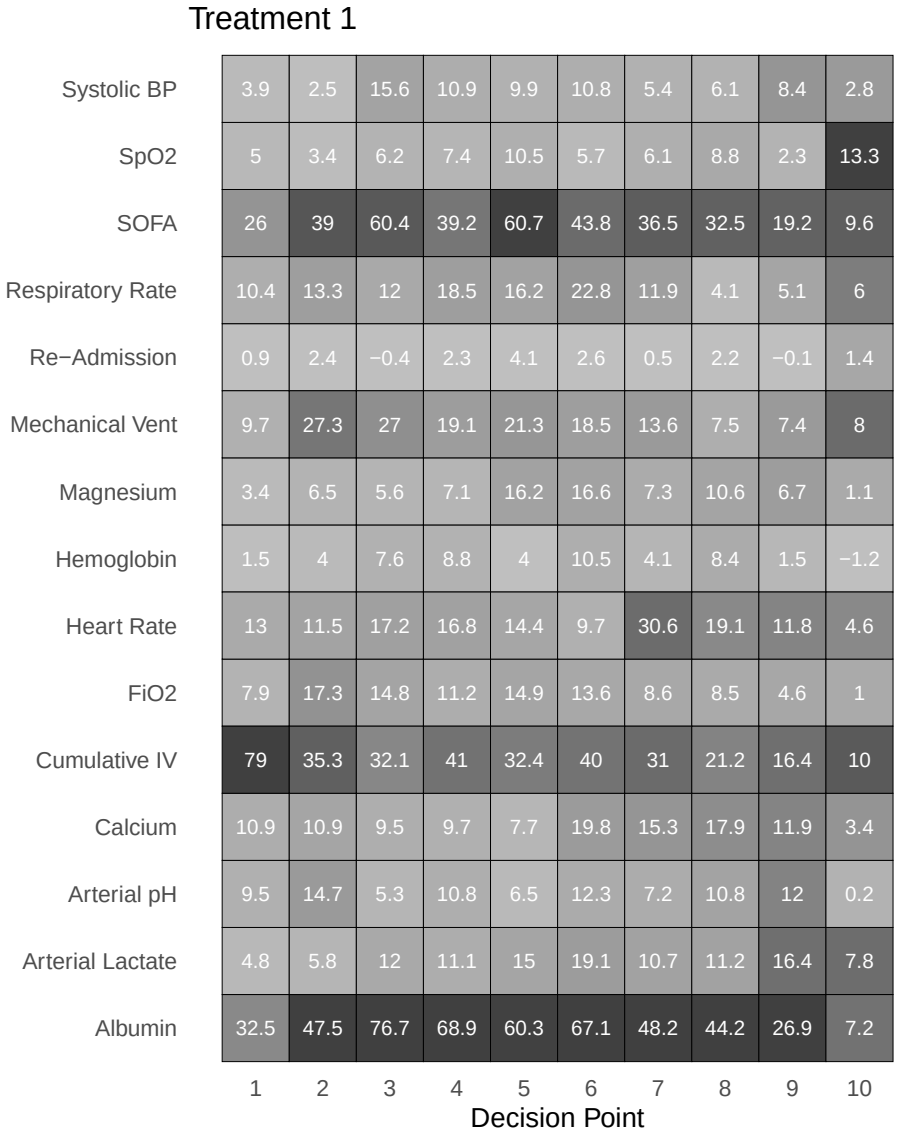

Figure 3.5: Variable importance measures from the random forest for $\mu_{1k}^{d^{opt}}(\mathbf{h}_k)$ at decision points $k = 1,\ldots,10$ for the baseline and longitudinal covariates studied in the MIMIC-III application. At each decision point, a color gradient is used to illustrate the relative magnitude of the variable importance measures.

than the observed treatment regime or the no treatment regime. Thus, we can expect the sum of restricted residual life across all decision points at which a patient is alive to be higher on average for patients who follow the estimated optimal treatment regime, as opposed to patients who follow the observed treatment regime or the no treatment regime. Thus, ReLiVE-Q leverages patient history to estimate personalized treatment regimes that maximize a clinically meaningful function of expected residual life. Because septic patients experience varying responses to treatment, this finding is especially significant for patients diagnosed with the life-threatening condition sepsis (László et al. (2015)).

Furthermore, we demonstrate that synthesizing longitudinal covariates with context vectors in ReLiVE-Q leads to improved estimation of an optimal treatment regime, as compared to simpler summary statistics. Further research is required to determine how alternative, sophisticated methods for summarizing longitudinal trajectories would perform.

In this study, we conduct ReLiVE-Q using random forests to estimate the Q-functions. Further research needs to be conducted to evaluate the performance of alternative estimation techniques. Moreover, we limit our study to clinical settings with two treatment options that are feasible for all patients at all decision points. Further research is required to determine how ReLiVE-Q performs in other treatment settings.

Lastly, we assume censoring is independent of treatment assignment, patient characteristics, restricted lifetime, and potential outcomes, and we estimate the censoring weights via the Kaplan-Meier method. However, it may be possible to relax this assumption and take censoring to be noninformative in the sense that the cause-specific hazard of censoring as a function of patient history and potential outcomes does not depend on the potential outcomes. In this case, a time-dependent Cox model could be used to incorporate covariate information into the censoring weight estimates. See Chapter 8 of Tsiatis et al. (2020) for details.

REFERENCES

Aggarwal, C. (2018). *Neural Networks and Deep Learning*. Springer International Publishing, Cham, Switzerland, 1st edition.

Bai, X., Tsiatis, A. A., Lu, W., and Song, R. (2017). Optimal treatment regimes for survival endpoints using a locally-efficient doubly-robust estimator from a classification perspective. *Lifetime Data Analysis*, 23(4):585–604.

Bates, D., Mächler, M., Bolker, B., and Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1):1–48.

Bellman, R. (1957). *Dynamic Programming*. Princeton University Press, Princeton, NJ.

Breiman, L. (1984). *Classification and Regression Trees*. Routledge, New York, NY.

Breiman, L. (2001). Random forests. *Machine Learning*, 45:5–32.

Chen, Y. Q. (2007). Additive expectancy regression. *Journal of the American Statistical Association*, 102(477):153–166.

Chollet, F. (2015). Keras. https://keras.io.

Evans, L., Rhodes, A., Alhazzani, W., Antonelli, M., Coopersmith, C. M., French, C., Machado, F. R., Mcintyre, L., Ostermann, M., Prescott, H. C., Schorr, C., Simpson, S., Wiersinga, W. J., Alshamsi, F., Angus, D. C., Arabi, Y., Azevedo, L., Beale, R., Beilman, G., Belley-Cote, E., Burry, L., Cecconi, M., Centofanti, J., Coz Yataco, A., De Waele, J., Dellinger, R. P., Doi, K., Du, B., Estenssoro, E., Ferrer, R., Gomersall, C., Hodgson, C., Hylander Møller, M., Iwashyna, T., Jacob, S., Kleinpell, R., Klompas, M., Koh, Y., Kumar, A., Kwizera, A., Lobo, S., Masur, H., McGloughlin, S., Mehta, S., Mehta, Y., Mer, M., Nunnally, M., Oczkowski, S., Osborn, T., Papathanassoglou, E., Perner, A., Puskarich, M., Roberts, J., Schweickert, W., Seckel, M., Sevransky, J., Sprung, C. L., Welte, T., Zimmerman, J., and Levy, M. (2021). Surviving sepsis campaign: International guidelines for management of sepsis and septic shock 2021. *Critical Care Medicine*, 49(11):e1063–e1143.

Friedman, J., Hastie, T., and Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1):1–22.

Gajardo, A., Bhattacharjee, S., Carroll, C., Chen, Y., Dai, X., Fan, J., Hadjipantelis, P. Z., Han, K., Ji, H., Zhu, C., Müller, H.-G., and Wang, J.-L. (2021). fdapace: Functional data analysis and empirical dynamics. https://CRAN.R-project.org/package=fdapace.

Gogia, S. (2019). *Fundamentals of Telemedicine and Telehealth*. Academic Press, 1st edition.

Goldberg, Y. and Kosorok, M. R. (2012). Q-learning with censored data. *Annals of Statistics*, 40(1):529–560.

Hager, R., Tsiatis, A. A., and Davidian, M. (2018). Optimal two-stage dynamic treatment regimes from a classification perspective with censored survival data. *Biometrics*, 74(4):1180–1192.

Harrell, F. E., Lee, K. L., and Mark, D. B. (1996). Tutorial in biostatistics: Multivariable prognostic models: Issues in developing models, evaluating assumptions and adequacy, and measuring and reducing errors. *Statistics in Medicine*, 15:361–387.

Hickey, G. L., Philipson, P., Jorgensen, A., and Kolamunnage-Dona, R. (2016). Joint modelling of time-to-event and multivariate longitudinal outcomes: Recent developments and issues. *BMC Medical Research Methodology*, 16(1):1–15.

Jiang, R., Lu, W., Song, R., and Davidian, M. (2017). On estimation of optimal treatment regimes for maximizing t-year survival probability. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 79(4):1165–1185.

Johnson, A. E. W., Pollard, T. J., Shen, L., Lehman, L. H., Feng, M., Ghassemi, M., Moody, B., Szolovits, P., Celi, L. A., and Mark, R. G. (2016). MIMIC-III, a freely accessible critical care database. *Scientific Data*, 3. 160035.

Kaplan, E. L. and Meier, P. (1958). Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association*, 53(282):457–481.

Kingma, D. P. and Ba, J. (2017). Adam: A method for stochastic optimization. arXiv:1412.6980 [cs.LG].

Komorowski, M. (2019). AI Clinician. https://github.com/matthieukomorowski/AI_Clinician.

László, I., Trásy, D., Molnár, Z., and Fazakas, J. (2015). Sepsis: From pathophysiology to individualized patient care. *Journal of Immunology Research*, 2015. 510436.

Liaw, A. and Wiener, M. (2002). Classification and regression by randomForest. *R News*, 2(3):18–22.

Lin, X., Lu, T., Yan, F., Li, R., and Huang, X. (2018). Mean residual life regression with functional principal component analysis on longitudinal data for dynamic prediction. *Biometrics*, 74(4):1482–1491.

Maguluri, G. and Zhang, C.-H. (1994). Estimation in the mean residual life regression model. *Journal of the Royal Statistical Society: Series B (Methodological)*, 56(3):477–489.

Murphy, S. A. (2003). Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 65(2):331–355.

Murphy, S. A. (2005). A generalization error for Q-learning. *Journal of Machine Learning Research*, 6:1073–1097.

O'Malley, T., Bursztein, E., Long, J., Chollet, F., Jin, H., and Invernizzi, L. (2019). KerasTuner. https://github.com/keras-team/keras-tuner.

Orellana, L., Rotnitzky, A., and Robins, J. M. (2010). Dynamic regime marginal structural mean models for estimation of optimal dynamic treatment regimes, Part I: Main content. *International Journal of Biostatistics*, 6(2). 8.

Rizopoulos, D., Molenberghs, G., and Lesaffre, E. M. E. H. (2017). Dynamic predictions with time-dependent covariates in survival analysis using joint modeling and landmarking. *Biometrical Journal*, 59(6):1261–1276.

Robins, J. M. (2004). Optimal structural nested models for optimal sequential decisions. In Lin, D. Y. and Heagerty, P. J., editors, *Proceedings of the Second Seattle Symposium in Biostatistics: Analysis of Correlated Data*, volume 179 of *Lecture Notes in Statistics*, pages 189–326. Springer, New York, NY.

Singer, M., Deutschman, C. S., Seymour, C. W., Shankar-Hari, M., Annane, D., Bauer, M., Bellomo, R., Bernard, G. R., Chiche, J.-D., Coopersmith, C. M., Hotchkiss, R. S., Levy, M. M., Marshall, J. C., Martin, G. S., Opal, S. M., Rubenfeld, G. D., van der Poll, T., Vincent, J.-L., and Angus, D. C. (2016). The third international consensus definitions for sepsis and septic shock (Sepsis-3). *Journal of the American Medical Association*, 315(8):801–810.

Strimbu, K. and Tavel, J. A. (2010). What are biomarkers? *Current opinion in HIV and AIDS*, 5(6):463–466.

Sun, L., Song, X., and Zhang, Z. (2012). Mean residual life models with time-dependent coefficients under right censoring. *Biometrika*, 99(1):185–197.

Sun, L. and Zhang, Z. (2009). A class of transformed mean residual life models with censored survival data. *Journal of the American Statistical Association*, 104(486):803–815.

Therneau, T. M. and Grambsch, P. M. (2000). *Modeling Survival Data: Extending the Cox Model*. Springer, New York, NY.

Tsiatis, A. A. and Davidian, M. (2004). Joint modeling of longitudinal and time-to-event data: An overview. In *Statistica Sinica*, volume 14, pages 809–834. Institute of Statistical Science, Academia Sinica.

Tsiatis, A. A., Davidian, M., Holloway, S. T., and Laber, E. B. (2020). *Dynamic Treatment Regimes: Statistical Methods for Precision Medicine*. Chapman and Hall/CRC, Boca Raton, FL, 1st edition.

van Houwelingen, H. and Putter, H. (2011). *Dynamic Prediction in Clinical Survival Analysis*. CRC Press, Boca Raton, FL, 1st edition.

van Houwelingen, H. C. (2007). Dynamic prediction by landmarking in event history analysis. *Scandinavian Journal of Statistics*, 34(1):70–85.

Werbos, P. J. (1990). Backpropagation through time: What it does and how to do it. *Proceedings of the IEEE*, 78(10):1550–1560.

Yao, F., Müller, H. G., and Wang, J. L. (2005). Functional data analysis for sparse longitudinal data. *Journal of the American Statistical Association*, 100(470):577–590.

Yuan, M. and Lin, Y. (2006). Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 68(1):49–67.

Zhang, B., Tsiatis, A. A., Laber, E. B., and Davidian, M. (2013). Robust estimation of optimal dynamic treatment regimes for sequential treatment decisions. *Biometrika*, 100(3):681–694.

Zhao, Y. Q., Zeng, D., Laber, E. B., and Kosorok, M. R. (2015a). New statistical learning methods for estimating optimal dynamic treatment regimes. *Journal of the American Statistical Association*, 110(510):583–598.

Zhao, Y. Q., Zeng, D., Laber, E. B., Song, R., Yuan, M., and Kosorok, M. R. (2015b). Doubly robust learning for estimating individualized treatment with censored data. *Biometrika*, 102(1):151–168.

Zheng, Y. and Heagerty, P. J. (2005). Partly conditional survival models for longitudinal data. *Biometrics*, 61(2):379–391.

Zhu, Y., Li, L., and Huang, X. (2019). Landmark linear transformation model for dynamic prediction with application to a longitudinal cohort study of chronic disease. *Journal of the Royal Statistical Society. Series C: Applied Statistics*, 68(3):771–791.

APPENDICES

APPENDIX

A

CHAPTER 2 SUPPLEMENTARY MATERIAL

A.1 LSTM-GLM Automated Hyperparameter Selection

We introduce an automated selection process for two important LSTM-GLM hyperparameters: the dimension of the context vectors, s, and the number of LSTM autoencoder training epochs, ep. The process is outlined below at a given prediction time τ.

1. Specify a set of D candidate hyperparameter settings $\mathscr{D} = \{(s_1, ep_1), \ldots, (s_D, ep_D)\}$.
2. For each $d = 1, \ldots, D$:
   (a) For each of the $k = 1, \ldots, p$ biomarkers:
      i. Train an LSTM autoencoder for ep_d epochs to construct s_d-dimensional, window-specific context vectors $\boldsymbol{\psi}_{ik}^{d}(\tau)$ for all i such that $Y_i > \tau$.
   (b) Define $\boldsymbol{\psi}_{i}^{d}(\tau) = \{\boldsymbol{\psi}_{i1}^{d}(\tau), \ldots, \boldsymbol{\psi}_{ip}^{d}(\tau)\}$.
3. Using R unique divisions of the data:
   (a) Divide the data into a training data set containing 100ω percent of at-risk patients and a testing data set containing the other $100(1-\omega)$ percent of at-risk patients, where $\omega \in (0,1)$. Stratify on the censoring indicator Δ_i.

(b) For each $d = 1, \ldots, D$:

i. Fit the LSTM-GLM on the baseline covariates $\mathbf{X}_i$ and the window-specific context vectors $\boldsymbol{\psi}_i^d(\tau)$ using patients i in the training data set.

ii. Compute the loss defined in Equation (A.1) on patients i in the testing data set.

$$\frac{1}{\sum_i \frac{\Delta_i I(Y_i>\tau)}{\widehat{G}(Y_i)}} \sum_{i=1}^{m} \frac{\Delta_i I(Y_i>\tau)}{\widehat{G}(Y_i)} \left[(Y_i-\tau) - g\{\eta(\tau)+\boldsymbol{\gamma}(\tau)^T\mathbf{X}_i+\boldsymbol{\alpha}(\tau)^T\boldsymbol{\psi}_i^d(\tau)\}\right]^2 \tag{A.1}$$

4. For each $d = 1, \ldots, D$:

(a) Calculate the median of the R testing losses computed in Step 3b. Denote the median testing loss for hyperparameter setting d as $m(d)$.

5. Define $d^{opt} = \text{argmin}_d\{m(d)\}$. Select $s = s_{d^{opt}}$ and $ep = ep_{d^{opt}}$.

The selection of R and ω should be guided by the size of the data set under study and the computational resources available. Note, the automated hyperparameter selection process must be repeated at each prediction time $\tau \in \mathscr{T}$.

A.2 Comparative Methods for Performance Evaluation

We dynamically predict MRL using the LSTM-GLM, the LSTM-NN, and six variations of the dynamic transformed MRL model, and we compare the prediction performance of the competing methods. For each of the six dynamic transformed MRL models, we define a distinct function of the history of longitudinal biomarker measurements, $\boldsymbol{\zeta}_i(\tau) = f\{\mathbf{Z}_i(t_{i1}), ..., \mathbf{Z}_i(t_{i\tau_i})\}$.

First, we specify $\boldsymbol{\zeta}_i^{(B)}(\tau) = \mathbf{Z}_i(t_{i1})$ to be the p-dimensional vector of baseline biomarker measurements collected on patient i. Second, we specify $\boldsymbol{\zeta}_i^{(L)}(\tau) = \mathbf{Z}_i(t_{i\tau_i})$ to be the p-dimensional vector of biomarker measurements collected most recently before prediction time τ on patient i. We refer to $\boldsymbol{\zeta}_i^{(L)}(\tau)$ as the "last-value carried forward" vector. Third, let

$$Z_{ik}^{avg}(\tau) = \left\{\sum_{j=1}^{n_i} I(t_{ij}<\tau)\right\}^{-1} \sum_{j=1}^{n_i} I(t_{ij}<\tau) Z_{ik}(t_{ij})$$

be the average value of biomarker k prior to time τ for patient i. We define the p-dimensional vector $\boldsymbol{\zeta}_i^{(A)}(\tau) = \{Z_{i1}^{avg}(\tau), \ldots, Z_{ip}^{avg}(\tau)\}$.

Next, we construct two formulations of $\boldsymbol{\zeta}_i(\tau)$ that contain the intercept and slope of each biomarker regressed against time. To maintain the dynamic nature of prediction, at each prediction time τ, we conduct regression using only biomarker measurements taken at times $t_{ij} < \tau$ on patients with $Y_i > \tau$. Specifically, for each biomarker $k = 1, \ldots, p$, at each prediction time τ, we fit the linear model

$$Z_{ik}(t_{ij}) = \beta_{0ik}^{(\tau)} + \beta_{1ik}^{(\tau)} t_{ij} + \epsilon_{ijk}^{(\tau)}, \tag{A.2}$$

where $t_{ij} \in [0, \tau)$, $\beta_{0ik}^{(\tau)}$ and $\beta_{1ik}^{(\tau)}$ are scalar parameters, and $\epsilon_{ijk}^{(\tau)}$ is a scalar error term.

First, we fit Equation (A.2) independently for each patient i and each biomarker k at each prediction time τ via linear regression. We assume $\epsilon_{ijk}^{(\tau)} \overset{ind}{\sim} \mathbb{N}(0, (\sigma_{ik}^{(\tau)})^2)$, where $\mathbb{N}(\mu, \Sigma)$ denotes a normal distribution with mean μ and variance-covariance Σ.

Second, we frame Equation (A.2) as a linear mixed effects model by defining

$$\beta_{0ik}^{(\tau)} = \eta_{0k}^{(\tau)} + b_{0ik}^{(\tau)}, \qquad \beta_{1ik}^{(\tau)} = \eta_{1k}^{(\tau)} + b_{1ik}^{(\tau)},$$

where $\eta_{0k}^{(\tau)}$ and $\eta_{1k}^{(\tau)}$ are scalar fixed parameters, and $b_{0ik}^{(\tau)}$ and $b_{1ik}^{(\tau)}$ are scalar random effects. Let $\mathbf{b}_{ik}^{(\tau)} = (b_{0ik}^{(\tau)}, b_{1ik}^{(\tau)})$. We assume $\mathbf{b}_{ik}^{(\tau)} \sim \mathbb{N}(\mathbf{0}, \mathbf{D}_k^{(\tau)})$ and $\epsilon_{ijk}^{(\tau)} \overset{ind}{\sim} \mathbb{N}(0, (\sigma_k^{(\tau)})^2)$. Additionally, we assume the random effects $\mathbf{b}_{ik}^{(\tau)}$ and the errors $\epsilon_{ijk}^{(\tau)}$ are independent. We fit the linear mixed effects model independently for each biomarker k at each prediction time τ based on the data for all patients with $Y_i > \tau$ via maximum likelihood methods.

We then specify the "linear regression" vector $\boldsymbol{\zeta}_i^{(S)}(\tau)$ and the "mixed effects" vector $\boldsymbol{\zeta}_i^{(M)}(\tau)$ to be the $2p$-dimensional vector $(\widehat{\beta}_{0i1}^{(\tau)}, \widehat{\beta}_{1i1}^{(\tau)}, \ldots, \widehat{\beta}_{0ip}^{(\tau)}, \widehat{\beta}_{1ip}^{(\tau)})$, where the parameter estimates are obtained via the aforementioned methods.

Lastly, motivated by the work of Lin et al. (2018), we construct the sixth formulation of $\boldsymbol{\zeta}_i(\tau)$ using FPCA. At each prediction time τ, we approximate the measurement of biomarker k for patient i at time $t_{ij} \in [0, \tau)$ as

$$Z_{ik}(t_{ij}) \approx \mu_k^{(\tau)}(t_{ij}) + \sum_{l=1}^{L_k^{(\tau)}} \lambda_{ilk}^{(\tau)} \cdot \rho_{lk}^{(\tau)}(t_{ij}) + \epsilon_{ijk}^{(\tau)},$$

where $\mu_k^{(\tau)}(\cdot)$ is the mean function, $\lambda_{ilk}^{(\tau)}$ is the lth FPC score, $\rho_{lk}^{(\tau)}(\cdot)$ is the lth eigenfunction, and $\epsilon_{ijk}^{(\tau)} \overset{ind}{\sim} \mathbb{N}(0, (\sigma_k^{(\tau)})^2)$.

Define $L_k^{(\tau)}$ to be the minimum number of FPC scores required to explain 99 percent

of the total variance of biomarker k with respect to prediction time τ. We then specify $\zeta_i^{(F)}(\tau) = (\lambda_{i11}^{(\tau)}, \ldots, \lambda_{iL_1^{(\tau)}1}^{(\tau)}, \ldots, \lambda_{i1p}^{(\tau)}, \ldots, \lambda_{iL_p^{(\tau)}p}^{(\tau)})$ to be a vector of the $L_k^{(\tau)}$ FPC scores for each of the $k = 1, \ldots, p$ biomarkers. Then the dimension of $\zeta_i^{(F)}(\tau)$ is $\sum_{k=1}^{p} L_k^{(\tau)}$.

We estimate $\mu_k^{(\tau)}(\cdot)$, $\rho_{lk}^{(\tau)}(\cdot)$, and $\lambda_{ilk}^{(\tau)}$ independently for each biomarker k at each prediction time τ using the principal components analysis through conditional estimation (PACE) algorithm (Yao et al. (2005)). To uphold the dynamic nature of prediction, we conduct estimation at prediction time τ using only biomarker measurements collected at times $t_{ij} < \tau$ on patients with $Y_i > \tau$.

A.3 Simulated Survival Times

We conduct two simulation studies to assess the prediction performance of the LSTM-GLM and the LSTM-NN. In each study, we generate the survival times T_i using a distinct model.

In the first study, we generate T_i according to the accelerated failure time (AFT) model

$$\nu_i = \int_0^{T_i} \exp\{\beta_1 B_i(s) + \beta_2 X_i\} ds,$$

where $\beta_1 = \beta_2 = 1$. For each patient, we generate a random $\nu_i = \exp(\theta_i)$, where $\theta_i \overset{iid}{\sim} \mathbb{N}(3,1)$.

In the second study, we generate T_i according to the Cox proportional hazards model

$$h_i\{t \mid B_i(t), X_i\} = \lambda \exp\{\beta_1 B_i(t) + \beta_2 X_i\},$$

where $\beta_1 = \beta_2 = 1$ and $\lambda = 0.05$. For each patient, we generate a random $\nu_i = -log(1-\theta_i)$, where $\theta_i \overset{iid}{\sim} \mathbb{U}(0,1)$ and $\mathbb{U}$ represents the uniform distribution.

Define $\lambda = 1$ for the AFT model. Then for both the AFT and Cox models,

$$\nu_i = \int_0^{T_i} \lambda \exp\{\beta_1 B_i(s) + \beta_2 X_i\} ds.$$

Define $K(t) = 1 + \sum_{j=1}^{8} j\, I(j < t \le j+1) + 9I(t > 9)$. Then

$$\begin{aligned} \nu_i(t) &= \int_0^t \lambda \exp\{\beta_1 B_i(s) + \beta_2 X_i\} ds \\ &= \nu_i^*\{t; K(t)\} + I\{K(t) > 1\} \sum_{j=1}^{K(t)-1} \nu_i^*(j; j), \end{aligned}$$

where

$$\begin{aligned} \nu_i^*(t; K) &= \int_{K-1}^t \lambda \exp\{\beta_1 B_i(s) + \beta_2 X_i\} ds \\ &= \lambda \exp(\beta_2 X_i) \int_{K-1}^t \exp\left(\beta_1 \left[a + b_{i0} + \sum_{j=1}^{K} (c_j + b_{ij})\{s - (j-1)\} \right] \right) ds \\ &= \lambda \exp\left\{ \beta_2 X_i + \beta_1 (a + b_{i0}) - \beta_1 \sum_{j=1}^{K} (c_j + b_{ij})(j-1) \right\} \int_{K-1}^t \exp\left[s\beta_1 \left\{ \sum_{j=1}^{K} (c_j + b_{ij}) \right\} \right] ds \\ &= \frac{\lambda \exp\left\{ \beta_2 X_i + \beta_1 (a + b_{i0}) - \beta_1 \sum_{j=1}^{K} (c_j + b_{ij})(j-1) \right\}}{\beta_1 \sum_{j=1}^{K} (c_j + b_{ij})} \times \\ &\qquad \left[\exp\left\{ t\beta_1 \sum_{j=1}^{K} (c_j + b_{ij}) \right\} - \exp\left\{ (K-1)\beta_1 \sum_{j=1}^{K} (c_j + b_{ij}) \right\} \right]. \end{aligned}$$

We invert $\nu_i(t)$. Define

$$R\{\nu_i(t)\} = 1 + \sum_{j=1}^{8} j\, I\{\nu_i(j) < \nu_i(t) \le \nu_i(j+1)\} + 9I\{\nu_i(t) > \nu_i(9)\}.$$

To simplify notation, we suppress the dependence of R on $\nu_i(t)$. Let $\mathscr{G} = \beta_1 \sum_{j=1}^{R} (c_j + b_{ij})$. Then

$$\nu_i^{-1}(t) = \log \left[\frac{\left\{ \nu_i(t) - I(R > 1) \sum_{j=1}^{R-1} \nu_i^*(j; j) \right\} \mathscr{G}}{\lambda \exp\left\{ \beta_2 X_i + \beta_1 (a + b_{i0}) - \beta_1 \sum_{j=1}^{R} (c_j + b_{ij})(j-1) \right\}} + \exp\{(R-1)\mathscr{G}\} \right] \Big/ \mathscr{G}.$$

For each simulation, we compute the survival time $T_i = \nu_i^{-1}(t)$ using each patient's randomly generated ν_i. We administratively censor patients whose randomly generated ν_i is not in the range of $\int_0^{T_i} \lambda \exp\{\beta_1 B_i(s) + \beta_2 X_i\} ds$.

A.4 Supplemental Simulation Studies

We repeat the simulation studies described in Chapter 2, Section 2.4. To simulate a longitudinal biomarker measured using a precise instrument, we reduce both the measurement error and the variation in measurement times. Specifically, we generate the longitudinal biomarker $Z_i(t_{ij}) = B_i(t_{ij}) + \epsilon_{ij}$ at the patient-specific measurement times $t_{ij} = \min(0, \tau_j + \varepsilon_{ij})$, where $\epsilon_{ij} \overset{iid}{\sim} \mathbb{N}(0, 0.05^2)$ and $\varepsilon_{ij} \overset{iid}{\sim} \mathbb{N}(0, 0.01^2)$. With the reduced measurement error, we are less concerned with overfitting the LSTM-NN, so we train the network for 25,000 epochs. We define all other simulation settings to be identical to those described in Chapter 2, Section 2.4.

We plot the distributions of the 500 testing losses and 500 testing C-indexes for each of the eight studied dynamic MRL models in Figure (A.1). For both the AFT and Cox simulations, the LSTM-NN results in the lowest median testing loss, followed by the LSTM-GLM. The LSTM-NN and LSTM-GLM also result in the highest median testing C-indexes. The FPCA model consistently results in the third-best median testing loss and testing C-index. Compared to the simulation studies presented in Chapter 2, Section 2.4, these results demonstrate a more significant improvement in calibration and discrimination for the LSTM-GLM and the LSTM-NN relative to competing methods. Thus, these results suggest that the LSTM-GLM and the LSTM-NN are especially useful for producing accurate dynamic predictions of MRL in settings where the longitudinal biomarkers are measured using precise instruments.

A.5 MIMIC-III Data Application Hyperparameter Selection

In the MIMIC-III data application, we must specify the hyperparameter settings of the LSTM-GLM and the LSTM-NN. The LSTM autoencoders used to construct the window-specific context vectors for the LSTM-GLM and the LSTM-NN have two important hyperparameters: the dimension of the window-specific context vectors, s, and the number of training epochs, ep_a. As s increases, the total number of parameters in the autoencoder increases, so more epochs are needed to train the autoencoder. Accordingly, at each prediction time $\tau \in \mathcal{T}$, we construct four sets of window-specific context vectors using the hyperparameter settings $(s, ep_a) \in \{(3,150),(5,150),(5,300),(7,300)\}$.

We fit four LSTM-GLMs that each regress on one of the four sets of context vectors. The distribution of 100 testing losses for each LSTM-GLM is plotted at each prediction time in

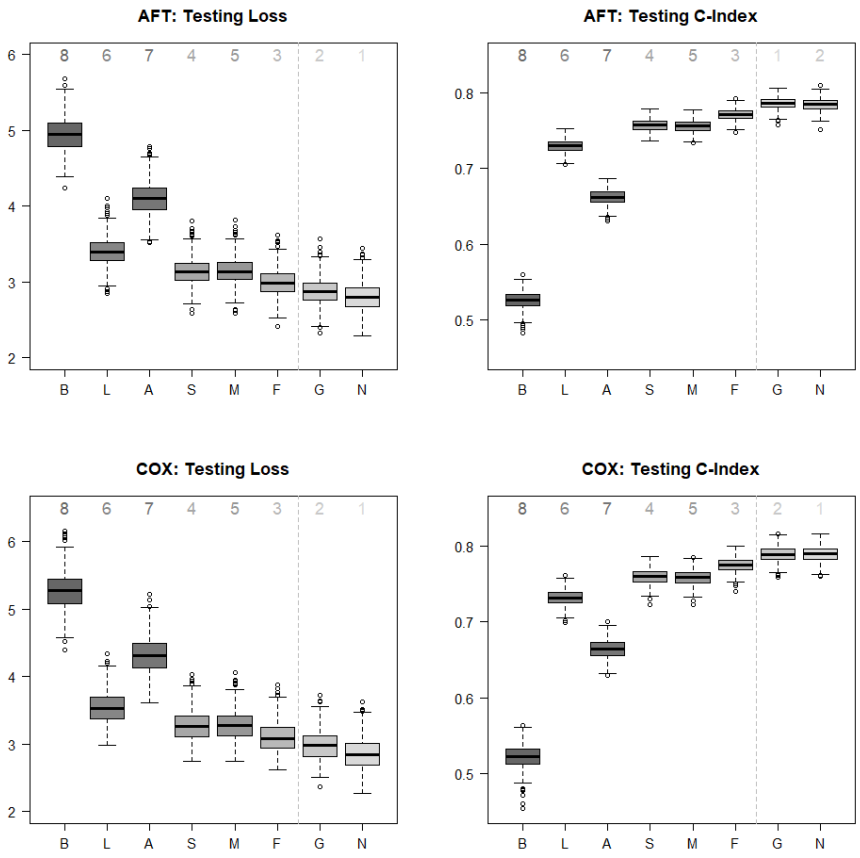

Figure A.1: Distribution of the 500 testing losses and 500 testing C-indexes for each of the eight dynamic prediction models. The models resulting in the lowest median testing loss and the highest median testing C-index are labelled 1. The models resulting in the highest median testing loss and the lowest median testing C-index are labelled 8. The six dynamic transformed MRL models are labelled according to their formulation of $\zeta_i(\tau)$. "B" represents the baseline vector. "L" represents the last-value carried forward vector. "A" represents the average vector. "S" represents the linear regression vector. "M" represents the mixed effects vector. "F" represents the FPCA vector. Furthermore, "G" represents the LSTM-GLM, and "N" represents the LSTM-NN.

Figure (A.2). For each $\tau \in \mathscr{T}$, we define the LSTM-GLM that results in the lowest median testing loss to be the "best" LSTM-GLM at time τ. The hyperparameter settings of the best LSTM-GLM at each $\tau \in \mathscr{T}$ can be seen in Table (2.1).

Additionally, we fit an "automated" LSTM-GLM. For each unique data division in the performance evaluation, we select the hyperparameter settings of the automated LSTM-GLM by conducting the automated hyperparameter selection process detailed in Section A.1 on the training data. We define the set of candidate hyperparameter settings to be $\mathscr{D} = \{(3,150),(5,150),(5,300),(7,300)\}$. We set the number of iterations to $R = 50$, and we set the proportion of patients in the sub-training data set to $\omega = 0.5$.

The LSTM-NN feed-forward neural network has three additional hyperparameters that influence prediction performance: the L_2-penalty tuning parameter, λ, the dimension of the parameter matrices, u, and the number of training epochs, ep_n. We train eight LSTM-NNs using all eight possible combinations of $\lambda \in \{0.005, 0.01\}$, $u \in \{1,2\}$, and $ep_n \in \{2000, 3000\}$.

It would be computationally expensive to tune λ, u, and ep_n with respect to all four sets of window-specific context vectors. Consequently, we train the LSTM-NNs using only the set of context vectors constructed with hyperparameter settings $(7,300)$. Because these context vectors have the largest dimension, they have the potential to retain the most information from the biomarker trajectories. Moreover, feed-forward neural networks like the LSTM-NN are capable of handling a large number of covariates.

The distribution of 100 testing losses for each LSTM-NN is plotted at each prediction time in Figure (A.3). For each $\tau \in \mathscr{T}$, we define the LSTM-NN that results in the lowest median testing loss to be the "best" LSTM-NN at time τ. The hyperparameter settings of the best LSTM-NN at each $\tau \in \mathscr{T}$ can be seen in Table (2.1).

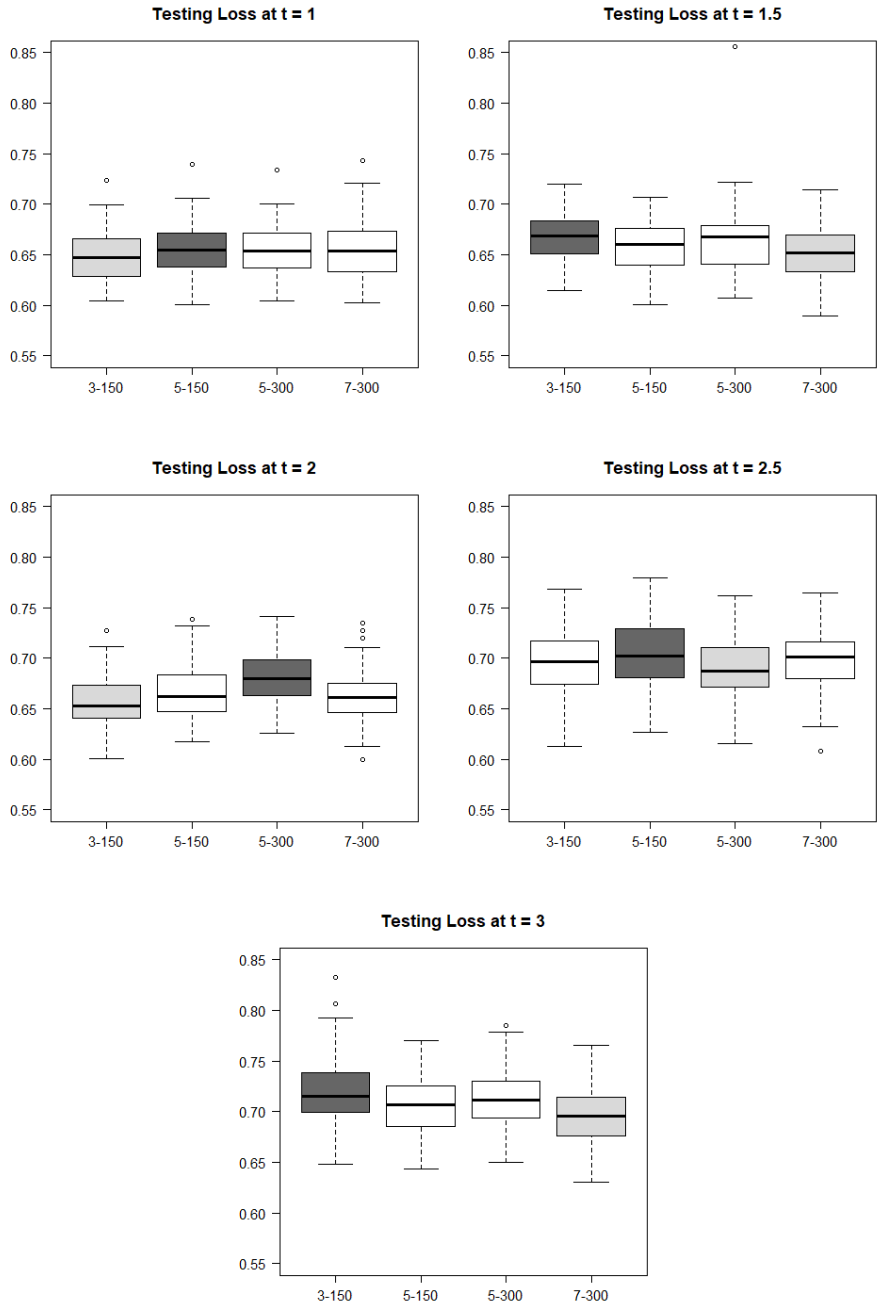

Figure A.2: Distribution of the 100 testing losses for each of the four LSTM-GLMs fit using a different set of window-specific context vectors at prediction times $\tau \in \mathscr{T} = \{1, 1.5, 2, 2.5, 3\}$. Let s be the dimension of the window-specific context vectors, and let ep_a be the number of epochs used to train the LSTM autoencoder. The x-axis is labelled with the hyperparameter settings "$s - ep_a$."

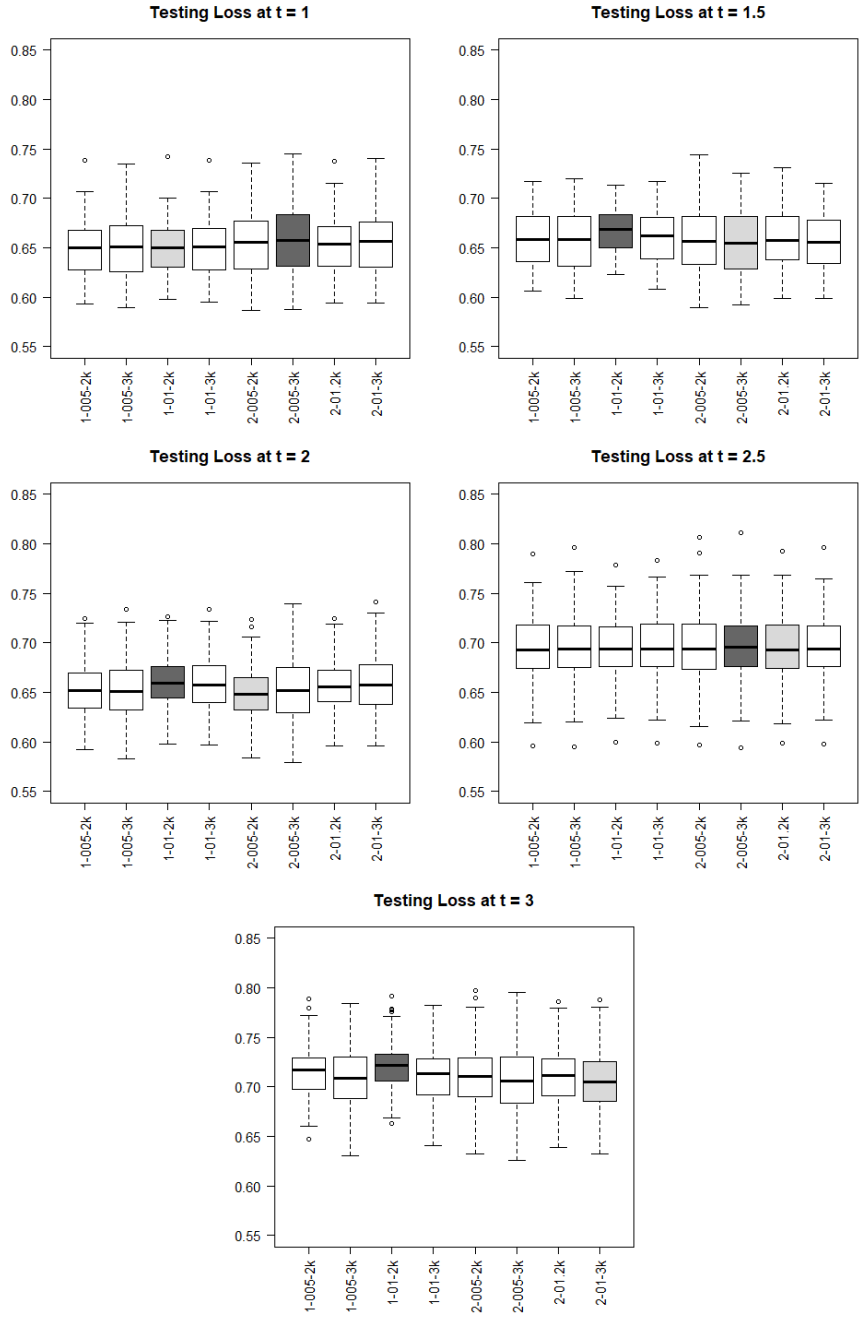

Figure A.3: Distribution of the 100 testing losses for each of the eight LSTM-NNs trained using different hyperparameter settings at prediction times $\tau \in \mathscr{T} = \{1, 1.5, 2, 2.5, 3\}$. Let λ be the L_2-penalty tuning parameter in the LSTM-NN objective function. Let u be the dimension of the LSTM-NN parameter matrices $\mathbf{W}_1$ and $\mathbf{W}_2$. Let ep_n be the number of epochs used to train the LSTM-NN. The x-axis is labelled with the hyperparameter settings "$u - \lambda - ep_n$."

APPENDIX

B

CHAPTER 3 SUPPLEMENTARY MATERIAL

B.1 Proof of Censoring Function

We show that $\mathscr{K}_k(x|\mathbf{h}_k, a_k) = P(C > x)/P(C > \tau_k)$ for $x > \tau_k$, $k = 1, \ldots, K$.

$$
\begin{aligned}
\mathscr{K}_k(x|\mathbf{h}_k, a_k) &= P(C > x \mid U^L > \tau_k, \mathbf{H}_k = \mathbf{h}_k, A_k = a_k) \\
&= P(C > x \mid \min\{T^*(\overline{a}_k), L, C\} > \tau_k, \mathbf{H}_k = \mathbf{h}_k, A_k = a_k) \text{ by SUTVA} \\
&= P(C > x \mid T^*(\overline{a}_k) > \tau_k, L > \tau_k, C > \tau_k, \mathbf{H}_k = \mathbf{h}_k, A_k = a_k) \\
&= P(C > x \mid C > \tau_k) \text{ by independent censoring assumptions} \\
&= \frac{P(C > x \cap C > \tau_k)}{P(C > \tau_k)} \\
&= \frac{P(C > x)}{P(C > \tau_k)} \text{ by } x > \tau_k
\end{aligned}
$$

B.2 Proof of Unbiased Value Estimate

Using mathematical induction, we prove

$$\mathcal{V}(d) = E\left\{\sum_{j=1}^{K} I[\min\{T^*(d), L\} > \tau_j] \cdot \left[\min\{T^*(d), L\} - \tau_j\right]\right\} = E\{V_1^d(\mathbf{h}_1)\}. \tag{B.1}$$

We note that

$$\mathcal{V}(d) = E\left[E\left\{\sum_{j=1}^{K} I[\min\{T^*(d), L\} > \tau_j] \cdot \left[\min\{T^*(d), L\} - \tau_j\right] \,\middle|\, U^L > \tau_1,\ \mathbf{H}_1 = \mathbf{h}_1\right\}\right].$$

For $k = 1, \ldots, K$, we prove the induction hypothesis

$$V_k^d(\mathbf{h}_k) = E\left\{\sum_{j=k}^{K} I\left[\min\{T^*(\overline{a}_{k-1}, d_k, \ldots, d_K), L\} > \tau_j\right] \cdot \left[\min\{T^*(\overline{a}_{k-1}, d_k, \ldots, d_K), L\} - \tau_j\right] \,\middle|\, U^L > \tau_k,\ \mathbf{H}_k = \mathbf{h}_k\right\}.$$

First, we show that the induction hypothesis holds at Decision K.

$$
\begin{aligned}
V_K^d(\mathbf{h}_K) &= E\left[\frac{\Delta^L\{\min(U,L)-\tau_K\}}{\mathscr{K}_K\{\min(U,L)|\mathbf{h}_K,d_K\}} \;\middle|\; U^L>\tau_K, \mathbf{H}_K=\mathbf{h}_K, A_K=d_K(\mathbf{h}_K)\right] \\
&= E\left\{ \frac{\Delta^L[\min\{T^*(\overline{a}_{K-1},d_K),L\}-\tau_K]}{\mathscr{K}_K[\min\{T^*(\overline{a}_{K-1},d_K),L\}\mid\mathbf{h}_K,d_K]} \;\middle|\; U^L>\tau_K, \mathbf{H}_K=\mathbf{h}_K, A_K=d_K(\mathbf{h}_K)\right\} \text{ by SUTVA} \\
&= E\left[E\left\{ \frac{\Delta^L[\min\{T^*(\overline{a}_{K-1},d_K),L\}-\tau_K]}{\mathscr{K}_K[\min\{T^*(\overline{a}_{K-1},d_K),L\}\mid\mathbf{h}_K,d_K]} \;\middle|\; U^L>\tau_K, \mathbf{H}_K=\mathbf{h}_K, A_K=d_K(\mathbf{h}_K), W^*\right\} \;\middle|\; \right. \\
&\qquad\qquad \left. U^L>\tau_K, \mathbf{H}_K=\mathbf{h}_K, A_K=d_K(\mathbf{h}_K)\right] \\
&= E\left[\frac{[\min\{T^*(\overline{a}_{K-1},d_K),L\}-\tau_K]}{\mathscr{K}_K[\min\{T^*(\overline{a}_{K-1},d_K),L\}\mid\mathbf{h}_K,d_K]} E\left\{\Delta^L \mid U^L>\tau_K, \mathbf{H}_K=\mathbf{h}_K, A_K=d_K(\mathbf{h}_K), W^*\right\} \;\middle|\; \right. \\
&\qquad\qquad \left. U^L>\tau_K, \mathbf{H}_K=\mathbf{h}_K, A_K=d_K(\mathbf{h}_K)\right] \\
&= E\left[\min\{T^*(\overline{a}_{K-1},d_K),L\}-\tau_K \mid U^L>\tau_K, \mathbf{H}_K=\mathbf{h}_K, A_K=d_K(\mathbf{h}_K)\right] \text{ by Aside 1} \\
&= E\left[I\{\min(T,L)>\tau_K\}\cdot[\min\{T^*(\overline{a}_{K-1},d_K),L\}-\tau_K] \mid U^L>\tau_K, \mathbf{H}_K=\mathbf{h}_K, A_K=d_K(\mathbf{h}_K)\right] \\
&= E\left[I[\min\{T^*(\overline{a}_{K-1},d_K),L\}>\tau_K]\cdot[\min\{T^*(\overline{a}_{K-1},d_K),L\}-\tau_K] \mid U^L>\tau_K, \mathbf{H}_K=\mathbf{h}_K, A_K=d_K(\mathbf{h}_K)\right] \\
&\qquad\qquad \text{by SUTVA} \\
&= E\left[I[\min\{T^*(\overline{a}_{K-1},d_K),L\}>\tau_K]\cdot[\min\{T^*(\overline{a}_{K-1},d_K),L\}-\tau_K] \mid U^L>\tau_K, \mathbf{H}_K=\mathbf{h}_K\right] \text{ by SRA}
\end{aligned}
$$

Aside 1:

$$
\begin{aligned}
&E\left\{\Delta^L \mid U^L>\tau_K, \mathbf{H}_K=\mathbf{h}_K, A_K=d_K(\mathbf{h}_K), W^*\right\} \\
&\qquad = P\{C>\min(T,L) \mid U^L>\tau_K, \mathbf{H}_K=\mathbf{h}_K, A_K=d_K(\mathbf{h}_K), W^*\} \\
&\qquad = P[C>\min\{T^*(\overline{a}_{K-1},d_K),L\} \mid U^L>\tau_K, \mathbf{H}_K=\mathbf{h}_K, A_K=d_K(\mathbf{h}_K), W^*] \text{ by SUTVA} \\
&\qquad = P[C>\min\{T^*(\overline{a}_{K-1},d_K),L\} \mid U^L>\tau_K, \mathbf{H}_K=\mathbf{h}_K, A_K=d_K(\mathbf{h}_K)] \text{ by independent censoring} \\
&\qquad = \mathscr{K}_K[\min\{T^*(\overline{a}_{K-1},d_K),L\} \mid \mathbf{h}_K,d_K]
\end{aligned}
$$

Assume the induction hypothesis holds at $k+1$. That is, assume

$$
V_{k+1}^d(\mathbf{h}_{k+1}) = E\left\{ \sum_{j=k+1}^{K} I[\min\{T^*(\overline{a}_k,d_{k+1},\ldots,d_K),L\}>\tau_j]\cdot\left[\min\{T^*(\overline{a}_k,d_{k+1},\ldots,d_K),L\}-\tau_j\right] \;\middle|\; U^L>\tau_{k+1},\ \mathbf{H}_{k+1}=\mathbf{h}_{k+1}\right\}.
$$

We show the induction hypothesis holds at Decision k. Recall

$$V_k^d(\mathbf{h}_k) = E\left\{ \frac{\Delta^L(U^L-\tau_k)}{\mathscr{K}_k(U^L|\mathbf{h}_k,d_k)} + \frac{I(U^L>\tau_{k+1})V_{k+1}^d(\mathbf{H}_{k+1})}{\mathscr{K}_k(\tau_{k+1}|\mathbf{h}_k,d_k)} \middle| U^L>\tau_k, \mathbf{H}_k=\mathbf{h}_k, A_k=d_k(\mathbf{h}_k)\right\} \tag{B.2}$$

$$= E\left\{ \frac{\Delta^L(U^L-\tau_k)}{\mathscr{K}_k(U^L|\mathbf{h}_k,d_k)} \middle| U^L>\tau_k, \mathbf{H}_k=\mathbf{h}_k, A_k=d_k(\mathbf{h}_k)\right\} \tag{B.3}$$

$$+ E\left\{ \frac{I(U^L>\tau_{k+1})V_{k+1}^d(\mathbf{H}_{k+1})}{\mathscr{K}_k(\tau_{k+1}|\mathbf{h}_k,d_k)} \middle| U^L>\tau_k, \mathbf{H}_k=\mathbf{h}_k, A_k=d_k(\mathbf{h}_k)\right\}. \tag{B.4}$$

Examine Equation (B.3).

$$E\left[\frac{\Delta^L\{\min(U,L)-\tau_k\}}{\mathscr{K}_k\{\min(U,L)|\mathbf{h}_k,d_k\}} \middle| U^L>\tau_k, \mathbf{H}_k=\mathbf{h}_k, A_k=d_k(\mathbf{h}_k)\right]$$

$$= E\left\{ \frac{\Delta^L[\min\{T^*(\overline{a}_{k-1},d_k,\ldots,d_K),L\}-\tau_k]}{\mathscr{K}_k[\min\{T^*(\overline{a}_{k-1},d_k,\ldots,d_K),L\}\mid \mathbf{h}_k,d_k]} \middle| U^L>\tau_k, \mathbf{H}_k=\mathbf{h}_k, A_k=d_k(\mathbf{h}_k)\right\} \text{ by SUTVA}$$

$$= E\left[E\left\{ \frac{\Delta^L[\min\{T^*(\overline{a}_{k-1},d_k,\ldots,d_K),L\}-\tau_k]}{\mathscr{K}_k[\min\{T^*(\overline{a}_{k-1},d_k,\ldots,d_K),L\}\mid \mathbf{h}_k,d_k]} \middle| U^L>\tau_k, \mathbf{H}_k=\mathbf{h}_k, A_k=d_k(\mathbf{h}_k), W^*\right\} \middle| \right.$$

$$\left. U^L>\tau_k, \mathbf{H}_k=\mathbf{h}_k, A_k=d_k(\mathbf{h}_k)\right]$$

$$= E\left[\frac{[\min\{T^*(\overline{a}_{k-1},d_k,\ldots,d_K),L\}-\tau_k]}{\mathscr{K}_k[\min\{T^*(\overline{a}_{k-1},d_k,\ldots,d_K),L\}\mid \mathbf{h}_k,d_k]} E\left\{\Delta^L \mid U^L>\tau_k, \mathbf{H}_k=\mathbf{h}_k, A_k=d_k(\mathbf{h}_k), W^*\right\} \middle| \right.$$

$$\left. U^L>\tau_k, \mathbf{H}_k=\mathbf{h}_k, A_k=d_k(\mathbf{h}_k)\right]$$

$$= E\left[\min\{T^*(\overline{a}_{k-1},d_k,\ldots,d_K),L\}-\tau_k \mid U^L>\tau_k, \mathbf{H}_k=\mathbf{h}_k, A_k=d_k(\mathbf{h}_k)\right] \text{ by Aside 2}$$

$$= E\left\{ I\{\min(T,L)>\tau_k\}\cdot[\min\{T^*(\overline{a}_{k-1},d_k,\ldots,d_K),L\}-\tau_k] \mid U^L>\tau_k, \mathbf{H}_k=\mathbf{h}_k, A_k=d_k(\mathbf{h}_k)\right\}$$

$$= E\left\{ I[\min\{T^*(\overline{a}_{k-1},d_k,\ldots,d_K),L\}>\tau_k]\cdot[\min\{T^*(\overline{a}_{k-1},d_k,\ldots,d_K),L\}-\tau_k] \mid U^L>\tau_k, \mathbf{H}_k=\mathbf{h}_k, A_k=d_k(\mathbf{h}_k)\right\}$$

$$\text{by SUTVA}$$

$$= E\left\{ I[\min\{T^*(\overline{a}_{k-1},d_k,\ldots,d_K),L\}>\tau_k]\cdot[\min\{T^*(\overline{a}_{k-1},d_k,\ldots,d_K),L\}-\tau_k] \mid U^L>\tau_k, \mathbf{H}_k=\mathbf{h}_k\right\} \text{ by SRA}$$

Aside 2:

$$
\begin{aligned}
&E\{\Delta^L \mid U^L > \tau_k, \mathbf{H}_k = \mathbf{h}_k, A_k = d_k(\mathbf{h}_k), W^*\} \\
&\qquad = P\{C > \min(T, L) \mid U^L > \tau_k, \mathbf{H}_k = \mathbf{h}_k, A_k = d_k(\mathbf{h}_k), W^*\} \\
&\qquad = P[C > \min\{T^*(\overline{a}_{k-1}, d_k, \ldots, d_K), L\} \mid U^L > \tau_k, \mathbf{H}_k = \mathbf{h}_k, A_k = d_k(\mathbf{h}_k), W^*] \text{ by SUTVA} \\
&\qquad = P[C > \min\{T^*(\overline{a}_{k-1}, d_k, \ldots, d_K), L\} \mid U^L > \tau_k, \mathbf{H}_k = \mathbf{h}_k, A_k = d_k(\mathbf{h}_k)] \\
&\qquad\qquad \text{by independent censoring} \\
&\qquad = \mathscr{K}_k[\min\{T^*(\overline{a}_{k-1}, d_k, \ldots, d_K), L\} \mid \mathbf{h}_k, d_k]
\end{aligned}
$$

Define

$$
S^*(\overline{a}_k, d_{k+1}, \ldots, d_K) = \sum_{j=k+1}^{K} I[\min\{T^*(\overline{a}_k, d_{k+1}, \ldots, d_K), L\} > \tau_j] \cdot \left[\min\{T^*(\overline{a}_k, d_{k+1}, \ldots, d_K), L\} - \tau_j\right].
$$

Examine Equation (B.4).

$$
\begin{aligned}
&E\left\{ \frac{I(U^L > \tau_{k+1}) V_{k+1}^d(\mathbf{H}_{k+1})}{\mathscr{K}_k(\tau_{k+1} | \mathbf{h}_k, d_k)} \middle| U^L > \tau_k, \mathbf{H}_k = \mathbf{h}_k, A_k = d_k(\mathbf{h}_k) \right\} \\
&= E\left[\frac{I(U^L > \tau_{k+1})}{\mathscr{K}_k(\tau_{k+1} | \mathbf{h}_k, d_k)} E\left\{ S^*(\overline{a}_k, d_{k+1}, \ldots, d_K) \,\middle|\, U^L > \tau_{k+1}, \mathbf{H}_{k+1} = \mathbf{h}_{k+1} \right\} \middle| U^L > \tau_k, \mathbf{H}_k = \mathbf{h}_k, A_k = d_k(\mathbf{h}_k) \right] \\
&= E\left[E\left\{ \frac{I(U^L > \tau_{k+1}) S^*(\overline{a}_k, d_{k+1}, \ldots, d_K)}{\mathscr{K}_k(\tau_{k+1} | \mathbf{h}_k, d_k)} \,\middle|\, \mathbf{H}_{k+1} = \mathbf{h}_{k+1} \right\} \middle| U^L > \tau_k, \mathbf{H}_k = \mathbf{h}_k, A_k = d_k(\mathbf{h}_k) \right] \text{ by Aside 3} \\
&= E\left\{ \frac{I(U^L > \tau_{k+1}) S^*(\overline{a}_k, d_{k+1}, \ldots, d_K)}{\mathscr{K}_k(\tau_{k+1} | \mathbf{h}_k, d_k)} \middle| U^L > \tau_k, \mathbf{H}_k = \mathbf{h}_k, A_k = d_k(\mathbf{h}_k) \right\} \\
&= E\left[E\left\{ \frac{I(U^L > \tau_{k+1}) S^*(\overline{a}_k, d_{k+1}, \ldots, d_K)}{\mathscr{K}_k(\tau_{k+1} | \mathbf{h}_k, d_k)} \middle| U^L > \tau_k, \mathbf{H}_k = \mathbf{h}_k, A_k = d_k(\mathbf{h}_k), W^* \right\} \middle| \right. \\
&\qquad\qquad \left. U^L > \tau_k, \mathbf{H}_k = \mathbf{h}_k, A_k = d_k(\mathbf{h}_k) \right] \\
&= E\left\{ \sum_{j=k+1}^{K} I[\min\{T^*(\overline{a}_{k-1}, d_k, \ldots, d_K), L\} > \tau_j] \cdot \left[\min\{T^*(\overline{a}_{k-1}, d_k, \ldots, d_K), L\} - \tau_j\right] \right| \\
&\qquad\qquad \left. U^L > \tau_k, \mathbf{H}_k = \mathbf{h}_k, A_k = d_k(\mathbf{h}_k) \right\} \text{ by Aside 4} \\
&= E\left\{ \sum_{j=k+1}^{K} I[\min\{T^*(\overline{a}_{k-1}, d_k, \ldots, d_K), L\} > \tau_j] \cdot \left[\min\{T^*(\overline{a}_{k-1}, d_k, \ldots, d_K), L\} - \tau_j\right] \middle| U^L > \tau_k, \mathbf{H}_k = \mathbf{h}_k \right\} \\
&\qquad\qquad \text{by SRA}
\end{aligned}
$$

Aside 3:

$$
\begin{aligned}
&E\left[\, E\left\{\frac{I(U^L>\tau_{k+1})S^*(\overline{a}_k,d_{k+1},\ldots,d_K)}{\mathscr{K}_k(\tau_{k+1}|\mathbf{h}_k,d_k)}\;\middle|\;\mathbf{H}_{k+1}=\mathbf{h}_{k+1}\right\}\middle|\, U^L>\tau_k,\mathbf{H}_k=\mathbf{h}_k,A_k=d_k(\mathbf{h}_k)\right]\\
&=E\Bigg(E\left[\, E\left\{\frac{I(U^L>\tau_{k+1})S^*(\overline{a}_k,d_{k+1},\ldots,d_K)}{\mathscr{K}_k(\tau_{k+1}|\mathbf{h}_k,d_k)}\;\middle|\;U^L>\tau_{k+1},\mathbf{H}_{k+1}=\mathbf{h}_{k+1}\right\}\middle|\,\mathbf{H}_{k+1}=\mathbf{h}_{k+1}\right]\Bigg|\\
&\qquad\qquad U^L>\tau_k,\mathbf{H}_k=\mathbf{h}_k,A_k=d_k(\mathbf{h}_k)\Bigg)\\
&=E\Bigg(E\left[\frac{I(U^L>\tau_{k+1})}{\mathscr{K}_k(\tau_{k+1}|\mathbf{h}_k,d_k)}\;E\left\{S^*(\overline{a}_k,d_{k+1},\ldots,d_K)\mid U^L>\tau_{k+1},\mathbf{H}_{k+1}=\mathbf{h}_{k+1}\right\}\middle|\,\mathbf{H}_{k+1}=\mathbf{h}_{k+1}\right]\Bigg|\\
&\qquad\qquad U^L>\tau_k,\mathbf{H}_k=\mathbf{h}_k,A_k=d_k(\mathbf{h}_k)\Bigg)\\
&=E\left[\frac{I(U^L>\tau_{k+1})}{\mathscr{K}_k(\tau_{k+1}|\mathbf{h}_k,d_k)}E\left\{S^*(\overline{a}_k,d_{k+1},\ldots,d_K)\mid U^L>\tau_{k+1},\mathbf{H}_{k+1}=\mathbf{h}_{k+1}\right\}\middle|\,U^L>\tau_k,\mathbf{H}_k=\mathbf{h}_k,A_k=d_k(\mathbf{h}_k)\right]
\end{aligned}
$$

Aside 4:

$$
\begin{aligned}
&E\left\{\frac{I(U^L>\tau_{k+1})S^*(\overline{a}_k,d_{k+1},\ldots,d_K)}{\mathscr{K}_k(\tau_{k+1}|\mathbf{h}_k,d_k)}\;\middle|\;U^L>\tau_k,\mathbf{H}_k=\mathbf{h}_k,A_k=d_k(\mathbf{h}_k),W^*\right\}\\
&=E\left[\frac{I\{\min(T,L,C)>\tau_{k+1}\}S^*(\overline{a}_k,d_{k+1},\ldots,d_K)}{\mathscr{K}_k(\tau_{k+1}|\mathbf{h}_k,d_k)}\;\middle|\;U^L>\tau_k,\mathbf{H}_k=\mathbf{h}_k,A_k=d_k(\mathbf{h}_k),W^*\right]\\
&=E\Bigg\{\frac{I[\min\{T^*(\overline{a}_{k-1},d_k,\ldots,d_K),L\}>\tau_{k+1}]I(C>\tau_{k+1})}{\mathscr{K}_k(\tau_{k+1}|\mathbf{h}_k,d_k)}\times\\
&\qquad\sum_{j=k+1}^{K}I[\min\{T^*(\overline{a}_k,d_{k+1},\ldots,d_K),L\}>\tau_j]\cdot\left[\min\{T^*(\overline{a}_k,d_{k+1},\ldots,d_K),L\}-\tau_j\right]\Bigg|\\
&\qquad\qquad U^L>\tau_k,\mathbf{H}_k=\mathbf{h}_k,A_k=d_k(\mathbf{h}_k),W^*\Bigg\}\text{ by SUTVA}\\
&=\frac{I[\min\{T^*(\overline{a}_{k-1},d_k,\ldots,d_K),L\}>\tau_{k+1}]}{\mathscr{K}_k(\tau_{k+1}|\mathbf{h}_k,d_k)}\times\\
&\qquad\sum_{j=k+1}^{K}I[\min\{T^*(\overline{a}_{k-1},d_k,\ldots,d_K),L\}>\tau_j]\cdot\left[\min\{T^*(\overline{a}_{k-1},d_k,\ldots,d_K),L\}-\tau_j\right]\times\\
&\qquad\qquad E\left\{I(C>\tau_{k+1})\mid U^L>\tau_k,\mathbf{H}_k=\mathbf{h}_k,A_k=d_k(\mathbf{h}_k),W^*\right\}\\
&=I[\min\{T^*(\overline{a}_{k-1},d_k,\ldots,d_K),L\}>\tau_{k+1}]\times\\
&\qquad\sum_{j=k+1}^{K}I[\min\{T^*(\overline{a}_{k-1},d_k,\ldots,d_K),L\}>\tau_j]\cdot\left[\min\{T^*(\overline{a}_{k-1},d_k,\ldots,d_K),L\}-\tau_j\right]\text{ by Aside 5}\\
&=\sum_{j=k+1}^{K}I[\min\{T^*(\overline{a}_{k-1},d_k,\ldots,d_K),L\}>\tau_j]\cdot\left[\min\{T^*(\overline{a}_{k-1},d_k,\ldots,d_K),L\}-\tau_j\right]
\end{aligned}
$$

Aside 5:

$$
\begin{aligned}
&E\{I(C>\tau_{k+1}) \mid U^L>\tau_k, \mathbf{H}_k=\mathbf{h}_k, A_k=d_k(\mathbf{h}_k), W^*\} \\
&\qquad = P\{C>\tau_{k+1} \mid U^L>\tau_k, \mathbf{H}_k=\mathbf{h}_k, A_k=d_k(\mathbf{h}_k), W^*\} \\
&\qquad = P\{C>\tau_{k+1} \mid U^L>\tau_k, \mathbf{H}_k=\mathbf{h}_k, A_k=d_k(\mathbf{h}_k)\} \text{ by independent censoring} \\
&\qquad = \mathscr{K}_k(\tau_{k+1} \mid \mathbf{h}_k, d_k)
\end{aligned}
$$

It follows that

$$
\begin{aligned}
V_k^d(\mathbf{h}_k) &= E\left\{ I[\min\{T^*(\overline{a}_{k-1}, d_k, \ldots, d_K), L\}>\tau_k] \cdot [\min\{T^*(\overline{a}_{k-1}, d_k, \ldots, d_K), L\}-\tau_k] \mid U^L>\tau_k, \mathbf{H}_k=\mathbf{h}_k \right\} \\
&\quad + E\left\{ \sum_{j=k+1}^{K} I[\min\{T^*(\overline{a}_{k-1}, d_k, \ldots, d_K), L\}>\tau_j] \cdot \left[\min\{T^*(\overline{a}_{k-1}, d_k, \ldots, d_K), L\}-\tau_j\right] \middle| U^L>\tau_k, \mathbf{H}_k=\mathbf{h}_k \right\} \\
&= E\left\{ \sum_{j=k}^{K} I[\min\{T^*(\overline{a}_{k-1}, d_k, \ldots, d_K), L\}>\tau_j] \cdot \left[\min\{T^*(\overline{a}_{k-1}, d_k, \ldots, d_K), L\}-\tau_j\right] \middle| U^L>\tau_k,\ \mathbf{H}_k=\mathbf{h}_k \right\}.
\end{aligned}
$$

Finally, we prove Equation (B.1).

$$
\begin{aligned}
E\{V_1^d(\mathbf{h}_1)\} &= E\left[E\left\{ \sum_{j=1}^{K} I[\min\{T^*(d), L\}>\tau_j] \cdot \left[\min\{T^*(d), L\}-\tau_j\right] \middle| U^L>\tau_1,\ \mathbf{H}_1=\mathbf{h}_1 \right\} \right] \\
&= E\left\{ \sum_{j=1}^{K} I[\min\{T^*(d), L\}>\tau_j] \cdot \left[\min\{T^*(d), L\}-\tau_j\right] \right\} \\
&= \mathscr{V}(d)
\end{aligned}
$$

B.3 Simulated Survival Times

We conduct a simulation study comprised of two analyses. In each analysis, we use a distinct survival time generation process. For $t>0$, define $A(t)=A_1 I(0<t\leq 3)+A_2 I(3<t\leq 6)+A_3 I(6<t\leq 9)+A_4 I(t>9)$. In the first analysis, we generate T according to the accelerated failure time (AFT) model

$$
\begin{aligned}
\nu = \int_0^T \exp\Big\{ &\beta_1 B_1(u)+\beta_2 B_2(u)+\beta_3 B_3(u)+\beta_4 S+\beta_5 A(u)+ \\
&\beta_6 B_1(u)A(u)+\beta_7 B_2(u)A(u)+\beta_8 B_3(u)A(u)+\beta_9 SA(u)\Big\} d\,u,
\end{aligned}
$$

where $(\beta_1,\ldots,\beta_9)=(0.25,-0.25,0.25,0.5,-1.5,-0.25,0.25,0.25,0.25)$. For each patient, we generate a random $\nu=\exp(\theta)$, where $\theta\sim\mathbb{N}(3,1)$ and $\mathbb{N}(\mu,\Sigma)$ denotes a normal distribution with mean μ and variance-covariance Σ.

In the second analysis, we generate T according to a Cox proportional hazards model with an exponential baseline hazard, specified as

$$h(t)=\lambda\exp\Big\{\beta_1B_1(t)+\beta_2B_2(t)+\beta_3B_3(t)+\beta_4S+\beta_5A(t)+ \\ \beta_6B_1(t)A(t)+\beta_7B_2(t)A(t)+\beta_8B_3(t)A(t)+\beta_9SA(t)\Big\},$$

where $(\beta_1,\ldots,\beta_9)=(0.25,-0.5,0.25,0.5,-1.5,-0.5,-0.5,0.25,-0.25)$ and $\lambda=0.1$. For each patient, we generate a random $\nu=-\log(1-\theta)$, where $\theta\sim\mathbb{U}(0,1)$ and $\mathbb{U}$ denotes the uniform distribution.

Define $\lambda=1$ for the AFT model. Then for both the AFT and Cox models,

$$\nu=\int_0^T\lambda\exp\Big\{\beta_1B_1(u)+\beta_2B_2(u)+\beta_3B_3(u)+\beta_4S+\beta_5A(u)+ \\ \beta_6B_1(u)A(u)+\beta_7B_2(u)A(u)+\beta_8B_3(u)A(u)+\beta_9SA(u)\Big\}du. \tag{B.5}$$

For $t>0$, define $\Psi(t)=1+\sum_{j=1}^{8}jI(j<t\le j+1)+9I(t>9)$. Furthermore, let $a_{(\Psi)}=a_1I\{\Psi\in(1,2,3)\}+a_2I\{\Psi\in(4,5,6)\}+a_3I\{\Psi\in(7,8,9)\}+a_4I(\Psi=10)$. Then

$$B_l(t)=c_0^l+\mathfrak{b}_0^l+\sum_{j=1}^{\Psi(t)}\{c_j^l+\mathfrak{b}_j^l+a_{(j)}(d_j^l+\mathfrak{f}_j^l)\}\cdot\{t-(j-1)\}.$$

We can write

$$\begin{aligned}\nu(t)=\int_0^t\lambda\exp\Big\{&\beta_1B_1(u)+\beta_2B_2(u)+\beta_3B_3(u)+\beta_4s+\beta_5a(u)+ \\ &\beta_6B_1(u)a(u)+\beta_7B_2(u)a(u)+\beta_8B_3(u)a(u)+\beta_9sa(u)\Big\}du \\ &=\nu^*\{t;\Psi(t)\}+I\{\Psi(t)>1\}\sum_{j=1}^{\Psi(t)-1}\nu^*(j;j),\end{aligned}$$

where

$$
\begin{aligned}
\nu^*(t;r) &= \int_{r-1}^{t} \lambda \exp\Big\{\beta_1 B_1(u)+\beta_2 B_2(u)+\beta_3 B_3(u)+\beta_4 s+\beta_5 a(u)+ \\
&\qquad \beta_6 B_1(u)a(u)+\beta_7 B_2(u)a(u)+\beta_8 B_3(u)a(u)+\beta_9 s a(u)\Big\} d\,u \\
&= \int_{r-1}^{t} \lambda \exp\Big[\{\beta_1+\beta_6 a(u)\}B_1(u)+\{\beta_2+\beta_7 a(u)\}B_2(u)+ \\
&\qquad \{\beta_3+\beta_8 a(u)\}B_3(u)+\{\beta_4+\beta_9 a(u)\}s+\beta_5 a(u)\Big] d\,u \\
&= \int_{r-1}^{t} \lambda \exp\Big[\{\beta_4+\beta_9 a_{(r)}\}s+\beta_5 a_{(r)}+\sum_{l=1}^{3}\{\beta_l+\beta_{l+5}a_{(r)}\}B_l(u)\Big] d\,u \\
&= \lambda \exp\Big[\{\beta_4+\beta_9 a_{(r)}\}s+\beta_5 a_{(r)}\Big]\times \\
&\qquad \int_{r-1}^{t} \exp\Bigg(\sum_{l=1}^{3}\{\beta_l+\beta_{l+5}a_{(r)}\}\Big[c_0^l+\mathfrak{b}_0^l+\sum_{j=1}^{r}\{c_j^l+\mathfrak{b}_j^l+a_{(j)}(d_j^l+\mathfrak{f}_j^l)\}\{u-(j-1)\}\Big]\Bigg) d\,u \\
&= \lambda \exp\Big[\{\beta_4+\beta_9 a_{(r)}\}s+\beta_5 a_{(r)}+\sum_{l=1}^{3}\{\beta_l+\beta_{l+5}a_{(r)}\}(c_0^l+\mathfrak{b}_0^l)- \\
&\qquad \sum_{l=1}^{3}\{\beta_l+\beta_{l+5}a_{(r)}\}\sum_{j=1}^{r}\{c_j^l+\mathfrak{b}_j^l+a_{(j)}(d_j^l+\mathfrak{f}_j^l)\}(j-1)\Big]\times \\
&\qquad \int_{r-1}^{t} \exp\Big[u\sum_{l=1}^{3}\{\beta_l+\beta_{l+5}a_{(r)}\}\sum_{j=1}^{r}\{c_j^l+\mathfrak{b}_j^l+a_{(j)}(d_j^l+\mathfrak{f}_j^l)\}\Big] d\,u \\
&= \lambda \exp\Big[\{\beta_4+\beta_9 a_{(r)}\}s+\beta_5 a_{(r)}+\sum_{l=1}^{3}\{\beta_l+\beta_{l+5}a_{(r)}\}(c_0^l+\mathfrak{b}_0^l)- \\
&\qquad \sum_{l=1}^{3}\{\beta_l+\beta_{l+5}a_{(r)}\}\sum_{j=1}^{r}\{c_j^l+\mathfrak{b}_j^l+a_{(j)}(d_j^l+\mathfrak{f}_j^l)\}(j-1)\Big]\times \\
&\qquad \left[\frac{\exp\Big[u\sum_{l=1}^{3}\{\beta_l+\beta_{l+5}a_{(r)}\}\sum_{j=1}^{r}\{c_j^l+\mathfrak{b}_j^l+a_{(j)}(d_j^l+\mathfrak{f}_j^l)\}\Big]}{\sum_{l=1}^{3}\{\beta_l+\beta_{l+5}a_{(r)}\}\sum_{j=1}^{r}\{c_j^l+\mathfrak{b}_j^l+a_{(j)}(d_j^l+\mathfrak{f}_j^l)\}}\right]_{r-1}^{t} \\
&= \Bigg\{\lambda \exp\Big[\{\beta_4+\beta_9 a_{(r)}\}s+\beta_5 a_{(r)}+\sum_{l=1}^{3}\{\beta_l+\beta_{l+5}a_{(r)}\}(c_0^l+\mathfrak{b}_0^l)- \\
&\qquad \sum_{l=1}^{3}\{\beta_l+\beta_{l+5}a_{(r)}\}\sum_{j=1}^{r}\{c_j^l+\mathfrak{b}_j^l+a_{(j)}(d_j^l+\mathfrak{f}_j^l)\}(j-1)\Big]\Big/ \\
&\qquad\qquad \Big[\sum_{l=1}^{3}\{\beta_l+\beta_{l+5}a_{(r)}\}\sum_{j=1}^{r}\{c_j^l+\mathfrak{b}_j^l+a_{(j)}(d_j^l+\mathfrak{f}_j^l)\}\Big]\Bigg\}\times \\
&\qquad \Bigg(\exp\Big[t\sum_{l=1}^{3}\{\beta_l+\beta_{l+5}a_{(r)}\}\sum_{j=1}^{r}\{c_j^l+\mathfrak{b}_j^l+a_{(j)}(d_j^l+\mathfrak{f}_j^l)\}\Big]- \\
&\qquad\qquad \exp\Big[(r-1)\sum_{l=1}^{3}\{\beta_l+\beta_{l+5}a_{(r)}\}\sum_{j=1}^{r}\{c_j^l+\mathfrak{b}_j^l+a_{(j)}(d_j^l+\mathfrak{f}_j^l)\}\Big]\Bigg).
\end{aligned}
$$

We invert $\nu(t)$. Define

$$\Phi\{\nu(t)\} = 1 + \sum_{j=1}^{8} j\, I\{\nu(j) < \nu(t) \leq \nu(j+1)\} + 9 I\{\nu(t) > \nu(9)\}.$$

To simplify notation, we suppress the dependence of $\Phi\{\nu(t)\}$ on $\nu(t)$. Additionally, define

$$\mathscr{G}_1 = \sum_{l=1}^{3} \{\beta_l + \beta_{l+5} a_{(\Phi)}\} \sum_{j=1}^{\Phi} \{c_j^l + \mathfrak{b}_j^l + a_{(j)}(d_j^l + \mathfrak{f}_j^l)\},$$

$$\mathscr{G}_2 = \sum_{l=1}^{3} \{\beta_l + \beta_{l+5} a_{(\Phi)}\} \sum_{j=1}^{\Phi} \{c_j^l + \mathfrak{b}_j^l + a_{(j)}(d_j^l + \mathfrak{f}_j^l)\}(j-1).$$

Then

$$\nu^{-1}(t) = \log\left(\frac{\left\{ \nu(t) - I(\Phi > 1) \sum_{j=1}^{\Phi-1} \nu^*(j;j) \right\} \mathscr{G}_1}{\lambda \exp\left[\{\beta_4 + \beta_9 a_{(\Phi)}\} s + \beta_5 a_{(\Phi)} + \sum_{l=1}^{3} \{\beta_l + \beta_{l+5} a_{(\Phi)}\}(c_0^l + \mathfrak{b}_0^l) - \mathscr{G}_2 \right]} + \exp\left\{ (\Phi - 1)\mathscr{G}_1 \right\} \right) \Big/ \mathscr{G}_1.$$

For each patient, we compute the survival time $T = \nu^{-1}(t)$ using the patient's randomly generated ν. We administratively censor patients whose randomly generated ν is outside of the range of Equation (B.5). The distribution of $U^L = \min(T, C, L)$ for the $m = 10{,}000$ generated patients is depicted in Figure (B.1) for the AFT and Cox analyses.

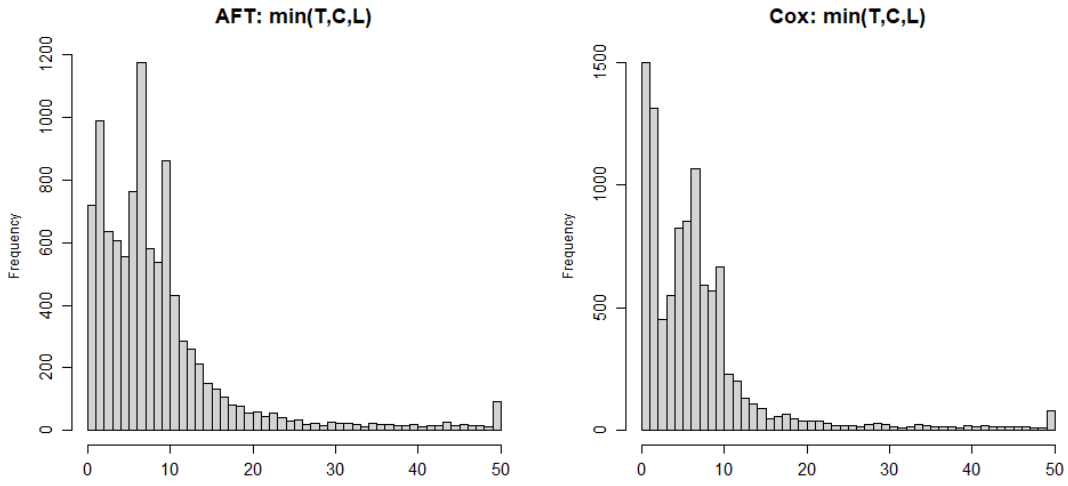

Figure B.1: The distribution of $U^L = \min(T, C, L)$ for the $m = 10,000$ patients generated in the simulation study. The distribution of U^L from the AFT survival time generation model is depicted on the left, and the distribution of U^L from the Cox survival time generation model is depicted on the right.