Title

Efficient Learning of Continuous-Time Hidden Markov Models with Discrete-Time Irregular Observations for Healthcare Intervention Planning

Author

Ghodsi, Saeed

Publication Date

2022

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

Efficient Learning of Continuous-Time Hidden Markov Models with Discrete-Time Irregular Observations for Healthcare Intervention Planning

A thesis submitted in partial satisfaction
of the requirements for the degree
Master of Science in Statistics

by

Saeed Ghodsi

2022

ABSTRACT OF THE THESIS

Efficient Learning of
Continuous-Time Hidden Markov Models
with Discrete-Time Irregular Observations
for Healthcare Intervention Planning

by

Saeed Ghodsi
Master of Science in Statistics
University of California, Los Angeles, 2022
Professor Yingnian Wu, Chair

The availability of vast amounts of electronic medical records data has inspired an increasing interest in data-driven healthcare intervention planning methods. Disease progression models provide a mechanism for understanding and predicting the impact of interventions on the health state of patients. Most traditional Markovian state-transition models perform poorly on real-world data since they are incapable of capturing complexities such as unobservability of the underlying health state and irregularity of visit times. Moreover, most of the existing frameworks are unable to explicitly model the effect of interventions on disease progression. CT-HMMs have recently attracted attention, as they are able to handle these complexities. Our main contribution is to propose a CT-HMM disease progression model, which incorporates the effect of interventions, and to present an efficient approach for learning the parameters of this model based on the EM algorithm. We demonstrate the effectiveness of our algorithm by performing experiments on synthetic data.

The thesis of Saeed Ghodsi is approved.

Reza H. Ahmadi

Arash Ali Amini

Yingnian Wu, Committee Chair

University of California, Los Angeles

2022

In dedication to my parents

for their endless love, support and encouragement.

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGMENTS

I would like to express my sincere appreciation to my committee chair, Professor Ying Nian Wu, for his continued support and encouragement as well as his compassionate and intuitive teaching style, which made every single class that I took with him a unique experience. Moreover, I would like to offer my deepest appreciation to my brilliant Management Ph.D. advisor and Statistics Master's committee member, Professor Reza Ahmadi, from whom I learned academic integrity, critical thinking, and persistence in research work. In addition, I am very grateful to Professor Sriram Dasu, who has played a key role in my entire academic journey, and made this thesis possible by his excellent mentorship and guidance. Finally, I would like to express my gratitude to my committee member and one of the best teachers that I've ever had, Professor Arash Amini, for his insightful comments that strengthened this thesis.

CHAPTER 1

Introduction

Disease progression modeling has been a hot research topic among the healthcare management, operations research, computer science, and statistics communities during the past couple of decades. The availability of vast amounts of healthcare data, such as Electronic Medical Records (EMR), enables data-driven modeling of disease progression over time and quantitative assessment of the impact of healthcare interventions on disease progression, which are key components of modern healthcare intervention planning systems. In particular, we're interested in developing mathematical models for studying the effect of healthcare interventions on medium-term and long-term progression of diseases. Such models would enable further development of decision-making frameworks for providing healthcare managers and physicians with guidelines regarding the population-level allocation of healthcare resources in order for improving the overall health state of patients and minimizing healthcare costs as a consequence. Even though the analytical framework that we present here is pretty general, our research was initially motivated by studying the case of behavioral healthcare systems, in which usually multiple intervention choices (e.g. group therapy, case management, medication prescription, etc.) are available for each patient.

Accurate modeling of disease progression is, though, difficult in its nature, as there are a variety of factors that need to be considered. In the following, we briefly discuss a few relevant aspects of the problem that make model development challenging [WSW14]:

- **Unobserved health condition**: The true underlying health state of patients is often unobserved and physicians only have access to a set of noisy signals that are correlated

with the patients' health condition. The model must be able to learn the progression of patients' health state based on the observed signals.

- **Discrete and irregularly-spaced observation times**: Although the true underlying progression of disease happens continuously over time, physicians observe the noisy signals only at certain discrete points in time. Moreover, the observations times are often irregularly spaced. In other words, physicians may have only a few observations over a certain period of time and more frequent observations during another period. Learning a continuous-time model based on these irregularly-spaced discrete-time observations is challenging in general.

- **The effect of interventions and covariates**: After each visit, the physician may prescribe medications or other interventions for the patient based on the clinical observations. Appropriate modeling of the effect of these interventions on disease progression is a key aspect of our work. Furthermore, there may be some covariates (e.g. age, gender, etc.) that affect the progression of disease.

Among all the approaches that have been proposed for modeling disease progression, State-Transition Models (STMs) that mostly impose Markovian assumptions on the structure of the time-series data are the most commonly used option. Earlier works often used Discrete-Time Markov Chains (DTMCs) and Continuous-Time Markov Chains (CTMCs). However, a recent trend in the research community is to consider the fact that the true underlying health state of patients is often unobserved by physicians. Therefore, latent-variable models have been suggested in the literature to account for the distinction between the physician observations and the true health states. Specifically, Hidden Markov Models (HMMs) as well as Partially-Observed Markov Decision Processes (POMDPs) are now preferred over fully-observed Markov chains. As we mentioned above, another important aspect of the problem that classical models usually do not consider is that in many cases visits happen irregularly over time, i.e. the time between consecutive visits can generally vary. Discrete-Time HMMs

(DT-HMMs) are inherently unable to capture this property, and traditional Continuous-Time HMMs (CT-HMMs) are often assumed to be observed continuously over time. Finally, most models assume no medical intervention can affect the course of disease progression. Although this assumption might be reasonable in certain cases in which there is no widely-approved treatment for the disease, it is restrictive for modeling more general situations. Therefore, our aim in this research project is to design a disease progression model that is able to capture the aforementioned complexities of the problem, and allows for designing a data-driven intervention planning strategy. In particular, we propose a CT-HMM model that is capable of handling unobserved health states and irregular visit times, and incorporates the effect of interventions into the model as well. Our main contribution is to propose an efficient approach for learning the parameters of this model based on the Expectation-Maximization (EM) algorithm. We demonstrate the effectiveness of our algorithm by performing simulation on synthetic data. We also present a Linear Programming (LP) formulation for a single-period intervention planning problem using our model.

In the next chapters, we first briefly review the recent disease progression modeling literature. Afterward, we will discuss recent developments for efficient parameter learning of CT-HMMs with discrete, irregular observations, that are appropriate for disease progression modeling due to their capability to model the complexities of real-world data. Our focus will be on *computational challenges* of learning CT-HMMs for real-world disease progression modeling applications that have discrete irregular observations, as explained. In particular, will go over an algorithm presented in [LLL15], that is based on the *matrix exponential* approach suggested by [HJ11] for evaluating a certain form of integrals. Then, we provide an extension of the model to incorporate the effect of *interventions* and *covariates* on disease progression by parameterizing the generator matrix of the CT-HMM. As part of the learning procedure, we present a forward-backward algorithm for efficient calculation of the posterior probabilities as well as an efficient method for calculating end-state conditioned expectations. Finally, we present an intervention planning framework based on the learned model.

CHAPTER 2

Literature Review

In this chapter, we first discuss about the application of state-transition models for disease progression modeling, and introduce some relevant recently developed Markovian models. Afterward, we briefly discuss about the intervention planning problem and go over some alternative Reinforcement Learning (RL) approaches as well. We refer the readers to [PFU18] for a review of the patient flow modeling approaches that we are not going to cover, such as queueing models which are commonly used in analyzing hospital operations.

2.1 State-Transition Models for Disease Progression

Linear regression models are among the simplest approaches that have been proposed for disease progression modeling. This approach is often suggested for situations in which the disease state is assumed to be directly related to certain biological indicators (e.g. pharmaceutical studies [CB16], [Mou12]). Multi-state models are more commonly used for disease progression modeling in the medium-term and long-term. Cohort-based state transition models take a closed group of individuals who share some specific characteristics or experiences (e.g. same disease, same risk factors, etc.) and analyze their evolution over time by running the cohort through a state-transition model. These models are attractive due to their transparency, efficiency, and ease of debugging. However, they are suitable for only situations in which a decision problem can be represented with a manageable number of health states that incorporate all the relevant characteristics [CBS12].

Among these models, DTMCs have been widely used for modeling behavioral conditions such as depression [AKG12]. On the other hand, CTMCs add more flexibility to the DTMC models by allowing transitions between states of disease to happen at arbitrary times [GLL94]. Each of these approaches have their own drawbacks and various extensions to these models have been proposed for overcoming these issues [LM13]. As we discussed beforehand, the true health state of patients may not necessarily be observable. Therefore, HMMs introduce a set of latent variables that are supposed to capture the underlying disease dynamics. The observations are then linked to these latent variables based on certain probability distributions. Being capable of learning arbitrary state transition times from irregularly observed discrete-time data is often mentioned as an advantage of CT-HMMs [LLL15].

Perhaps, one of the closest methods to our work that employs CT-HMMs for modeling disease progression using discrete-time irregularly-spaced observations is [LM13]. Specifically, the authors assume each physician observation can be mapped to a certain subset of the underlying health states, and a panel data of observations is available. They also use the Coxian proportional hazard approach for incorporating the effect of covariates on the disease progression. Afterward, the authors formulate an EM algorithm for estimating the parameters of their model. A key drawback of that approach is that it requires the generator matrix of the CT-HMM to be diagonalizable during the learning process, which usually doesn't hold. A similar approach for modeling the progression of Huntington's disease is presented in [SGL19]. We would like to mention a closely related work by [WSW14] as well, which employs CT-HMMs for modeling the time-series data of International Classification of Diseases - Version 9 (ICD-9) codes that are used for describing the disease category.

Some recent papers also consider relaxation of certain assumptions. For example, [AHS17] considers a situation in which the observation times are not random (aka informative sampling) and the physician decides the follow-up time based on the observations that s/he makes. On the other hand, STMs often assume that disease progression satisfies the Markov property, i.e. the future dynamics of the true underlying health state depends only on the current

value of the health state. However, in many situations, the duration of staying in a health state actually affects the probability of transitioning into a new health state in a given time period. Although there are some classical techniques for adapting Markovian STMs to modeling non-Markovian processes [SCS14], employing novel machine learning concepts, such as the attention mechanism, for handling non-Markovian cases is a hot research topic [AS19]. Semi-Markovian models have also been explored in the literature for handling violations of this assumption [AV18].

As we explained, most of the aforementioned papers ignore the effect of interventions on disease progression. The work by [SCS20] is an example of approaches that incorporate the effect of intervention into the model as well. In particular, the authors design a DT-HMM with a discrete set of possible hidden state and Normally distributed continuous observations. The effect of intervention is then modeled as linear shifts in the mean of the observation distribution. Thus, the actual disease progression is assumed to be independent of the intervention, and only the physician observation might be altered when the intervention changes. An advantage of their modeling framework is that it is also capable of accommodating heterogeneous effects via defining a patient-specific latent variable that is added to the observation model.

2.2 Healthcare Intervention Planning

The STM framework has been used in the past by healthcare management and operations research communities for modeling patients' disease progression, which is a crucial part of the intervention planning and resource allocation procedure. The model developed by [DIJ13] for capacity allocation in community-based healthcare delivery for chronic diseases is an example of such models. More specifically, the authors consider a homogeneous population of patients and define a representative patient whose disease progression is governed by a Markov process. Then they formulate a stochastic dynamic programming problem that is supposed to maximize the aggregate Quality-Adjusted Life-Years (QALYs) for the whole population

over the entire planning horizon. Similarly, [LLV19] developed an optimal screening model for Hepatocellular Carcinoma in which the disease progression is modeled at the individual level using a partially observable Markov decision process. A methodological framework for assessing the impact of system-level policy interventions on a set of health outcomes defined for certain chronic diseases is proposed in [KVM11]. Specifically, the authors adopted a Markovian STM for describing the pattern of visiting providers by patients over time.

Since most of the papers that incorporate the intervention effect into the model use Markov Decision Processes (MDPs) instead of HMMs, we provide a brief review of the relevant approaches in the following and discuss their drawbacks. Reinforcement learning has traditionally been used in healthcare management research for a variety of applications, including intervention planning. For example, many RL techniques have been developed in the past few years for intervention planning in critical care settings for hypotension treatment [FMD20], [ZWD21] and sepsis treatment [RKA17], [GJM18], [KCB18], [PDW18], [LSS20]. The Dynamic Treatment Regimes (DTR) framework is also known to be one of the traditional approaches for representing healthcare intervention planning as a sequential decision-making problem that aims to determine the optimal course of treatment actions [CM14]. We refer the readers to [SLL11], [GJM18], [YLN21] for more detailed discussions on applications of RL in healthcare management.

One of the most relevant RL papers to our work is [FHD20]. In that paper, the authors employ a discrete-time POMDP model with a finite number of underlying health states and Normally-distributed continuous physician observations. Assuming access to a set of trajectories of patient observations, they suggest using an EM algorithm for learning the parameters of the POMDP model similar to an Input-Output HMM [BF94]. A main advantage of their approach is that the effect of intervention is naturally incorporated into the model. The Point-Based Value Iteration (PBVI) algorithm is then used for planning based on beliefs over the latent state distribution.

Most of the traditional RL algorithms and framework have been designed for online

settings in which the agent (planner) is able to interact with the environment and explore the responses of the environment to different actions (interventions). Hence, a fundamental challenge of the so-called offline RL approaches that aim to plan using a fixed dataset without further exploration is that the distribution of the observational data might be different from the distribution induced by the actor during the exploration phase [LKT20]. In other words, the planner may explore regions of the action-state space for which sufficient data is not available in the data set. Therefore, a recent trend in the RL research community is to design algorithms that avoid exploring unseen regions. In particular, this issue is very serious in domains such as healthcare, in which minor inaccuracies in planning can potentially results in significant damages to patients. A parallel research track that studies safety in RL for healthcare applications is also closely related to the distribution shift and dataset coverage considerations that we discussed [FKS21].

Despite all the advantages that the aforementioned RL approaches offer, such as being able to naturally incorporate the effect of intervention on disease progression into the model, we decided not to use this framework for certain reasons. Specifically, most RL algorithms have been developed for discrete-time settings, while the focus of our work is on modeling a continuous-time process with discrete, irregularly-spaced observations. Even though some developments have been made in the recent years on analyzing continuous-time POMDPs [ASK20] and there are also some tricks that might help in representing our problem using discrete-time POMDPs, we've chosen to use CT-HMMs as they are more convenient to work with, and they have been explored for disease progression modeling applications in the past as well.

CHAPTER 3

Methodology

3.1 Efficient Learning of CT-HMM Parameters

In this section, we will first present a basic model setup for CT-HMMs with discrete-time and irregularly-spaced observations in the context of disease progression modeling. Afterward, we will go over the traditional learning approaches and discuss about their shortcomings. Finally, we briefly explain the idea presented in [LLL15] for efficiently learning the model parameters.

3.1.1 Model Setup

Assume that there are N patients in the healthcare system, indexed by $n = 1, \cdots, N$. Furthermore, suppose there are J possible values for the observations and I possible values for the unobserved underlying health state. Denote by $y_n^{\tau_{ns}} \in \{0, \cdots, J-1\}$ (for $s = 1, \cdots, T_n$) and $z_n^{\tau} \in \{1, \cdots, I\}$ (for the time index $\tau \in \mathbb{R}_+$), respectively, the observed value and the true disease state corresponding to patient n. In this notation, T_n number of observations are assumed to be available for patient n and $\tau_{n1}, \cdots, \tau_{nT_n}$ denote the discrete time points corresponding to these observations. Notice that we are assuming the underlying disease state $\mathbf{z}_n$ is a continuous-time random processes, and observations are made only at irregular discrete points in time (physicians are unable to observe $\mathbf{z}_n$ directly). Furthermore, we're assuming that the underlying disease progression models for all the patients have the same number of health states, which is equal to I. Define the $I \times I$ matrices $\mathcal{Q}$ and $\mathcal{P}(\Delta\tau)$, respectively, as the generator matrix and the transition probability function of the CT-HMM. Moreover, denote

by $\mathcal{E}$ the emission model, i.e. $\mathcal{E}(i,j) = \mathbb{P}\{\mathbf{y}_n^{\tau_{ns}} = j | \mathbf{z}_n^{\tau_{ns}} = i\}$ (for any $s = 1, \cdots, T_n$). For example, we can suppose the emission model associated with health state i is binomial with probability μ_i, which implies that $\mathcal{E}(i,j) = \binom{J-1}{j} \mu_i^j (1-\mu_i)^{J-1-j}$ for $1 \leq i \leq I, 0 \leq j \leq J-1$. Under the time-homogeneity assumption, we can then express the transition probability functions as $\mathcal{P}(\Delta\tau) = \exp(\Delta\tau\mathcal{Q}) = \sum_{l=0}^{\infty} (\Delta\tau)^l \frac{\mathcal{Q}^l}{l!}$, where the matrix exponential is defined based on power series.

3.1.2 Parameter Learning Using the Expectation Maximization Algorithm

The EM algorithm has traditionally been used for learning the parameters of HMMs from panel data based on the Maximum Likelihood Estimation (MLE) approach [Jac11]. The work by [MHS07] is one of the earliest successful attempts for designing an efficient learning procedure for CTMCs with discrete irregular observations. Their methodology was then extended to CT-HMMs by [LLL15]. In the following, we will briefly explain their formulation for the problem in the context of our model.

Assuming that the samples corresponding to different individuals are Independent and Identically Distributed (IID), consider the n-th patient ($1 \leq n \leq N$). Let the random variables $\tilde{\boldsymbol{\tau}}_{ns}$ (for $1 \leq s \leq \tilde{\mathbf{T}}_n$) indicate the points in time at which the true underlying state of the patient's disease (i.e. variable $\mathbf{z}_n^{\tau}$) changes. In this notation, $\tilde{\mathbf{T}}_n$ is a random variable which determines the number of times that the patient's true health state has changed. Given realizations of $\mathbf{y}_n^{\tau_{ns}}$ ($1 \leq s \leq T_n$) and $\mathbf{z}_n^{\tau}$ ($\tau_{n1} \leq \tau \leq \tau_{nT_n}$), we can represent the continuous-time complete-data likelihood function for the n-th patient as a function of $\mathbf{y}_n^{\tau_{n\,1:T_n}}$, $\mathbf{z}_n^{\tau}$, and the corresponding variables $\tilde{T}_n$ and $\tilde{\tau}_{ns}$ ($1 \leq s \leq \tilde{T}_n$) in the following form:

$$
\begin{aligned}
&\mathfrak{L}_n(\mathcal{Q}, \mathcal{E} | \mathbf{y}_n^{\tau_{n\,1:T_n}}, \mathbf{z}_n^{\tau}) \triangleq \mathbb{P}\{\mathbf{y}_n^{\tau_{n\,1:T_n}}, \mathbf{z}_n^{\tau} | \mathcal{Q}, \mathcal{E}\} \\
&= \prod_{\tilde{s}=1}^{\tilde{T}_n - 1} \left[\frac{\mathcal{Q}[z_n^{\tilde{\tau}_{n\tilde{s}}}, z_n^{\tilde{\tau}_{n\,\tilde{s}+1}}]}{-\mathcal{Q}[z_n^{\tilde{\tau}_{n\tilde{s}}}, z_n^{\tilde{\tau}_{n\tilde{s}}}]} \times (-\mathcal{Q}[z_n^{\tilde{\tau}_{n\tilde{s}}}, z_n^{\tilde{\tau}_{n\tilde{s}}}]) \, e^{\mathcal{Q}[z_n^{\tilde{\tau}_{n\tilde{s}}}, z_n^{\tilde{\tau}_{n\tilde{s}}}] (\tilde{\tau}_{n\,\tilde{s}+1} - \tilde{\tau}_{n\,\tilde{s}})} \right] \times \prod_{s=1}^{T_n} \mathcal{E}(z_n^{\tau_{ns}}, y_n^{\tau_{ns}})
\end{aligned}
$$

$$= \prod_{\tilde{s}=1}^{\tilde{T}_n - 1} \left[\mathcal{Q}[z_n^{\tilde{\tau}_{n\tilde{s}}}, z_n^{\tilde{\tau}_{n\,\tilde{s}+1}}] \times e^{\mathcal{Q}[z_n^{\tilde{\tau}_{n\tilde{s}}}, z_n^{\tilde{\tau}_{n\tilde{s}}}](\tilde{\tau}_{n\,\tilde{s}+1} - \tilde{\tau}_{n\,\tilde{s}})} \right] \times \prod_{s=1}^{T_n} \mathcal{E}(z_n^{\tau_{ns}}, y_n^{\tau_{ns}})$$

The term $\frac{\mathcal{Q}[z_n^{\tilde{\tau}_{n\tilde{s}}}, z_n^{\tilde{\tau}_{n\,\tilde{s}+1}}]}{-\mathcal{Q}[z_n^{\tilde{\tau}_{n\tilde{s}}}, z_n^{\tilde{\tau}_{n\tilde{s}}}]}$ is the probability that the embedded DTMC transitions from $z_n^{\tilde{\tau}_{n\tilde{s}}}$ to $z_n^{\tilde{\tau}_{n\,\tilde{s}+1}}$ and the term $-\mathcal{Q}[z_n^{\tilde{\tau}_{n\tilde{s}}}, z_n^{\tilde{\tau}_{n\tilde{s}}}]\, e^{\mathcal{Q}[z_n^{\tilde{\tau}_{n\tilde{s}}}, z_n^{\tilde{\tau}_{n\tilde{s}}}](\tilde{\tau}_{n\,\tilde{s}+1} - \tilde{\tau}_{n\,\tilde{s}})}$ indicates the PDF of the exponential distribution of the sojourn time corresponding to state $z_n^{\tilde{\tau}_{n\tilde{s}}}$. Notice that the parameters $\mathcal{Q}$ and $\mathcal{E}$ are assumed to be homogeneous across the patients. Now, let's define the random variables $\boldsymbol{\chi}_{n\,ik}$ $(1 \leq i, k \leq I, k \neq i)$ and $\boldsymbol{\psi}_{n\,i}$ $(1 \leq i \leq I)$, respectively, as the total number of transitions from state i to state k and the total amount of time spent in state i during the entire observation time. Then for realizations of $\mathbf{y}$ and $\mathbf{z}$, the likelihood function can be represented in the following way:

$$\mathfrak{L}_n(\mathcal{Q}, \mathcal{E} | \mathbf{y}_n^{\tau_{n\,1:T_n}}, \mathbf{z}_n^{\tau}) = \prod_{i=1}^{I} \prod_{k=1, k\neq i}^{I} \mathcal{Q}[i,k]^{\chi_{n\,ik}} \times \prod_{i=1}^{I} e^{\mathcal{Q}[i,i]\,\psi_{n\,i}} \times \prod_{s=1}^{T_n} \mathcal{E}(z_n^{\tau_{ns}}, y_n^{\tau_{ns}})$$

The likelihood for all the patients can be calculated as $\mathfrak{L} = \Pi_{n=1}^{N} \mathfrak{L}_n$. The function $\mathfrak{L}$ can not be maximized since the true underlying health state random processes $\mathbf{z}_n$ are not observable. Therefore, EM-based procedures have been proposed for estimation of the parameters. The EM algorithm determines the optimum parameters by iteratively maximizing $\mathbb{E}_{\tilde{\mathcal{Q}}, \tilde{\mathcal{E}}}[\log \mathfrak{L}(\mathcal{Q}, \mathcal{E} | \mathbf{y}_n^{\tau_{n\,1:T_n}})]$ at each step, where $\tilde{\mathcal{Q}}, \tilde{\mathcal{E}}$ are the parameter values determined at the previous iteration and $\mathcal{Q}, \mathcal{E}$ are the decision variables at the current step. In other words, the expected log-likelihood in the E-step takes the following form:

$$\begin{aligned}
\mathbb{E}_{\tilde{\mathcal{Q}}, \tilde{\mathcal{E}}}[\log \mathfrak{L}(\mathcal{Q}, \mathcal{E} | \mathbf{y}_n^{\tau_{n\,1:T_n}}] &= \mathbb{E}_{\tilde{\mathcal{Q}}, \tilde{\mathcal{E}}} \left[\sum_{n=1}^{N} \sum_{i=1}^{I} \sum_{k=1, k\neq i}^{I} \boldsymbol{\chi}_{n\,ik} \log(\mathcal{Q}[i,k]) \,\middle|\, \mathbf{y}_n^{\tau_{n\,1:T_n}} \right] \\
&+ \mathbb{E}_{\tilde{\mathcal{Q}}, \tilde{\mathcal{E}}} \left[\sum_{n=1}^{N} \sum_{i=1}^{I} \mathcal{Q}[i,i]\, \boldsymbol{\psi}_{n\,i} \,\middle|\, \mathbf{y}_n^{\tau_{n\,1:T_n}} \right] + \mathbb{E}_{\tilde{\mathcal{Q}}, \tilde{\mathcal{E}}} \left[\sum_{n=1}^{N} \sum_{s=1}^{T_n} \log(\mathcal{E}(\mathbf{z}_n^{\tau_{ns}}, \mathbf{y}^{\tau_{ns}})) \,\middle|\, \mathbf{y}_n^{\tau_{n\,1:T_n}} \right] \\
&= \sum_{n=1}^{N} \sum_{i=1}^{I} \sum_{k=1, k\neq i}^{I} \left(\log(\mathcal{Q}[i,k]) \times \mathbb{E}_{\tilde{\mathcal{Q}}, \tilde{\mathcal{E}}} \left[\boldsymbol{\chi}_{n\,ik} \,|\, \mathbf{y}_n^{\tau_{n\,1:T_n}} \right] \right)
\end{aligned}$$

$$+\sum_{n=1}^{N}\sum_{i=1}^{I}\left(\mathcal{Q}[i,i]\times\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\big[\boldsymbol{\psi}_{n\,i}\,|\,\mathbf{y}_n^{\tau_n\,1:T_n}\big]\right)+\sum_{n=1}^{N}\sum_{s=1}^{T_n}\left(\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\big[\log(\mathcal{E}(\mathbf{z}_n^{\tau_{ns}},y_n^{\tau_{ns}}))\,|\,\mathbf{y}_n^{\tau_n\,1:T_n}\big]\right)$$

This function will be maximized over $\mathcal{Q}$ and the parameters of $\mathcal{E}$ in the M-step and the following optimum parameters will be chosen (notice that $\mathcal{Q}[i,i]=-\sum_{k=1,k\neq i}^{I}\mathcal{Q}[i,k]$ holds for $1\leq i\leq I$):

$$\hat{\mathcal{Q}}[i,k]=\frac{\sum_{n=1}^{N}\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\big[\boldsymbol{\chi}_{n\,ik}\,|\,\mathbf{y}_n^{\tau_n\,1:T_n}\big]}{\sum_{n=1}^{N}\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\big[\boldsymbol{\psi}_{n\,i}\,|\,\mathbf{y}_n^{\tau_n\,1:T_n}\big]};\ (k\neq i)\qquad\hat{\mathcal{Q}}[i,i]=-\sum_{k=1,k\neq i}^{I}\hat{\mathcal{Q}}[i,k]$$

Hence, the M-step requires computing the values of $\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\big[\boldsymbol{\chi}_{n\,ik}\,|\,\mathbf{y}_n^{\tau_n\,1:T_n}\big]$ and $\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\big[\boldsymbol{\psi}_{n\,i}\,|\,\mathbf{y}_n^{\tau_n\,1:T_n}\big]$. At this stage, *computational complexity* of traditional model learning approaches makes them inappropriate for real-world applications. In particular, the Monte Carlo Expectation-Maximization (MCEM) algorithm approximates these expectations by sampling from the posterior distribution of $\mathbb{P}\{\mathbf{z}|\mathbf{y}\}$ using the current parameter estimates $\tilde{\mathcal{Q}},\tilde{\mathcal{E}}$ [LC01]. However, this approach is computationally expensive for real-world applications. In the case that performing eigendecomposition on the generator matrix is possible, the expectations can be calculated efficiently [MHS07, WSW14]. However, during the learning procedure, the generator matrix is often not diagonalizable. Hence, we need an alternative efficient approach for calculating these expectations for general cases. In the next section, we'll discuss about recent developments that have made this possible.

3.1.3 Efficient Calculation of the Expectations

An efficient mechanism for computing these quantities has recently been proposed by [LLL15]. More specifically, the authors expand the expectations in the following way:

$$\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\big[\boldsymbol{\chi}_{n\,ik}\,|\,\mathbf{y}_n^{\tau_n\,1:T_n}\big]=\sum_{s=1}^{T_n-1}\sum_{\tilde{i}=1}^{I}\sum_{\tilde{k}=1}^{I}\mathbb{P}\{\mathbf{z}_n^{\tau_n\,s}=\tilde{i},\mathbf{z}_n^{\tau_n\,s+1}=\tilde{k}\,|\,\mathbf{y}_n^{\tau_n\,1:T_n}\}$$
$$\times\,\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\big[\boldsymbol{\chi}_{n\,ik}^{\tau_n(s)}\,|\mathbf{z}_n^{\tau_n\,s}=\tilde{i},\,\mathbf{z}_n^{\tau_n\,s+1}=\tilde{k}\big]$$

$$\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\big[\boldsymbol{\psi}_{n\,i}\,|\,\mathbf{y}_n^{\tau_n\,1:T_n}\big] = \sum_{s=1}^{T_n-1}\sum_{\tilde{i}=1}^{I}\sum_{\tilde{k}=1}^{I}\mathbb{P}\{\mathbf{z}_n^{\tau_n\,s}=\tilde{i},\mathbf{z}_n^{\tau_n\,s+1}=\tilde{k}\,|\,\mathbf{y}_n^{\tau_n\,1:T_n}\}$$
$$\times\,\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\big[\boldsymbol{\psi}_{n\,i}^{\tau_n\,(s)}\,|\mathbf{z}_n^{\tau_n\,s}=\tilde{i},\,\mathbf{z}_n^{\tau_n\,s+1}=\tilde{k}\big]$$

where $\boldsymbol{\chi}_{n\,ik}^{\tau_n\,(s)}$ and $\boldsymbol{\psi}_{n\,i}^{\tau_n\,(s)}$ are the number of transitions and the total time spent in different states by variable $\mathbf{z}$ during the period $\tau_{n\,(s)} = (\tau_{ns}, \tau_{n\,s+1}]$. Then they suggest efficient procedures for calculating $\mathbb{P}\{\mathbf{z}_n^{\tau_n\,s}=\tilde{i},\mathbf{z}_n^{\tau_n\,s+1}=\tilde{k}\,|\,\mathbf{y}_n^{\tau_n\,1:T_n}\}$, $\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\big[\boldsymbol{\chi}_{n\,ik}^{\tau_n\,(s)}\,|\mathbf{z}_n^{\tau_n\,s}=\tilde{i},\,\mathbf{z}_n^{\tau_n\,s+1}=\tilde{k}\big]$, and $\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\big[\boldsymbol{\psi}_{n\,i}^{\tau_n\,(s)}\,|\mathbf{z}_n^{\tau_n\,s}=\tilde{i},\,\mathbf{z}_n^{\tau_n\,s+1}=\tilde{k}\big]$. For calculating the probability, they construct an inhomogeneous DT-HMM for which the transition matrix is constant between any two consecutive observations, but may vary after an observation has been made. In mathematical terms, the transition matrix associated with time period $[\tau_s, \tau_{s+1})$ is defined as $\mathcal{P}(\tau_{s+1}-\tau_s) = \exp((\tau_{s+1}-\tau_s)\tilde{\mathcal{Q}})$. The problem of finding the probability $\mathbb{P}\{\mathbf{z}_n^{\tau_n\,s}=\tilde{i},\mathbf{z}_n^{\tau_n\,s+1}=\tilde{k}\,|\,\mathbf{y}_n^{\tau_n\,1:T_n}\}$ can then be re-represented as the problem of finding the underlying state sequence of a DT-HMM given a set of observations, which can be solved using the forward-backward algorithm.

For calculating the end-state conditioned expectations, the authors suggest to use a method called "*Expm*", which performs matrix exponential on a double-sized auxiliary matrix. The integral form of the expectations have been derived by [HJ05], [HJ11] as:

$$\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\big[\boldsymbol{\chi}_{n\,ik}^{\tau_n\,(s)}\,|\mathbf{z}^{\tau_s}=\tilde{i},\,\mathbf{z}^{\tau_{s+1}}=\tilde{k}\big] = \frac{\tilde{\mathcal{Q}}_{ik}}{(e^{(\tau_{s+1}-\tau_s)\tilde{\mathcal{Q}}})_{\tilde{i}\tilde{k}}}\int_0^{(\tau_{s+1}-\tau_s)}(e^{x\tilde{\mathcal{Q}}})_{\tilde{i}i}(e^{(\tau_{s+1}-\tau_s-x)\tilde{\mathcal{Q}}})_{k\tilde{k}}\,dx \quad (3.1)$$

$$\mathbb{E}_{\tilde{\mathcal{Q}},\tilde{\mathcal{E}}}\big[\boldsymbol{\psi}_{n\,i}^{\tau_n\,(s)}\,|\mathbf{z}^{\tau_s}=\tilde{i},\,\mathbf{z}^{\tau_{s+1}}=\tilde{k}\big] = \frac{1}{(e^{(\tau_{s+1}-\tau_s)\tilde{\mathcal{Q}}})_{\tilde{i}\tilde{k}}}\int_0^{(\tau_{s+1}-\tau_s)}(e^{x\tilde{\mathcal{Q}}})_{\tilde{i}i}(e^{(\tau_{s+1}-\tau_s-x)\tilde{\mathcal{Q}}})_{i\tilde{k}}\,dx \quad (3.2)$$

Therefore, the problem reduces to efficient calculation of the above integral. Consider the general problem of evaluating an integral of the form $\int_0^t e^{x\tilde{\mathcal{Q}}}\,\mathcal{B}\,e^{(t-x)\tilde{\mathcal{Q}}}\,dx$, where $\mathcal{B}$ is a matrix of the same dimensions as $\tilde{\mathcal{Q}}$. It is known that the result of this integral is equal to the upper-right corner of the matrix exponential $e^{t\mathcal{A}}$, where $\mathcal{A} = \begin{pmatrix}\tilde{\mathcal{Q}} & \mathcal{B}\\ 0 & \tilde{\mathcal{Q}}\end{pmatrix}$ [Van78]. Setting the elements of $\mathcal{B}$ equal to zero everywhere except for $\mathcal{B}_{ik}=1$, the integral will result in a matrix

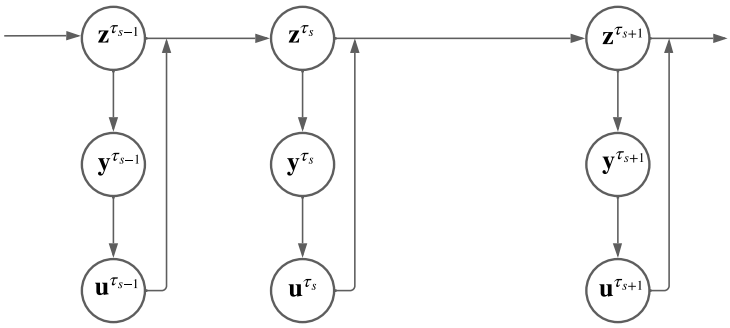

Figure 3.1: CT-HMM model with interventions at the discrete and irregular observation times. The main difference between this model and the original model is that we have an additional layer for intervention variables. The intervention decision at visit s affects the disease progression between visits s and $s+1$ (i.e. the generator matrix is not constant).

with elements $(\tilde{i}, \tilde{k})$ corresponding to the values of indices in the original integral.

3.2 Incorporating the Effect of Intervention

In this section, we reformulate the model to incorporating the effect of interventions and covariates on disease progression. We first suggest a parameterization of the generator matrix that allows for making the transition rates a function of intervention variables and covariates. We then assume that at the training time, we have access to the observation and intervention variables. In the context of disease progression modeling, observations can be the results of a set of diagnostic questions that the physician asks from their patients, and interventions can be the drugs that are prescribe based on the patient's answers and symptoms. Specifically, we consider the model structure presented in figure 3.1, and develop an expectation maximization algorithm, based on the techniques we discussed earlier, for learning the model parameters.

3.2.1 Model Setup

Suppose we have N patients in the system. We consider a modified version of the model for patient $1 \leq n \leq N$ that is described in the following. There is an underlying true health state

variable denoted by $\mathbf{z}_n^\tau$ that is continuous-time. We denote by $\mathbf{z}_n^{\tau_{ns}}$ the value of $\mathbf{z}_n^\tau$ at time τ_{ns} and by the continuous-time variable $\mathbf{z}_n^{\tau_{n\,(s)}}$ the values of $\mathbf{z}_n^\tau$ between times $(\tau_{ns}, \tau_{n\,s+1}]$. A physician observes a noisy version of this variable at discrete points in time that are denoted by $\mathbf{y}_n^{\tau_{ns}}$, where $s = 1, \cdots, T_n$, and T_n indicates the number of observations for patient n. A treatment option $\mathbf{u}_n^{\tau_{ns}}$ is then assigned to patient based on the observation. We assume that in the parameter estimation phase, variables $\mathbf{z}$ are hidden and we only have access to historical recordings of $\mathbf{y}$ and $\mathbf{u}$. Suppose the variables $\mathbf{z}$, $\mathbf{y}$, and $\mathbf{u}$ are all discrete-valued and take one of the values $\{1, \cdots, I\}$, $\{0, \cdots, J-1\}$, and $\{0, \cdots, L-1\}$ respectively. The generator matrix of the underlying CTMC that governs the dynamics of process $\mathbf{z}$ is denoted by $\mathcal{Q} \in \mathbb{R}^{I \times I}$. The emission distribution $\mathcal{E}(i,j) = \mathbb{P}(\mathbf{y}_n^{\tau_{ns}} = j \,|\, \mathbf{z}_n^{\tau_{ns}} = i)$ and intervention distribution $\mathcal{G}(j,l) = \mathbb{P}(\mathbf{u}_n^{\tau_{ns}} = l \,|\, \mathbf{y}_n^{\tau_{ns}} = j)$ $(1 \leq s \leq T_n)$ are assumed to be binomial with parameters μ_i $(1 \leq i \leq I)$ and η_j $(0 \leq j \leq J-1)$, respectively. Finally, we denote by $\pi \in \mathbb{R}^I$ the vector of initial probabilities of $\mathbf{z}$.

3.2.2 Parameterizing the Generator Matrix

In practice, many covariates (e.g. age, sex, etc.) may affect the patients' disease progression. In the disease progression modeling literature, these covariates are either ignored or modeled in an explicit manner. The most common approach for incorporating covariates into the model is to define parameters according to the Cox proportional hazards function [Cox72]. [KL85] formulated the problem of learning CTMCs using irregularly observed panel data for the first time. In their model, the authors defined the elements of the generator matrix as exponential functions of linear transformations of the vector of covariates. The same approach has been widely adopted by others as well. Despite the prevalence of this modeling approach, we have not found any clinical evidence that justifies the suitability of this functional form.

We consider the Cox's proportional hazards model for incorporating the effects of covariates

and treatment on the progression of disease:

$$[\mathcal{Q}]_{ikl} = \delta_{ikl} \exp(\rho'_{ikl} w_n); \quad 1 \le i,k \le I,\ k \ne i,\ 0 \le l \le L-1 \tag{3.3}$$

$$[\mathcal{Q}]_{iil} = -\sum_{k=1,k\ne i}^{I} [\mathcal{Q}]_{ikl}; \quad 1 \le i \le I,\ 0 \le l \le L-1 \tag{3.4}$$

where $w_n \in \mathbb{R}^D$ is the vector of covariates associated with patient n, and l is the intervention variable. Specifically, assume that in each period (s) (i.e., between $(\tau_s, \tau_{s+1}]$ for a fixed $1 \le s \le T_n$) the generator matrix is constant and denoted by $\mathcal{Q}_{n\,(s)}$ for patient n. By setting $l = \mathbf{u}_n^{\tau_{ns}}$, the off-diagonal (i,k)'th element of $\mathcal{Q}_{n\,(s)}$ can be modeled as $\delta_{ik\,\mathbf{u}_n^{\tau_{ns}}} \exp(\rho'_{ik\,\mathbf{u}_n^{\tau_{ns}}} w_n)$.

3.2.3 Learning the Model Parameters via the EM Algorithm

Define $\Theta \triangleq (\pi, \rho, \delta, \mu, \eta)$. In this section, we provide the basic formulation for estimating Θ based on the EM algorithm. A detailed description of the algorithm will be later presented in the appendix. The complete-data likelihood for patient n can be written as:

$$\begin{aligned}\mathcal{L}_n(\Theta) = \mathbb{P}(\mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}, \mathbf{z}_n^{\tau_n}|\Theta) &= \mathbb{P}(\mathbf{u}_n^{\tau_n T_n}|\mathbf{y}_n^{\tau_n T_n}, \eta)\,\mathbb{P}(\mathbf{y}_n^{\tau_n T_n}|\mathbf{z}_n^{\tau_n T_n}, \mu) \\ &\times \prod_{s=1}^{T_n-1} \left[\mathbb{P}(\mathbf{z}_n^{\tau_n\,(s)}|\mathbf{z}_n^{\tau_{ns}}, \mathbf{u}_n^{\tau_{ns}}, \rho, \delta)\,\mathbb{P}(\mathbf{u}_n^{\tau_{ns}}|\mathbf{y}_n^{\tau_{ns}}, \eta)\,\mathbb{P}(\mathbf{y}_n^{\tau_{ns}}|\mathbf{z}_n^{\tau_{ns}}, \mu)\right] \times \mathbb{P}(\mathbf{z}_n^{\tau_{n1}})\end{aligned}$$

The complete-data log-likelihood will then be:

$$\begin{aligned}\log\,\mathcal{L}_n(\Theta) = \log\mathbb{P}(\mathbf{z}_n^{\tau_{n1}}) &+ \sum_{s=1}^{T_n-1}\left[\log\mathbb{P}(\mathbf{z}_n^{\tau_n\,(s)}|\mathbf{z}_n^{\tau_{ns}}, \mathbf{u}_n^{\tau_{ns}}, \rho, \delta)\right] \\ &+ \sum_{s=1}^{T_n}\left[\log\mathbb{P}(\mathbf{u}_n^{\tau_{ns}}|\mathbf{y}_n^{\tau_{ns}}, \eta) + \log\mathbb{P}(\mathbf{y}_n^{\tau_{ns}}|\mathbf{z}_n^{\tau_{ns}}, \mu)\right]\end{aligned}$$

However, we cannot calculate the value of $\log\,\mathcal{L}_n(\Theta)$ since the $\mathbf{z}$ variables are not observed. Therefore, we start with a set of initial parameters denoted by $\tilde{\Theta} \triangleq (\tilde{\pi}, \tilde{\rho}, \tilde{\delta}, \tilde{\mu}, \tilde{\eta})$ and find the posterior distribution of $\mathbf{z}$. Then we calculate the expected value of $\log\,\mathcal{L}_n(\Theta)$ (i.e.

$\mathbb{E}[\log \mathcal{L}_n(\Theta)|\mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}, \tilde{\Theta}]$ over the conditional distribution of $\mathbf{z}$) and find the value of Θ that maximize this expectation. This procedure is repeated iteratively until convergence. For a given realization of $\mathbf{z}$, the function $\log \mathcal{L}_n(\Theta)$ can be simplified as we'll discuss in the following.

Define $\chi_{n\,ik}^{\tau_n\,(s)}$ as the total number of transition from state i to state k during the time period $(\tau_{ns}, \tau_{n\,s+1}]$. Moreover, define $\psi_{n\,i}^{\tau_n\,(s)}$ as the total amount of time that $\mathbf{z}$ has spent in state i during this period. Remember that $([\mathcal{Q}_{n\,(s)}]_{ik} \,|\, \mathbf{u}_n^{\tau_{ns}}, \rho, \delta) = \delta_{ik\,\mathbf{u}_n^{\tau_{ns}}} \exp(\rho'_{ik\,\mathbf{u}_n^{\tau_{ns}}} w_n)$ for $k \neq i$. So given $\mathbf{u}_n^{\tau_{ns}}$, ρ, and δ, the generator matrix $Q_{n\,(s)}$ is known deterministically. Therefore:

$$\mathbb{P}(\mathbf{z}_n^{\tau_n\,(s)}|\mathbf{z}_n^{\tau_{ns}}, \mathbf{u}_n^{\tau_{ns}}, \rho, \delta) = \prod_{i=1}^{I} \prod_{\substack{k=1 \\ k\neq i}}^{I} \left[Q_{n\,(s)}\right]_{ik}^{\chi_{n\,ik}^{\tau_n\,(s)}} \times \prod_{i=1}^{I} \exp\left([\mathcal{Q}_{n\,(s)}]_{ii}\, \psi_{n\,i}^{\tau_n\,(s)}\right)$$

$$\log \mathbb{P}(\mathbf{z}_n^{\tau_n\,(s)}|\mathbf{z}_n^{\tau_{ns}}, \mathbf{u}_n^{\tau_{ns}}, \rho, \delta) = \sum_{i=1}^{I} \sum_{\substack{k=1 \\ k\neq i}}^{I} \left(\chi_{n\,ik}^{\tau_n\,(s)} \log([\mathcal{Q}_{n\,(s)}]_{ik}) - \psi_{n\,i}^{\tau_n\,(s)}\, [\mathcal{Q}_{n\,(s)}]_{ik}\right)$$

where we used the fact that $[\mathcal{Q}_{n\,(s)}]_{ii} = -\sum_{k=1 k\neq i}^{I} [\mathcal{Q}_{n\,(s)}]_{ik}$. Furthermore, we have:

$$\log \mathbb{P}(\mathbf{y}_n^{\tau_{ns}}|\mathbf{z}_n^{\tau_{ns}}, \mu) = \sum_{i=1}^{I} \left[\log \binom{J-1}{\mathbf{y}_n^{\tau_{ns}}} + \mathbf{y}_n^{\tau_{ns}} \log \mu_i + (J - 1 - \mathbf{y}_n^{\tau_{ns}}) \log(1-\mu_i)\right] \mathbb{I}\{\mathbf{z}_n^{\tau_{ns}} = i\}$$

$$\log \mathbb{P}(\mathbf{u}_n^{\tau_{ns}}|\mathbf{y}_n^{\tau_{ns}}, \eta) = \sum_{j=0}^{J-1} \left[\log \binom{L-1}{\mathbf{u}_n^{\tau_{ns}}} + \mathbf{u}_n^{\tau_{ns}} \log \eta_j + (L - 1 - \mathbf{u}_n^{\tau_{ns}}) \log(1-\eta_j)\right] \mathbb{I}\{\mathbf{y}_n^{\tau_{ns}} = j\}$$

where $\mathbb{I}$ is the indicator function. Finally, notice that $\log \mathbb{P}(\mathbf{z}_n^{\tau_{n1}}) = \sum_{i=1}^{I} \log \pi_i\, \mathbb{I}\{\mathbf{z}_n^{\tau_{n1}} = i\}$. The complete-data log-likelihood for all the patients will then be:

$$\begin{aligned}\log \mathcal{L}(\Theta) = \sum_{n=1}^{N} \log\, \mathcal{L}_n(\Theta) &= \sum_{n=1}^{N} \sum_{i=1}^{I} \log(\pi_i)\, \mathbb{I}\{\mathbf{z}_n^{\tau_{n1}} = i\} \\ &+ \sum_{n=1}^{N} \sum_{s=1}^{T_n-1} \sum_{i=1}^{I} \sum_{\substack{k=1 \\ k\neq i}}^{I} \left[\chi_{n\,ik}^{\tau_n\,(s)} \left(\log(\delta_{ik\,\mathbf{u}_n^{\tau_{ns}}}) + \rho'_{ik\,\mathbf{u}_n^{\tau_{ns}}} w_n\right) - \psi_{n\,i}^{\tau_n\,(s)}\, \delta_{ik\,\mathbf{u}_n^{\tau_{ns}}} \exp(\rho'_{ik\,\mathbf{u}_n^{\tau_{ns}}} w_n)\right]\end{aligned}$$

$$+\sum_{n=1}^{N}\sum_{s=1}^{T_n}\sum_{i=1}^{I}\left[\log\binom{J-1}{\mathbf{y}_n^{\tau_{ns}}}+\mathbf{y}_n^{\tau_{ns}}\log(\mu_i)+(J-1-\mathbf{y}_n^{\tau_{ns}})\log(1-\mu_i)\right]\mathbb{I}\{\mathbf{z}_n^{\tau_{ns}}=i\}$$
$$+\sum_{n=1}^{N}\sum_{s=1}^{T_n}\sum_{j=0}^{J-1}\left[\log\binom{L-1}{\mathbf{u}_n^{\tau_{ns}}}+\mathbf{u}_n^{\tau_{ns}}\log(\eta_j)+(L-1-\mathbf{u}_n^{\tau_{ns}})\log(1-\eta_j)\right]\mathbb{I}\{\mathbf{y}_n^{\tau_{ns}}=j\}$$

The above expression cannot be calculated due to the uncertainty of $\mathbf{z}$, as explained. Therefore, our strategy is to find the expected value of $\log\mathcal{L}$ over the conditional distribution of $\mathbf{z}$ (i.e., $\mathbf{z}_n^{\tau}|\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta}$).

$$\begin{aligned}
&\mathbb{E}[\log\,\mathcal{L}(\Theta)|\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta}]\\
&=\sum_{n=1}^{N}\sum_{i=1}^{I}\log(\pi_i)\,\mathbb{P}(\mathbf{z}_n^{\tau_n 1}=i\,|\,\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta})\\
&+\sum_{n=1}^{N}\sum_{s=1}^{T_n-1}\sum_{i=1}^{I}\sum_{\substack{k=1\\k\neq i}}^{I}\sum_{l=0}^{L-1}\Big[\mathbb{E}\big[\chi_{n\,ik}^{\tau_n(s)}|\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta}\big]\big(\log(\delta_{ikl})+\rho_{ikl}^{'}w_n\big)\\
&\qquad\qquad-\mathbb{E}[\psi_{n\,i}^{\tau_n(s)}|\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta}]\,\delta_{ikl}\,\exp(\rho_{ikl}^{'}w_n)\Big]\mathbb{I}\{\mathbf{u}_n^{\tau_{ns}}=l\}\\
&+\sum_{n=1}^{N}\sum_{s=1}^{T_n}\sum_{i=1}^{I}\left[\log\binom{J-1}{\mathbf{y}_n^{\tau_{ns}}}+\mathbf{y}_n^{\tau_{ns}}\log(\mu_i)+(J-1-\mathbf{y}_n^{\tau_{ns}})\log(1-\mu_i)\right]\\
&\qquad\qquad\times\mathbb{P}(\mathbf{z}_n^{\tau_{ns}}=i\,|\,\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta})\\
&+\sum_{n=1}^{N}\sum_{s=1}^{T_n}\sum_{j=0}^{J-1}\left[\log\binom{L-1}{\mathbf{u}_n^{\tau_{ns}}}+\mathbf{u}_n^{\tau_{ns}}\log(\eta_j)+(L-1-\mathbf{u}_n^{\tau_{ns}})\log(1-\eta_j)\right]\mathbb{I}\{\mathbf{y}_n^{\tau_{ns}}=j\}
\end{aligned}$$

For the moment, assume we have calculated the above conditional expectations and probabilities. We will then be able to update the π,ρ,δ,μ,η parameters by maximizing the objective with respect to these parameters in the M-step. In particular, define $\gamma_n^s(i)=\mathbb{P}(\mathbf{z}_n^{\tau_{ns}}=i\,|\,\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta})$ (for $1\leq s\leq T_n$) and $\nu_n^s(i,k)=\mathbb{P}(\mathbf{z}_n^{\tau_{ns}}=i,\mathbf{z}_n^{\tau_n\,s+1}=k\,|\,\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta})$ (for $1\leq s\leq T_n-1$). Now, let's assume we replace $\pi_I=1-\sum_{i'=1}^{I-1}\pi_{i'}$ in the objective function

and take the derivative with respect to some π_i:

$$\frac{\partial}{\partial \pi_i}\mathbb{E}\big[\log\,\mathcal{L}(\Theta)|\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta}\big]=\frac{\partial}{\partial \pi_i}\sum_{i'=1}^{I-1}\log(\pi_{i'})\left(\sum_{n=1}^{N}\gamma_n^1(i')\right)$$
$$+\frac{\partial}{\partial \pi_i}\log(1-\sum_{i'=1}^{I-1}\pi_{i'})\left(\sum_{n=1}^{N}\gamma_n^1(I)\right)=\frac{1}{\pi_i}\left(\sum_{n=1}^{N}\gamma_n^1(i)\right)-\frac{1}{\pi_I}\left(\sum_{n=1}^{N}\gamma_n^1(I)\right)$$

By setting the derivative to zero, we get $\hat{\pi}_i=\frac{\sum_{n=1}^{N}\gamma_n^1(i)}{\sum_{n=1}^{N}\gamma_n^1(I)}\,\hat{\pi}_I$ (for $1\leq i\leq I-1$). The constraint $\sum_{i=1}^{I}\hat{\pi}_i=1$ simplifies to $\hat{\pi}_I=\frac{\sum_{n=1}^{N}\gamma_n^1(I)}{\sum_{i'=1}^{I}\sum_{n=1}^{N}\gamma_n^1(i')}$, which implies $\hat{\pi}_i=\frac{\sum_{n=1}^{N}\gamma_n^1(i)}{\sum_{i'=1}^{I}\sum_{n=1}^{N}\gamma_n^1(i')}$. On the other hand, taking the derivative with respect to μ_i gives us:

$$\frac{\partial}{\partial \mu_i}\mathbb{E}\big[\log\,\mathcal{L}(\Theta)|\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta}\big]=\sum_{n=1}^{N}\sum_{s=1}^{T_n}\left[\frac{1}{\mu_i}\mathbf{y}_n^{\tau_{ns}}-\frac{1}{1-\mu_i}(J-1-\mathbf{y}_n^{\tau_{ns}})\right]\gamma_n^s(i)$$
$$=\frac{1}{\mu_i}\left(\sum_{n=1}^{N}\sum_{s=1}^{T_n}\mathbf{y}_n^{\tau_{ns}}\gamma_n^s(i)\right)-\frac{1}{1-\mu_i}\left(\sum_{n=1}^{N}\sum_{s=1}^{T_n}(J-1-\mathbf{y}_n^{\tau_{ns}})\gamma_n^s(i)\right)$$

The optimal parameters will then $\hat{\mu}_i=\frac{\sum_{n=1}^{N}\sum_{s=1}^{T_n}\mathbf{y}_n^{\tau_{ns}}\gamma_n^s(i)}{(J-1)\sum_{n=1}^{N}\sum_{s=1}^{T_n}\gamma_n^s(i)}$ (for $1\leq i\leq I$). Similarly, we can set the derivative with respect to η_j equal to zero and obtain $\hat{\eta}_j=\frac{\sum_{n=1}^{N}\sum_{s=1}^{T_n}\mathbf{u}_n^{\tau_{ns}}\mathbb{I}\{\mathbf{y}_n^{\tau_{ns}}=j\}}{(L-1)\sum_{n=1}^{N}\sum_{s=1}^{T_n}\mathbb{I}\{\mathbf{y}_n^{\tau_{ns}}=j\}}$ (for $0\leq j\leq J-1$).

We can find the optimum parameters $\hat{\delta}$ and $\hat{\rho}$ by applying the Newton's method for optimizing the objective function. Specifically, the derivative of the expected complete-data log-likelihood with respect to δ_{ikl} is:

$$\frac{\partial}{\partial\,\delta_{ikl}}\mathbb{E}\big[\log\,\mathcal{L}(\Theta)|\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta}\big]=\frac{1}{\delta_{ikl}}\sum_{n=1}^{N}\sum_{s=1}^{T_n-1}\mathbb{E}\big[\chi_{n\,ik}^{\tau_n\,(s)}|\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta}\big]\mathbb{I}\{\mathbf{u}_n^{\tau_{ns}}=l\}$$
$$-\sum_{n=1}^{N}\sum_{s=1}^{T_n-1}\mathbb{E}[\psi_{n\,i}^{\tau_n\,(s)}|\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta}]\,\exp(\rho_{ikl}^{'}\,w_n)\mathbb{I}\{\mathbf{u}_n^{\tau_{ns}}=l\}$$

Thus, having the optimal $\hat{\rho}$ parameters, the solution will be:

$$\hat{\delta}_{ikl} = \frac{\sum_{n=1}^{N}\sum_{s=1}^{T_n-1}\mathbb{E}\big[\chi_{n\,ik}^{\tau_n(s)}|\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta}\big]\mathbb{I}\{\mathbf{u}_n^{\tau_{ns}}=l\}}{\sum_{n=1}^{N}\sum_{s=1}^{T_n-1}\mathbb{E}[\psi_{n\,i}^{\tau_n(s)}|\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta}]\,\exp(\hat{\rho}_{ikl}^{'}\,w_n)\mathbb{I}\{\mathbf{u}_n^{\tau_{ns}}=l\}} \tag{3.5}$$

Let's calculate the first and second derivatives of the objective function with respect to ρ_{ikl}:

$$\frac{\partial}{\partial\,\rho_{ikl}}\mathbb{E}[\log\,\mathcal{L}(\Theta)|\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta}] = \sum_{n=1}^{N}\sum_{s=1}^{T_n-1}\mathbb{I}\{\mathbf{u}_n^{\tau_{ns}}=l\}$$
$$\Big(\mathbb{E}[\chi_{n\,ik}^{\tau_n(s)}|\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta})] - \mathbb{E}[\psi_{n\,i}^{\tau_n(s)}|\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta}]\,\hat{\delta}_{ikl}\,\exp(\rho_{ikl}^{'}w_n)\Big)w_n \tag{3.6}$$

$$\frac{\partial^2}{\partial\,\rho_{ikl}^2}\mathbb{E}[\log\,\mathcal{L}(\Theta)|\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta}] = -\sum_{n=1}^{N}\sum_{s=1}^{T_n-1}\mathbb{I}\{\mathbf{u}_n^{\tau_{ns}}=l\}$$
$$\Big(\mathbb{E}[\psi_{n\,i}^{\tau_n(s)}|\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta}]\,\hat{\delta}_{ikl}\,\exp(\rho_{ikl}^{'}w_n)\Big)w_n w_n^{'} \tag{3.7}$$

We suggest iteratively updating $\hat{\delta}$ according to equation 3.5 and $\hat{\rho}$ using the Newton's method based on equations 3.6 and 3.7. Therefore, the M-step can be completed efficiently if we calculate the posterior probabilities $\gamma_n^s(\cdot)$ and $\nu_n^s(\cdot,\cdot)$ along with the conditional expectations $\mathbb{E}[\chi_{n\,ik}^{\tau_n(s)}|\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta})]$ and $\mathbb{E}[\psi_{n\,i}^{\tau_n(s)}|\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta}]$ in the E-step.

The expectations can be further expanded by conditioning on the true underlying health states at the consecutive observation times:

$$\mathbb{E}[\chi_{n\,ik}^{\tau_n(s)}|\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta})] = \sum_{i'=1}^{I}\sum_{k'=1}^{I}\mathbb{E}[\chi_{n\,ik}^{\tau_n(s)}|\mathbf{z}_n^{\tau_{ns}}=i^{'},\mathbf{z}_n^{\tau_n\,s+1}=k^{'},\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta}]\,\nu_n^s(i',k')$$

$$\mathbb{E}[\psi_{n\,i}^{\tau_n(s)}|\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta}] = \sum_{i'=1}^{I}\sum_{k'=1}^{I}\mathbb{E}[\psi_{n\,i}^{\tau_n(s)}|\mathbf{z}_n^{\tau_{ns}}=i^{'},\mathbf{z}_n^{\tau_n\,s+1}=k^{'},\mathbf{u}_n^{\tau_n\,1:T_n},\mathbf{y}_n^{\tau_n\,1:T_n},\tilde{\Theta}]\,\nu_n^s(i',k')$$

Now, the end state-conditioned expectations can be calculated by the method that [LLL15] proposed according to equations 3.1 and 3.2, as we briefly discussed in the previous section. The main difference here would be that we need to construct the generator matrix $\tilde{\mathcal{Q}}_{n\,(s)}$, associated with period $(\tau_{ns},\tau_{n\,s+1}]$, based on equations 3.3 and 3.4 by setting $l=\mathbf{u}_n^{\tau_{ns}}$ and

using parameters $\tilde{\delta}$, $\tilde{\rho}$. We'll later provide the detailed formulas for evaluating the expectations in the appendix.

3.2.4 The Forward-Backward Algorithm

At this point, the problem has been reduced to efficient calculation of $\gamma_n^s(i) = \mathbb{P}(\mathbf{z}_n^{\tau_{ns}} = i|\mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}, \tilde{\Theta})$ and $\nu_n^s(i,k) = \mathbb{P}(\mathbf{z}_n^{\tau_{ns}} = i, \mathbf{z}_n^{\tau_n\,s+1} = k|\mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}, \tilde{\Theta})$. To calculate these probabilities, we construct a time-inhomogeneous DT-HMM from our CT-HMM, by defining the transition probability function of the DT-HMM to be equal to $[\tilde{\mathcal{P}}_{n\,(s)}]_{ik} = \mathbb{P}(\tilde{\mathbf{z}}_n^{s+1} = k|\tilde{\mathbf{z}}_n^s = i, \mathbf{u}_n^{\tau_{ns}}, \tilde{\Theta}) = \left[e^{(\tau_n\,s+1 - \tau_{ns})\tilde{\mathcal{Q}}_{n\,(s)}}\right]_{ik}$ where $\tilde{\mathbf{z}}$ is the discrete-time latent variable of our DT-HMM and $\tilde{\mathcal{Q}}_{n\,(s)}$ is the generator matrix, associated with interval $(\tau_{ns}, \tau_{n\,s+1}]$, as we explained above.

We start by calculating $\gamma_n^s(i) = \mathbb{P}(\tilde{\mathbf{z}}_n^s = i|\mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}, \tilde{\Theta})$. The forward and backward probabilities will be $\alpha_n^s(i) = \mathbb{P}(\tilde{\mathbf{z}}_n^s = i, \mathbf{u}_n^{\tau_n\,1:s}, \mathbf{y}_n^{\tau_n\,1:s}|\tilde{\Theta})$ and $\beta_n^s(i) = \mathbb{P}(\mathbf{u}_n^{\tau_n\,s+1:T_n}, \mathbf{y}_n^{\tau_n\,s+1:T_n}|\tilde{\mathbf{z}}_n^s = i, \mathbf{u}_n^{\tau_{ns}}, \tilde{\Theta})$, respectively. The dynamic programming equations for calculating the forward and backward variables are:

$$
\begin{aligned}
\alpha_n^s(i) &= \sum_{k=1}^{I} \mathbb{P}(\tilde{\mathbf{z}}_n^s = i, \tilde{\mathbf{z}}_n^{s-1} = k, \mathbf{u}_n^{\tau_n\,1:s}, \mathbf{y}_n^{\tau_n\,1:s}|\tilde{\Theta}) \\
&= \sum_{k=1}^{I} \mathbb{P}(\mathbf{u}_n^{\tau_{ns}}|\mathbf{y}_n^{\tau_{ns}}, \tilde{\Theta}) \mathbb{P}(\mathbf{y}_n^{\tau_{ns}}|\tilde{\mathbf{z}}_n^s = i, \tilde{\Theta})\, \mathbb{P}(\tilde{\mathbf{z}}_n^s = i|\tilde{\mathbf{z}}_n^{s-1} = k, \mathbf{u}_n^{\tau_n\,s-1}, \tilde{\Theta}) \\
&\qquad \mathbb{P}(\tilde{\mathbf{z}}_n^{s-1} = k, \mathbf{u}_n^{\tau_n\,1:s-1}, \mathbf{y}_n^{\tau_n\,1:s-1}|\tilde{\Theta}) \\
&= \sum_{k=1}^{I} \tilde{\mathcal{G}}(\mathbf{y}_n^{\tau_{ns}}, \mathbf{u}_n^{\tau_{ns}})\, \tilde{\mathcal{E}}(i, \mathbf{y}_n^{\tau_{ns}})\, \left[\tilde{\mathcal{P}}_{n\,(s-1)}\right]_{ki} \alpha_n^{s-1}(k) \qquad (3.8) \\
\beta_n^s(i) &= \sum_{k=1}^{I} \mathbb{P}(\tilde{\mathbf{z}}_n^{s+1} = k, \mathbf{u}_n^{\tau_n\,s+1:T_n}, \mathbf{y}_n^{\tau_n\,s+1:T_n}|\tilde{\mathbf{z}}_n^s = i, \mathbf{u}_n^{\tau_{ns}}, \tilde{\Theta}) \\
&= \sum_{k=1}^{I} \mathbb{P}(\mathbf{u}_n^{\tau_n\,s+2:T_n}, \mathbf{y}_n^{\tau_n\,s+2:T_n}|\tilde{\mathbf{z}}_n^{s+1} = k, \mathbf{u}_n^{\tau_n\,s+1}, \tilde{\Theta})\, \mathbb{P}(\mathbf{u}_n^{\tau_n\,s+1}|\mathbf{y}_n^{\tau_n\,s+1}, \tilde{\Theta}) \\
&\qquad \mathbb{P}(\mathbf{y}_n^{\tau_n\,s+1}|\tilde{\mathbf{z}}_n^{s+1} = k, \tilde{\Theta})\, \mathbb{P}(\tilde{\mathbf{z}}_n^{s+1} = k|\tilde{\mathbf{z}}_n^s = i, \mathbf{u}_n^{\tau_{ns}}, \tilde{\Theta})
\end{aligned}
$$

$$= \sum_{k=1}^{I} \tilde{\mathcal{G}}(\mathbf{y}_n^{\tau_{n\,s+1}}, \mathbf{u}_n^{\tau_{n\,s+1}})\, \tilde{\mathcal{E}}(k, \mathbf{y}_n^{\tau_{n\,s+1}})\, \left[\tilde{\mathcal{P}}_{n\,(s)}\right]_{ik} \beta_n^{s+1}(k) \tag{3.9}$$

where $\tilde{\mathcal{E}}(\cdot,\cdot)$ and $\tilde{\mathcal{G}}(\cdot,\cdot)$ are the emission and intervention probabilities using the $\tilde{\mu}$ and $\tilde{\eta}$ parameters, respectively. Notice that $\tilde{\mathcal{P}}_{n\,(s-1)} = e^{(\tau_{ns} - \tau_{n\,s-1})\tilde{Q}_{n\,(s-1)}}$ and $\tilde{\mathcal{P}}_{n\,(s)} = e^{(\tau_{n\,s+1} - \tau_{ns})\tilde{Q}_{n\,(s)}}$ depend on matrices $\tilde{Q}_{n\,(s)}$ and $\tilde{Q}_{n\,(s-1)}$ that are known as $\tilde{\delta}$, $\tilde{\rho}$, $\mathbf{u}_n^{\tau_{ns}}$, and $\mathbf{u}_n^{\tau_{n\,s-1}}$ are given. The boundary cases for the forward and backward equations are:

$$\alpha_n^1(i) = \tilde{\mathcal{G}}(\mathbf{y}_n^{\tau_{n1}}, \mathbf{u}_n^{\tau_{n1}})\, \tilde{\mathcal{E}}(i, \mathbf{y}_n^{\tau_{n1}}) \tilde{\pi}_i \tag{3.10}$$

$$\beta_n^{T_n-1}(i) = \sum_{k=1}^{I} \tilde{\mathcal{G}}(\mathbf{y}_n^{\tau_{n\,T_n}}, \mathbf{u}_n^{\tau_{n\,T_n}})\, \tilde{\mathcal{E}}(k, \mathbf{y}_n^{\tau_{n\,T_n}})\, \left[\tilde{\mathcal{P}}_{n\,(T_n-1)}\right]_{ik} \tag{3.11}$$

The posterior probability of $\tilde{\mathbf{z}}_n^s$ can then be written as $\gamma_n^s(i) = \frac{\alpha_n^s(i)\beta_n^s(i)}{\sum_{i'=1}^{I} \alpha_n^s(i')\beta_n^s(i')}$. To calculate $\nu_n^s(\cdot,\cdot)$, we need to first calculate $\mathbb{P}(\tilde{\mathbf{z}}_n^s = i, \tilde{\mathbf{z}}_n^{s+1} = k, \mathbf{u}_n^{\tau_{n\,1:T_n}}, \mathbf{y}_n^{\tau_{n\,1:T_n}} | \tilde{\Theta})$:

$$\mathbb{P}(\tilde{\mathbf{z}}_n^s = i, \tilde{\mathbf{z}}_n^{s+1} = k, \mathbf{u}_n^{\tau_{n\,1:T_n}}, \mathbf{y}_n^{\tau_{n\,1:T_n}} | \tilde{\Theta}) = \mathbb{P}(\mathbf{u}_n^{\tau_{n\,s+2:T_n}}, \mathbf{y}_n^{\tau_{n\,s+2:T_n}} | \tilde{\mathbf{z}}_n^{s+1} = k, \mathbf{u}_n^{\tau_{n\,s+1}}, \tilde{\Theta}) \times$$
$$\mathbb{P}(\mathbf{u}_n^{\tau_{n\,s+1}} | \mathbf{y}_n^{\tau_{n\,s+1}}, \tilde{\Theta})\, \mathbb{P}(\mathbf{y}_n^{\tau_{n\,s+1}} | \tilde{\mathbf{z}}_n^{s+1} = k, \tilde{\Theta})\, \mathbb{P}(\tilde{\mathbf{z}}_n^{s+1} = k | \tilde{\mathbf{z}}_n^s = i, \mathbf{u}_n^{\tau_{n\,s}}, \tilde{\Theta})\, \mathbb{P}(\tilde{\mathbf{z}}_n^s = i, \mathbf{u}_n^{\tau_{n\,1:s}}, \mathbf{y}_n^{\tau_{n\,1:s}}, \tilde{\Theta})$$
$$= \beta_n^{s+1}(k) \tilde{\mathcal{G}}(\mathbf{y}_n^{\tau_{n\,s+1}}, \mathbf{u}_n^{\tau_{n\,s+1}})\, \tilde{\mathcal{E}}(k, \mathbf{y}_n^{\tau_{n\,s+1}}) \left[\tilde{\mathcal{P}}_{n\,(s)}\right]_{ik} \alpha_n^s(i)$$

The boundary case will be:

$$\mathbb{P}(\tilde{\mathbf{z}}_n^{T_n-1} = i, \tilde{\mathbf{z}}_n^{T_n} = k, \mathbf{u}_n^{\tau_{n\,1:T_n}}, \mathbf{y}_n^{\tau_{n\,1:T_n}} | \tilde{\Theta}) = \tilde{\mathcal{G}}(\mathbf{y}_n^{\tau_{n\,T_n}}, \mathbf{u}_n^{\tau_{n\,T_n}})\, \tilde{\mathcal{E}}(k, \mathbf{y}_n^{\tau_{n\,T_n}}) \left[\tilde{\mathcal{P}}_{n\,(T_n-1)}\right]_{ik} \alpha_n^{T_n-1}(i)$$

The posterior probability $\mathbb{P}(\tilde{\mathbf{z}}_n^s = i, \tilde{\mathbf{z}}_n^{s+1} = k | \mathbf{u}_n^{\tau_{n\,1:T_n}}, \mathbf{y}_n^{\tau_{n\,1:T_n}}, \tilde{\Theta})$ will then be:

$$\nu_n^s(i,k) = \frac{\tilde{\mathcal{G}}(\mathbf{y}_n^{\tau_{n\,s+1}}, \mathbf{u}_n^{\tau_{n\,s+1}})\, \tilde{\mathcal{E}}(k, \mathbf{y}_n^{\tau_{n\,s+1}}) \left[\tilde{\mathcal{P}}_{n\,(s)}\right]_{ik} \alpha_n^s(i)\, \beta_n^{s+1}(k)}{\sum_{i'=1}^{I} \sum_{k'=1}^{I} \tilde{\mathcal{G}}(\mathbf{y}_n^{\tau_{n\,s+1}}, \mathbf{u}_n^{\tau_{n\,s+1}})\, \tilde{\mathcal{E}}(k', \mathbf{y}_n^{\tau_{n\,s+1}}) \left[\tilde{\mathcal{P}}_{n\,(s)}\right]_{i'\,k'} \alpha_n^s(i')\, \beta_n^{s+1}(k')} \tag{3.12}$$

The boundary case can be calculated similarly. Hence, all the probabilities that we need in the E-step can be computed efficiently.

3.2.5 Intervention Planning

Now, suppose we have a population of patients under support (e.g. through the capitation healthcare payment system) and we would like to come up with a strategy for intervention planning using the learned disease progression model from the historical data. In particular, define Ω as a measure of disutility that depends on the patient's underlying unobserved health state $\mathbf{z}$. We would like to choose an intervention planning policy $\mathcal{G}$ that provides a guideline for selecting an appropriate intervention option based on physician observation in a way that $\mathbb{E}[\Omega]$ for the entire population is minimized. We are considering a probabilistic policy model $\mathbb{P}(\mathbf{u} \,|\, \mathbf{y})$, since observing $\mathbf{y}$ does not contain all the information about the exact value of the underlying state $\mathbf{z}$. Furthermore, a stochastic policy is more appropriate from a practical perspective as it accounts for variations in intervention selection based on physician's opinion. In an ideal world that there are no resource constraints, the optimization problem can be solved for each patient independently. However, in practice each intervention requires a certain amount of resources, and the total available resources are often limited by external factors. Hence, we present an optimization problem that minimizes the expected disutility while incorporating the resource constraints. In the following, we consider the simplest version of the intervention planning problem, that is a single-period optimization problem. Specifically, we assume to have a planning period of fixed length $\Delta\tau$ and try to minimize the expected disutility at the end of the planning period.

Suppose $\mathbb{P}(\mathbf{y}, w)$ indicates the joint probability of having observation $\mathbf{y}$ and covariates w for an arbitrary patient selected from the population of patients under support. From an intuitive perspective, this probability measures the joint distribution of covariates (e.g. age) and healthcare observations (e.g. severely sick) in the population. Notice that this probability depends indirectly on the current health state of the population. The single-period intervention planning problem is equivalent to finding the optimal policy matrix $\mathcal{G}$, that indicates the probability of assigning each treatment option $\mathbf{u}$ given the observed value of $\mathbf{y}$.

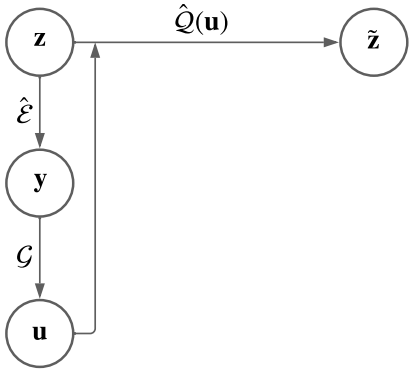

Figure 3.2: Given an observation $\mathbf{y}$, the policy matrix $\mathcal{G}$ determines intervention $\mathbf{u}$, which affects the generator matrix $\hat{\mathcal{Q}}(\mathbf{u})$. The disutility Ω is modeled as a function of the patient's underlying health state at the end of the planning period, denoted by $\tilde{\mathbf{z}}$.

Following the law of iterated expectations, we can represent the expected cost as:

$$\mathbb{E}[\Omega] = \mathbb{E}\big[\mathbb{E}[\Omega|\mathbf{y}, w]\big] = \sum_{w} \sum_{j=0}^{J-1} \mathbb{E}[\Omega|\mathbf{y} = j, w]\, \mathbb{P}(\mathbf{y} = j, w)$$

where the inner expectation is the expected cost for a patient with observation $\mathbf{y} = j$ and covariates w and the outer expectation corresponds to the entire population. Assuming that the distribution $\mathbb{P}(\mathbf{y}, w)$ is known or can be approximated reasonably accurately, one can observe that the population expected cost is essentially a deterministic linear function of the conditional expected cost $\mathbb{E}[\Omega|\mathbf{y}, w]$. Therefore, we focus on evaluating this expectation for given values of $\mathbf{y}$ and w.

We can apply law of iterated expectations one more time to further simplify the conditional expectation as $\mathbb{E}[\Omega|\mathbf{y} = j, w] = \mathbb{E}\big[\mathbb{E}[\Omega|\mathbf{z}, \mathbf{u}] \,|\, \mathbf{y} = j, w\big]$, where the outer expectation is calculated over the conditional distribution of $\mathbb{P}(\mathbf{z}, \mathbf{u}|\mathbf{y} = j)$. We assume that, given $\mathbf{y}$, there is no other factor that affects both the physician's intervention decision and the underlying health state. Therefore, $\mathbf{z} \perp \mathbf{u} \,|\, \mathbf{y}$, which implies that $\mathbb{P}(\mathbf{z}, \mathbf{u}|\mathbf{y}) = \mathbb{P}(\mathbf{z}|\mathbf{y}) \times \mathbb{P}(\mathbf{u}|\mathbf{y})$. Clearly, we have $\mathbb{P}(\mathbf{u} = l|\mathbf{y} = j) = \mathcal{G}(j, l)$. The posterior probability of the underlying health state $\mathbf{z}$

given observation $\mathbf{y}$ is equal to:

$$\mathbb{P}(\mathbf{z}=i|\mathbf{y}=j)=\frac{\mathbb{P}(\mathbf{y}=j|\mathbf{z}=i)\mathbb{P}(\mathbf{z}=i)}{\sum_{i'=1}^{I}\mathbb{P}(\mathbf{y}=j|\mathbf{z}=i')\mathbb{P}(\mathbf{z}=i')}=\frac{\hat{\mathcal{E}}(i,j)\times\hat{\pi}_i}{\sum_{i'=1}^{I}\hat{\mathcal{E}}(i',j)\times\hat{\pi}_{i'}}$$

Hence, the expected disutility can be expanded as:

$$\begin{aligned}\mathbb{E}[\Omega|\mathbf{y}=j,w]&=\sum_{i=1}^{I}\sum_{l=0}^{L-1}\mathbb{E}[\Omega|\mathbf{z}=i,\mathbf{u}=l,w]\,\mathbb{P}(\mathbf{z}=i,\mathbf{u}=l|\mathbf{y}=j)\\&=\sum_{i=1}^{I}\sum_{l=0}^{L-1}\mathbb{E}[\Omega|\mathbf{z}=i,\mathbf{u}=l,w]\,\frac{\hat{\mathcal{E}}(i,j)\times\hat{\pi}_i}{\sum_{i'=1}^{I}\hat{\mathcal{E}}(i',j)\times\hat{\pi}_{i'}}\,\mathcal{G}(j,l)\end{aligned}\tag{3.13}$$

Suppose we have estimated the parameter set $\hat{\Theta}=(\hat{\pi},\hat{\rho},\hat{\delta},\hat{\mu},\hat{\eta})$. The generator matrix for applying intervention $\mathbf{u}=l$ on a given patient with covariates w can be expressed as $[\mathcal{Q}(l)]_{ik}=\hat{\delta}_{ikl}\,\exp(\hat{\rho}_{ikl}^{'}w)$. Given $\mathbf{u}=l$ and w, the estimated generator matrix can be used for estimating the transition probability function as $\mathcal{P}_l(\Delta\tau)=\exp(\Delta\tau\mathcal{Q}(l))=\sum_{m=0}^{\infty}(\Delta\tau)^m\,\frac{\mathcal{Q}(l)^m}{m!}$. Starting from state $\mathbf{z}=i$, the probability distribution of the underlying health state at the end of planning period will then be $\mathbb{P}(\tilde{\mathbf{z}}=\tilde{i}\,|\,\mathbf{z}=i,\mathbf{u}=l)=\big[\mathcal{P}_l(\Delta\tau)\big]_{i\tilde{i}}$. For any given value of the underlying health state $\tilde{\mathbf{z}}=\tilde{i}$, it seems reasonable to approximate Ω with a known constant $\Omega_{\tilde{i}}=\mathbb{E}[\Omega|\tilde{\mathbf{z}}=\tilde{i}]$. Hence, the above expectation can be expressed as:

$$\mathbb{E}[\Omega|\mathbf{z}=i,\mathbf{u}=l,w]=\sum_{\tilde{i}=1}^{I}\Omega_{\tilde{i}}\,\mathbb{P}(\mathbf{z}\,|\,\mathbf{z}=\tilde{i},\mathbf{u}=l=\tilde{i})=\sum_{\tilde{i}=1}^{I}\Omega_{\tilde{i}}\,\big[\mathcal{P}_l(\Delta\tau)\big]_{i\tilde{i}}\tag{3.14}$$

where w implicitly affects $\mathcal{P}_l(\Delta\tau)$. Replacing equation 3.14 back in equation 3.13 will give us a simplification of the expected disutility given y and w:

$$\mathbb{E}[\Omega|\mathbf{y}=j,w]=\sum_{l=0}^{L-1}\left[\sum_{i=1}^{I}\sum_{\tilde{i}=1}^{I}\Omega_{\tilde{i}}\,\big[\mathcal{P}_l(\Delta\tau)\big]_{i\tilde{i}}\,\frac{\hat{\mathcal{E}}(i,j)\times\hat{\pi}_i}{\sum_{i'=1}^{I}\hat{\mathcal{E}}(i',j)\times\hat{\pi}_{i'}}\right]\mathcal{G}(j,l)\tag{3.15}$$

Finally, the expected disutility for the entire population under support will be:

$$\mathbb{E}[\Omega] = \sum_{j=0}^{J-1}\sum_{l=0}^{L-1}\left[\sum_{w}\sum_{i=1}^{I}\sum_{\tilde{i}=1}^{I}\Omega_{\tilde{i}}\left[\mathcal{P}_l(\Delta\tau)\right]_{i\tilde{i}}\frac{\hat{\mathcal{E}}(i,j)\times\hat{\pi}_i}{\sum_{i'=1}^{I}\hat{\mathcal{E}}(i',j)\times\hat{\pi}_{i'}}\mathbb{P}(\mathbf{y}=j,w)\right]\mathcal{G}(j,l) \qquad (3.16)$$

For any given pair of indices (j,l) the expression inside the brackets is deterministically known. Therefore, the expected disutility is essentially a linear function of the decision variables $\mathcal{G}(j,l)$ $(1\leq j\leq J,\, 1\leq l\leq L)$. It seems reasonable to model resource consumption as a linear function of the intervention variables as well. In particular, let $\mathcal{C}^r_{jl}\in\mathbb{R}$ denote the amount of resource r $(1\leq r\leq R)$ required for providing intervention $\mathbf{u}=l$ to a patient with clinical observation $\mathbf{y}=j$. Moreover, define $\mathcal{B}^r\in\mathbb{R}$ as the total amount of available resource r. The planning problem can now be formulated based on equation 3.16 as an LP problem:

$$\begin{aligned}
\min_{\mathcal{G}\geq 0}\quad & \mathbb{E}[\Omega] = \sum_{j=0}^{J-1}\sum_{l=0}^{L-1}\left[\sum_{w}\sum_{i=1}^{I}\sum_{\tilde{i}=1}^{I}\Omega_{\tilde{i}}\left[\mathcal{P}_l(\Delta\tau)\right]_{i\tilde{i}}\frac{\hat{\mathcal{E}}(i,j)\times\hat{\pi}_i}{\sum_{i'=1}^{I}\hat{\mathcal{E}}(i',j)\times\hat{\pi}_{i'}}\mathbb{P}(\mathbf{y}=j,w)\right]\mathcal{G}(j,l)\\
s.t.\quad & \sum_{j=0}^{J-1}\sum_{l=0}^{L-1}\mathcal{C}^r_{jl}\,\mathcal{G}(j,l)\leq\mathcal{B}^r; \qquad (0\leq j\leq J-1, 0\leq l\leq L-1, 1\leq r\leq R)\\
& \sum_{l=0}^{L-1}\mathcal{G}(j,l)=1; \qquad (0\leq j\leq J-1, 0\leq l\leq L-1)
\end{aligned}$$

which is solvable using off-the-shelf solvers in polynomial time.

CHAPTER 4

Experimental Results

In this section, we first describe the data that we used in our experiments. Afterward, we present the results of our analysis and provide interpretations. We've implemented our simulations using *R*, and performed the visualizations using *Python*. All the codes, along with comments, are publicly available on *GitHub*[1].

4.1 Synthetic Data Generation

We use synthetic data in our simulations. In the following, we'll describe the data generation process and present some relevant figures. More details on implementation will be later provided in the appendix.

Specifically, we generate N IID samples from the graphical model presented in figure 3.1. For each sample, we simulate the number of visits according to a rounded Normal distribution $T_n \sim round(\mathcal{N}(\mu_T, \sigma_T^2))$. We assume that we have at least two visits per patient, and eliminate samples with less than two visits. The time between consecutive visits is also assumed to be distributed as $(\tau_{n\,s+1} - \tau_{n\,s}) \sim \mathcal{N}(\mu_\tau, \sigma_\tau^2)$ for $1 \leq s \leq T_n$. Moreover, we generate an age variable for each patient that will be used in our analysis as the covariate. The age distribution is assumed to be a mixture of two Normal distributions as $a_n \sim p_a\,\mathcal{N}(\mu_{a,y}, \sigma_{a,y}^2) + (1-p_a)\,\mathcal{N}(\mu_{a,o}, \sigma_{a,o}^2)$ that has one peak for young patients and another peak for old patients.

[1] Refer to the CT-HMM directory in *https://github.com/saeedghodsi93/Disease_Progression_Modeling_HMM*

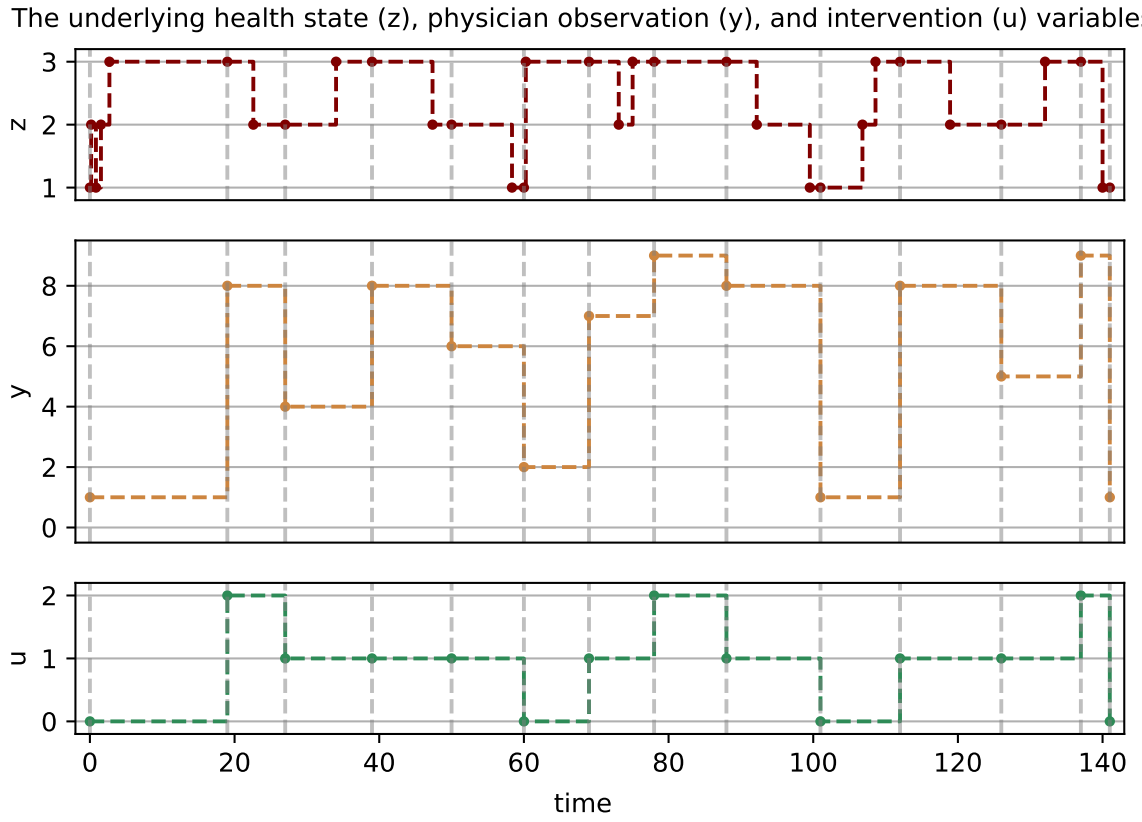

Figure 4.1: An example of the generated synthetic data. The dashed vertical lines indicate visit times. As the figure demonstrates, the underlying health state variable $\mathbf{z}$ may change between the visits, and the physician only observes a noisy signal of the snapshots of the underlying process at the visit times.

We assume to have $I = 3$ underlying health states, $J = 10$ possible values for the observation, and $L = 3$ different intervention options. We will later provide details regarding the choice of the true parameters and initial parameters of the EM algorithm in the appendix. Without loss of generality, we assume $\tau_{n1} = 0$ for all the patients. The initial underlying health state node is generated as $\mathbf{z}_n^{\tau_{n1}} \sim Cat(\pi^*)$. The first physician observation variable will then be generated according to $\mathbf{y}_n^{\tau_{n1}} \sim Cat(\mathcal{E}^*(\mathbf{z}_n^{\tau_{n1}}, \cdot))$. Similarly, we generate the first intervention variable as $\mathbf{u}_n^{\tau_{n1}} \sim Cat(\mathcal{G}^*(\mathbf{y}_n^{\tau_{n1}}, \cdot))$.

Denote the transition probability matrix of the embedded DTMC and the mean sojourn times by ζ^* and ι^*, respectively. Between two consecutive visits in any given period $s \in$

$\{1, \cdots, T_n - 1\}$, the generator matrix $\mathcal{Q}^*_{\cdot,\cdot,\mathbf{u}_n^{\tau_{ns}}}$ is known, and the parameters $\zeta^*_{\cdot,\cdot,\mathbf{u}_n^{\tau_{ns}}}$ and $\iota^*_{\cdot,\mathbf{u}_n^{\tau_{ns}}}$ can be extracted from it. Therefore, we iterate over index $\tilde{s}$, while first generating a sojourn time variable according to $(\tilde{\tau}_{n\,\tilde{s}+1} - \tilde{\tau}_{n\tilde{s}}) \sim Exp(\iota^*_{\mathbf{z}_n^{\tilde{\tau}_{n\tilde{s}}},\mathbf{u}_n^{\tau_{ns}}})$ and then generating the corresponding next state as $\mathbf{z}_n^{\tilde{\tau}_{n\,\tilde{s}+1}} \sim Cat(\zeta^*_{\mathbf{z}_n^{\tilde{\tau}_{n\tilde{s}}},\cdot,\mathbf{u}_n^{\tau_{ns}}})$ each time. This iterative procedure continues until $\tilde{\tau}_{n\,\tilde{s}+1} \geq \tau_{ns}$ for some $\tilde{s}$. At this point, we sample the next physician observation variable according to $\mathbf{y}_n^{\tau_{n\,s+1}} \sim Cat(\mathcal{E}^*(\mathbf{z}_n^{\tau_{n\,s+1}}, \cdot))$, and then generate the corresponding intervention variable as $\mathbf{u}_n^{\tau_{n\,s+1}} \sim Cat(\mathcal{G}^*(\mathbf{y}_n^{\tau_{n\,s+1}}, \cdot))$. Figure 4.1 presents an example of the samples that we've generated following this procedure.

4.2 Results

In this section, we report the results of our experiments on a synthetic dataset that we generated using the above approach. We'll later present a detailed version of the estimation algorithm as well as the true and initial model parameters in the appendix.

To measure the performance of our algorithm, we take the estimated parameters in each iteration and calculate the Root Mean Squared Error (RMSE) for each parameter based on the distance between the estimated parameter and the true parameter value. We plot RMSE for different sample sizes as a function of the EM iteration in figure 4.2. Although there is an initial jump in the error terms associated with some of the parameters, we observe that all the estimation errors would eventually converge to zero if the sample size and the number of iterations are large enough. Notice that the scale of the errors vary due to the difference between the scale of the parameters and the difference between the quality of the initial points that we've chosen.

As the results indicate, $\hat{\pi}$ and $\hat{\mu}$ converge to their true values even for relatively small sample sizes. For the intervention model, we estimate its parameter values as $\hat{\eta}_j = \frac{\sum_{n=1}^{N}\sum_{s=1}^{T_n}\mathbf{u}_n^{\tau_{ns}}\mathbb{I}\{\mathbf{y}_n^{\tau_{ns}}=j\}}{(L-1)\sum_{n=1}^{N}\sum_{s=1}^{T_n}\mathbb{I}\{\mathbf{y}_n^{\tau_{ns}}=j\}}$, which only depends on variables $\mathbf{y}$ and $\mathbf{u}$ that are both observed. In other words, we practically don't update this parameter in the EM iterations and just use the

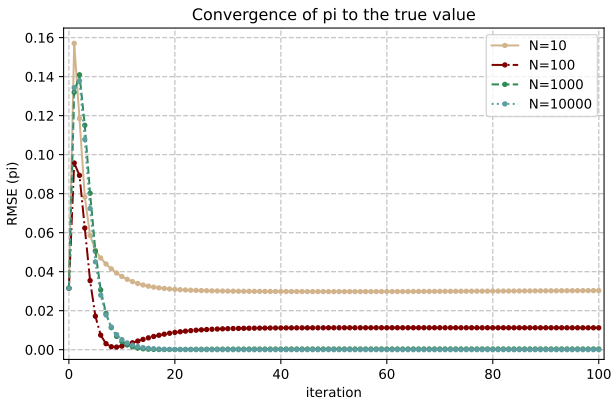

(a) Convergence of $\hat{\pi}$

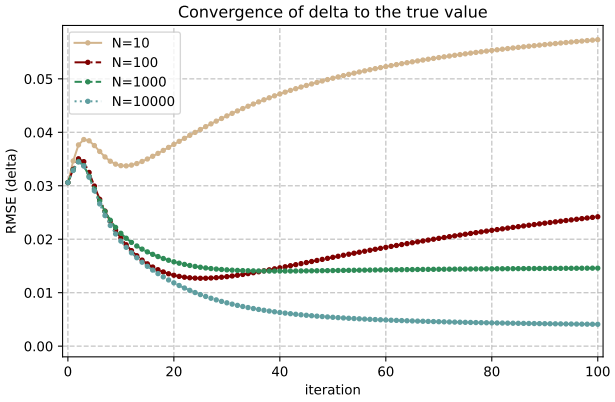

(b) Convergence of $\hat{\delta}$

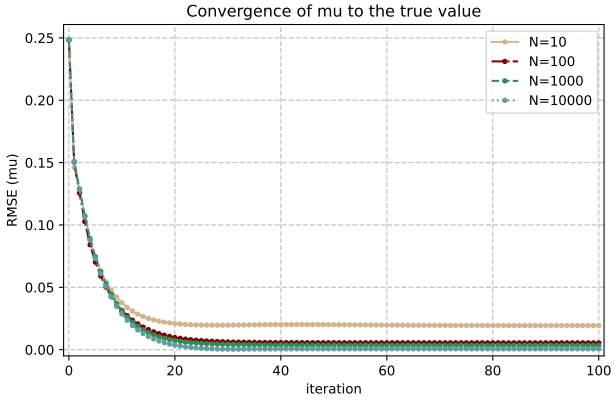

(c) Convergence of $\hat{\mu}$

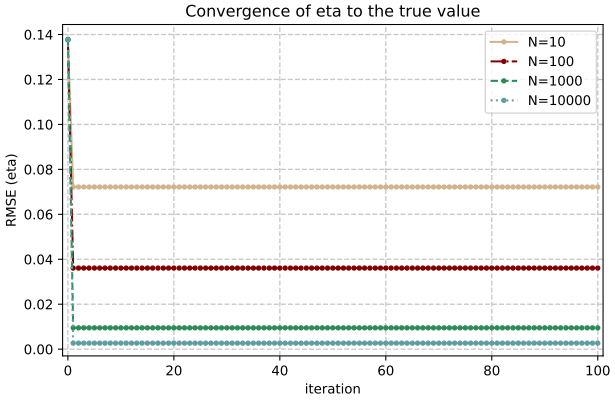

(d) Convergence of $\hat{\eta}$

Figure 4.2: Convergence of the EM estimated parameters to their true values for different sample sizes.

standard mean estimator for any given j. Hence, we observe that the $\hat{\eta}_j$ errors are constant as we make more iterations and the quality of the model fit depends only on the number of samples. For the generator matrix parameters, the algorithm seems not to be initially converging to the true values when the sample size is too small. However, having a large dataset, the parameter values converge to the true values after enough number of iterations have been passed. Needless to explain, there are many different methods for improving the convergence rates of the EM algorithm that we've not discussed here [VR08].

CHAPTER 5

Conclusion

In this thesis, we provided a model for disease progression that incorporates the effect of interventions, and presented an efficient approach for learning its parameters. In particular, we first introduced the disease progression modeling and data-driven intervention planning problems and discussed the motivations behind our research study in Chapter 1. We then explained the challenges and complexities of disease progression modeling as well as different model classes that have been proposed for this purpose in the past, including DTMCs, CTMCs, HMMs, and POMDPs. We claimed that CT-HMMs are a good choice for disease progression modeling, due to their capability of handling discrete-time irregularly-spaced observations and modeling latent state spaces. Afterward, we briefly reviewed the relevant literature and justified our choice of modeling framework in Chapter 2. In Chapter 3, we first provided a basic version of the CT-HMM model and discussed a recently proposed approach for estimating the parameters. We then modified the model in order to incorporate the effect of interventions and presented our estimation algorithm. Finally, we explained our synthetic data generation approach and presented the experimental results in Chapter 4.

Parameter estimation for CT-HMMs with discrete-time and irregular observations is computationally challenging, even when potential interventions do not affect the generator matrix. As a baseline approach, the MCEM algorithm can be used for estimating the model parameters. More specifically, we need to generate a set of samples from the joint distribution of our graphical model (e.g. using the Gibbs sampling algorithm), and then we can approximate the expected number of transitions and expected sojourn times that we need in the E-step by

their sample averages. However, this approach is computationally challenging as the number of Monte Carlo samples that we need to generate in order to have an approximation with small variance is usually very large in this case. As we explained, evaluation of the aforementioned expectations can be performed by applying eigendecomposition on the generator matrix. However, the generator matrix is often not diagonalizable during the learning procedure. Therefore, we decided to build our algorithm based on the Expm method that [LLL15] has recently proposed.

As we explained, our main contribution is to present a modified CT-HMM disease progression model that naturally incorporates the effect of interventions and allows for performing intervention planning using the learned model. Specifically, we parameterize the elements of the generator matrix according to the Cox proportional hazards function to model the impact of the intervention variables and covariates on them. We then derived an EM algorithm for iteratively estimating the model parameters. In the E-step, we decompose the complete-data log-likelihood and represent it in terms of a set of posterior probabilities and end state conditioned expectations. The probabilities are then evaluated using a forward-backward algorithm, and the expectations are calculated using the aforementioned approach. We demonstrate the effectiveness of our algorithm by experimenting on a synthetically generated dataset.

As the last part of our work, we developed an optimization problem for intervention planning using our model. In particular, we assumed that the emission and transition models are estimated using data, and try to find an intervention model that minimizes the overall system costs subject to a set of resource constraints. We assume that the cost associated with being in any certain underlying health state is fixed and known. Afterward, we calculate the probability of each underlying health state as a function of model parameters and further estimate the total cost. We then show that if the resource consumption functions are linear, the optimization problem is an LP, which is solvable in polynomial time.

Appendix A

Implementation Details

In this section, we present the details of our simulations as well as the algorithm for estimating the model parameters, based on our discussions in the previous chapters. We ran the experiments on a 2.2 GHz Intel Core-i7 8750H microprocessor with 6 cores, and the computations were completed in less than one hour on our machine for the moderate sample size of $N = 1,000$. For the largest dataset that we used in our experiments (i.e. $N = 10,000$), the entire 100 iterations of the EM algorithm were finished in almost eight hours.

A.1 Model Parameters

First, we choose the mean and standard deviations of the number of visits per patient, and the between visit times as $\mu_T = 10$, $\sigma_T = 3$, $\mu_\tau = 10$, $\sigma_\tau = 3$. Similarly, we set the age parameters to $\mu_{a,y} = 25$, $\sigma_{a,y} = 5$, $\mu_{a,o} = 50$, $\sigma_{a,o} = 10$, $p_a = 0.25$. We then set the true initial health state distribution parameter as $\pi^* = (0.25, 0.45, 0.30)$. To set the generator matrix, we construct the transition probability matrix of the embedded DTMC in the following way:

$$\zeta^*_{\cdot,\cdot,0} = \begin{bmatrix} 0.00 & 0.60 & 0.40 \\ 0.20 & 0.00 & 0.80 \\ 0.10 & 0.90 & 0.00 \end{bmatrix}, \qquad \zeta^*_{\cdot,\cdot,1} = \begin{bmatrix} 0.00 & 0.80 & 0.20 \\ 0.60 & 0.00 & 0.40 \\ 0.30 & 0.70 & 0.00 \end{bmatrix}, \qquad \zeta^*_{\cdot,\cdot,2} = \begin{bmatrix} 0.00 & 0.90 & 0.10 \\ 0.80 & 0.00 & 0.20 \\ 0.50 & 0.50 & 0.00 \end{bmatrix}$$

where the first two indices iterate over $i \in \{1, \cdots, I\}$, $k \in \{1, \cdots, I\}$, and the last index iterates over the set of intervention options $l \in \{0, \cdots, L-1\}$. We also set the mean sojourn

times associated with each of the states as:

$$\iota^*_{\cdot,0} = \begin{bmatrix} 4 & 6 & 20 \end{bmatrix}, \qquad \iota^*_{\cdot,1} = \begin{bmatrix} 9 & 12 & 10 \end{bmatrix}, \qquad \iota^*_{\cdot,2} = \begin{bmatrix} 11 & 10 & 7 \end{bmatrix}$$

where the first index iterates over $i \in \{1, \cdots, I\}$. The true baseline generator matrix elements can then be constructed as:

$$\delta^*_{\cdot,\cdot,0} = \begin{bmatrix} - & 0.15 & 0.10 \\ 0.03 & - & 0.13 \\ 0.00 & 0.04 & - \end{bmatrix}, \quad \delta^*_{\cdot,\cdot,1} = \begin{bmatrix} - & 0.09 & 0.02 \\ 0.05 & - & 0.03 \\ 0.03 & 0.07 & - \end{bmatrix}, \quad \delta^*_{\cdot,\cdot,2} = \begin{bmatrix} - & 0.08 & 0.01 \\ 0.08 & - & 0.02 \\ 0.07 & 0.07 & - \end{bmatrix}$$

We set the true covariate coefficients in the following way:

$$\rho^*_{\cdot,\cdot,0} = \rho^*_{\cdot,\cdot,1} = \rho^*_{\cdot,\cdot,2} = \begin{bmatrix} - & 0.020 & 0.020 \\ 0.015 & - & 0.015 \\ 0.010 & 0.010 & - \end{bmatrix}$$

By choosing these parameters, we're effectively making the sojourn times associated with healthier states shorter as age increases, while keeping the sojourn times of the less healthy states relatively constant. Although our formulation allows for modeling more complex relationships, we decided to choose this certain functional form due to its practical relevance. For example, for a 50-years old patient, the true generator matrix will be:

$$\mathcal{Q}^*_{\cdot,\cdot,0} = \begin{bmatrix} -0.68 & 0.41 & 0.27 \\ 0.07 & -0.35 & 0.28 \\ 0.01 & 0.07 & -0.08 \end{bmatrix}, \qquad \mathcal{Q}^*_{\cdot,\cdot,1} = \begin{bmatrix} -0.30 & 0.24 & 0.06 \\ 0.11 & -0.18 & 0.07 \\ 0.05 & 0.12 & -0.16 \end{bmatrix}$$

$$\mathcal{Q}^*_{\cdot,\cdot,2} = \begin{bmatrix} -0.25 & 0.22 & 0.02 \\ 0.17 & -0.21 & 0.04 \\ 0.12 & 0.12 & -0.24 \end{bmatrix}$$

Furthermore, we set the emission probabilities as $\mu^* = (0.1, 0.5, 0.9)$. Ultimately, the intervention probabilities, associated with different observations will be set to:

$$\eta^* = (0.04, 0.15, 0.26, 0.32, 0.43, 0.51, 0.62, 0.77, 0.81, 0.90)$$

As our initial parameters for the EM algorithm, we set $\tilde{\pi} = (0.33, 0.33, 0.33)$. Moreover, we choose the transition probability matrix of the embedded DTMC as well as the mean sojourn times in the following way:

$$\tilde{\zeta}_{\cdot,\cdot,0} = \begin{bmatrix} 0.00 & 0.50 & 0.50 \\ 0.30 & 0.00 & 0.70 \\ 0.20 & 0.80 & 0.00 \end{bmatrix}, \qquad \tilde{\zeta}_{\cdot,\cdot,1} = \begin{bmatrix} 0.00 & 0.60 & 0.40 \\ 0.40 & 0.00 & 0.60 \\ 0.40 & 0.60 & 0.00 \end{bmatrix}, \qquad \tilde{\zeta}_{\cdot,\cdot,2} = \begin{bmatrix} 0.00 & 0.70 & 0.30 \\ 0.60 & 0.00 & 0.40 \\ 0.40 & 0.60 & 0.00 \end{bmatrix}$$

$$\tilde{\iota}_{\cdot,0} = \begin{bmatrix} 5 & 8 & 12 \end{bmatrix}, \qquad \tilde{\iota}_{\cdot,1} = \begin{bmatrix} 6 & 10 & 10 \end{bmatrix}, \qquad \tilde{\iota}_{\cdot,2} = \begin{bmatrix} 9 & 8 & 4 \end{bmatrix}$$

Hence, the corresponding baseline generator matrix elements will be:

$$\tilde{\delta}_{\cdot,\cdot,0} = \begin{bmatrix} - & 0.10 & 0.10 \\ 0.04 & - & 0.09 \\ 0.02 & 0.07 & - \end{bmatrix}, \quad \tilde{\delta}_{\cdot,\cdot,1} = \begin{bmatrix} - & 0.10 & 0.07 \\ 0.04 & - & 0.06 \\ 0.04 & 0.06 & - \end{bmatrix}, \quad \tilde{\delta}_{\cdot,\cdot,2} = \begin{bmatrix} - & 0.08 & 0.03 \\ 0.07 & - & 0.05 \\ 0.10 & 0.15 & - \end{bmatrix}$$

Moreover, we start with the following covariate coefficient that are uninformative about the difference between the health states:

$$\tilde{\rho}_{\cdot,\cdot,0} = \tilde{\rho}_{\cdot,\cdot,1} = \tilde{\rho}_{\cdot,\cdot,2} = \begin{bmatrix} - & 0.01 & 0.01 \\ 0.01 & - & 0.01 \\ 0.01 & 0.01 & - \end{bmatrix}$$

Finally, set $\tilde{\mu} = (0.45, 0.65, 0.7)$ and $\tilde{\eta} = (0.01, 0.03, 0.06, 0.08, 0.45, 0.48, 0.51, 0.55, 0.91, 0.97)$.

A.2 The Parameter Estimation Algorithm

In the following, we provide the sketch of the algorithm that we've implemented for estimating the model parameters.

Algorithm 1: The EM algorithm for learning the CT-HMM parameters

Input: Initial parameters $\tilde{\Theta} = (\tilde{\pi}, \tilde{\delta}, \tilde{\rho}, \tilde{\mu}, \tilde{\eta})$
Output: Estimated parameters $\hat{\Theta} = (\hat{\pi}, \hat{\delta}, \hat{\rho}, \hat{\mu}, \hat{\eta})$
Data: A set of N IID samples, each including $(\mathbf{y}_n^{\tau_{ns}}, \mathbf{u}_n^{\tau_{ns}})\big|_{1\leq s\leq T_n}$ for $1 \leq n \leq N$
while *not converged* **do**
E-step: for $1 \leq n \leq N$

- Calculate $\alpha_n^s(i)$ using equations 3.8, 3.10 and $\beta_n^s(i)$ using equations 3.9, 3.11
- Calculate $\gamma_n^s(i) = \frac{\alpha_n^s(i)\beta_n^s(i)}{\sum_{i'=1}^{I}\alpha_n^s(i')\beta_n^s(i')}$ and $\nu_n^s(i,k)$ based on equation 3.12
- Calculate $\mathbb{E}[\chi_{n\,ik}^{\tau_n(s)}|\mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}, \tilde{\Theta})]$ and $\mathbb{E}[\psi_{n\,i}^{\tau_n(s)}|\mathbf{u}_n^{\tau_n\,1:T_n}, \mathbf{y}_n^{\tau_n\,1:T_n}, \tilde{\Theta}]$

M-step:

- Set $\tilde{\pi}_i = \frac{\sum_{n=1}^{N}\gamma_n^1(i)}{\sum_{i'=1}^{I}\sum_{n=1}^{N}\gamma_n^1(i')}$, $\tilde{\mu}_i = \frac{\sum_{n=1}^{N}\sum_{s=1}^{T_n}\mathbf{y}_n^{\tau_{ns}}\gamma_n^s(i)}{(J-1)\sum_{n=1}^{N}\sum_{s=1}^{T_n}\gamma_n^s(i)}$, $\tilde{\eta}_j = \frac{\sum_{n=1}^{N}\sum_{s=1}^{T_n}\mathbf{u}_n^{\tau_{ns}}\mathbb{I}\{\mathbf{y}_n^{\tau_{ns}}=j\}}{(L-1)\sum_{n=1}^{N}\sum_{s=1}^{T_n}\mathbb{I}\{\mathbf{y}_n^{\tau_{ns}}=j\}}$
- Iteratively update $\tilde{\delta}_{ikl}$ using equation 3.5 and $\tilde{\rho}_{ikl}$ using equations 3.6, 3.7

return $\hat{\pi} = \tilde{\pi}$, $\hat{\delta} = \tilde{\delta}$, $\hat{\rho} = \tilde{\rho}$, $\hat{\mu} = \tilde{\mu}$, $\hat{\eta} = \tilde{\eta}$

As part of the E-step, we need to calculate the end state-conditioned expectations, as we discussed in the previous chapters. For a specific patient n and time period $(\tau_{ns}, \tau_{n\,s+1}]$, the generator matrix $\tilde{\mathcal{Q}}_{n\,(s)}$ can be constructed by setting $l = \mathbf{u}_n^{\tau_{ns}}$ in equations 3.3, 3.4. Therefore, for all the pairs (i,k), we evaluate the integral $\xi_{ik} = \int_0^{\tau_{n\,s+1}-\tau_{ns}} \exp\left(x\tilde{\mathcal{Q}}_{n\,(s)}\right)\mathcal{B}_{ik}\exp\left((\tau_{n\,s+1}-\tau_{ns}-x)\tilde{\mathcal{Q}}_{n\,(s)}\right)dx$, where all the elements of $\mathcal{B}_{ik}$ are zero except for the (i,k)-th element. Specifically, we define a new $2I \times 2I$ matrix as $\mathcal{A}_{ik} = \begin{pmatrix} \tilde{\mathcal{Q}}_{n\,(s)} & \mathcal{B}_{ik} \\ 0 & \tilde{\mathcal{Q}}_{n\,(s)} \end{pmatrix}$ and evaluate the matrix exponential $\exp\left((\tau_{n\,s+1}-\tau_{ns})\mathcal{A}_{ik}\right)$. The upper right $I \times I$ corner of this matrix will be equal to the matrix ξ_{ik}. Assuming the end states are $\tilde{i}$ and $\tilde{k}$, the (i,k)-th expectations in equations 3.1, 3.2 can be determined by using the $(\tilde{i},\tilde{k})$-th element of the matrix ξ_{ik}.

REFERENCES

[AHS17] Ahmed M Alaa, Scott Hu, and Mihaela Schaar. "Learning from clinical judgments: Semi-markov-modulated marked hawkes processes for risk prognosis." In *International Conference on Machine Learning*, pp. 60–69. PMLR, 2017.

[AKG12] Hossein Haji Ali Afzali, Jonathan Karnon, and Jodi Gray. "A critical review of model-based economic studies of depression." *Pharmacoeconomics*, **30**(6):461–482, 2012.

[AS19] Ahmed M Alaa and Mihaela van der Schaar. "Attentive State-Space Modeling of Disease Progression." *Advances in Neural Information Processing Systems*, **32**:11338–11348, 2019.

[ASK20] Bastian Alt, Matthias Schultheis, and Heinz Koeppl. "POMDPs in continuous time and discrete spaces." *Advances in Neural Information Processing Systems*, **33**:13151–13162, 2020.

[AV18] Ahmed M Alaa and Mihaela Van Der Schaar. "A hidden absorbing semi-Markov model for informatively censored temporal data: Learning and inference." *The Journal of Machine Learning Research*, **19**(1):108–169, 2018.

[BF94] Yoshua Bengio and Paolo Frasconi. "An input output HMM architecture." *Advances in neural information processing systems*, **7**, 1994.

[CB16] Sarah F Cook and Robert R Bies. "Disease progression modeling: Key concepts and recent developments." *Current pharmacology reports*, **2**(5):221–230, 2016.

[CBS12] J Jaime Caro, Andrew H Briggs, Uwe Siebert, and Karen M Kuntz. "Modeling good research practices—overview: a report of the ISPOR-SMDM Modeling Good Research Practices Task Force–1." *Medical Decision Making*, **32**(5):667–677, 2012.

[CM14] Bibhas Chakraborty and Susan A Murphy. "Dynamic treatment regimes." *Annual review of statistics and its application*, **1**:447–464, 2014.

[Cox72] David R Cox. "Regression models and life-tables." *Journal of the Royal Statistical Society: Series B (Methodological)*, **34**(2):187–202, 1972.

[DIJ13] Sarang Deo, Seyed Iravani, Tingting Jiang, Karen Smilowitz, and Stephen Samuelson. "Improving health outcomes through better capacity allocation in a community-based chronic care model." *Operations Research*, **61**(6):1277–1294, 2013.

[FHD20] Joseph Futoma, Michael Hughes, and Finale Doshi-Velez. "POPCORN: Partially Observed Prediction Constrained Reinforcement Learning." In *International Conference on Artificial Intelligence and Statistics*, pp. 3578–3588. PMLR, 2020.

[FKS21] Mehdi Fatemi, Taylor W Killian, Jayakumar Subramanian, and Marzyeh Ghassemi. "Medical Dead-ends and Learning to Identify High-risk States and Treatments." *Advances in Neural Information Processing Systems*, **34**, 2021.

[FMD20] Joseph Futoma, Muhammad A Masood, and Finale Doshi-Velez. "Identifying distinct, effective treatments for acute hypotension with SODA-RL: safely optimized diverse accurate reinforcement learning." *AMIA Summits on Translational Science Proceedings*, **2020**:181, 2020.

[GJM18] Omer Gottesman, Fredrik Johansson, Joshua Meier, Jack Dent, Donghun Lee, Srivatsan Srinivasan, Linying Zhang, Yi Ding, David Wihl, Xuefeng Peng, et al. "Evaluating reinforcement learning algorithms in observational health settings." *arXiv preprint arXiv:1805.12298*, 2018.

[GLL94] RC Gentleman, JF Lawless, JC Lindsey, and P Yan. "Multi-state Markov models for analysing incomplete disease history data with illustrations for HIV disease." *Statistics in medicine*, **13**(8):805–821, 1994.

[HJ05] Asger Hobolth and Jens Ledet Jensen. "Statistical inference in evolutionary models of DNA sequences via the EM algorithm." *Statistical applications in genetics and molecular biology*, **4**(1), 2005.

[HJ11] Asger Hobolth and Jens Ledet Jensen. "Summary statistics for endpoint-conditioned continuous-time Markov chains." *Journal of applied probability*, **48**(4):911–924, 2011.

[Jac11] Christopher Jackson. "Multi-state models for panel data: the msm package for R." *Journal of statistical software*, **38**(1):1–28, 2011.

[KCB18] Matthieu Komorowski, Leo A Celi, Omar Badawi, Anthony C Gordon, and A Aldo Faisal. "The artificial intelligence clinician learns optimal treatment strategies for sepsis in intensive care." *Nature medicine*, **24**(11):1716–1720, 2018.

[KL85] JD Kalbfleisch and Jerald Franklin Lawless. "The analysis of panel data under a Markov assumption." *Journal of the american statistical association*, **80**(392):863–871, 1985.

[KVM11] Beste Kucukyazici, Vedat Verter, and Nancy E Mayo. "An analytical framework for designing community-based care for chronic diseases." *Production and Operations Management*, **20**(3):474–488, 2011.

[LC01] Richard A Levine and George Casella. "Implementations of the Monte Carlo EM algorithm." *Journal of Computational and Graphical Statistics*, **10**(3):422–439, 2001.

[LKT20] Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. "Offline reinforcement learning: Tutorial, review, and perspectives on open problems." *arXiv preprint arXiv:2005.01643*, 2020.

[LLL15] Yu-Ying Liu, Shuang Li, Fuxin Li, Le Song, and James M Rehg. "Efficient Learning of Continuous-Time Hidden Markov Models for Disease Progression." *Advances in Neural Information Processing Systems*, **28**:3600–3608, 2015.

[LLV19] Elliot Lee, Mariel S Lavieri, and Michael Volk. "Optimal screening for hepatocellular carcinoma: A restless bandit model." *Manufacturing & Service Operations Management*, **21**(1):198–212, 2019.

[LM13] Jane M Lange and Vladimir N Minin. "Fitting and interpreting continuous-time latent Markov models for panel data." *Statistics in medicine*, **32**(26):4581–4595, 2013.

[LSS20] MingYu Lu, Zachary Shahn, Daby Sow, Finale Doshi-Velez, and Li-wei H Lehman. "Is Deep Reinforcement Learning Ready for Practical Applications in Healthcare? A Sensitivity Analysis of Duel-DDQN for Hemodynamic Management in Sepsis Patients." In *AMIA Annual Symposium Proceedings*, volume 2020, p. 773. American Medical Informatics Association, 2020.

[MHS07] Philipp Metzner, Illia Horenko, and Christof Schütte. "Generator estimation of Markov jump processes based on incomplete observations nonequidistant in time." *Physical Review E*, **76**(6):066702, 2007.

[Mou12] DR Mould. "Models for disease progression: new approaches and uses." *Clinical Pharmacology & Therapeutics*, **92**(1):125–131, 2012.

[PDW18] Xuefeng Peng, Yi Ding, David Wihl, Omer Gottesman, Matthieu Komorowski, Li-wei H Lehman, Andrew Ross, Aldo Faisal, and Finale Doshi-Velez. "Improving sepsis treatment strategies by combining deep and kernel-based reinforcement learning." In *AMIA Annual Symposium Proceedings*, volume 2018, p. 887. American Medical Informatics Association, 2018.

[PFU18] Ryan Palmer, Naomi J Fulop, and Martin Utley. "A systematic literature review of operational research methods for modelling patient flow and outcomes within community healthcare and other settings." *Health Systems*, **7**(1):29–50, 2018.

[RKA17] Aniruddh Raghu, Matthieu Komorowski, Imran Ahmed, Leo Celi, Peter Szolovits, and Marzyeh Ghassemi. "Deep reinforcement learning for sepsis treatment." *arXiv preprint arXiv:1711.09602*, 2017.

[SCS14] Lachlan Standfield, Tracy Comans, and Paul Scuffham. "Markov modeling and discrete event simulation in health care: a systematic comparison." *International journal of technology assessment in health care*, **30**(2):165, 2014.

[SCS20] Kristen A Severson, Lana M Chahine, Luba Smolensky, Kenney Ng, Jianying Hu, and Soumya Ghosh. "Personalized Input-Output Hidden Markov Models for Disease Progression Modeling." In *Machine Learning for Healthcare Conference*, pp. 309–330. PMLR, 2020.

[SGL19] Zhaonan Sun, Soumya Ghosh, Ying Li, Yu Cheng, Amrita Mohan, Cristina Sampaio, and Jianying Hu. "A probabilistic disease progression modeling approach and its application to integrated Huntington's disease observational data." *JAMIA open*, **2**(1):123–130, 2019.

[SLL11] Susan M Shortreed, Eric Laber, Daniel J Lizotte, T Scott Stroup, Joelle Pineau, and Susan A Murphy. "Informing sequential clinical decision-making through reinforcement learning: an empirical study." *Machine learning*, **84**(1-2):109–136, 2011.

[Van78] Charles Van Loan. "Computing integrals involving the matrix exponential." *IEEE transactions on automatic control*, **23**(3):395–404, 1978.

[VR08] Ravi Varadhan and Christophe Roland. "Simple and globally convergent methods for accelerating the convergence of any EM algorithm." *Scandinavian Journal of Statistics*, **35**(2):335–353, 2008.

[WSW14] Xiang Wang, David Sontag, and Fei Wang. "Unsupervised learning of disease progression models." In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 85–94, 2014.

[YLN21] Chao Yu, Jiming Liu, Shamim Nemati, and Guosheng Yin. "Reinforcement learning in healthcare: A survey." *ACM Computing Surveys (CSUR)*, **55**(1):1–36, 2021.

[ZWD21] Kristine Zhang, Yuanheng Wang, Jianzhun Du, Brian Chu, Leo Anthony Celi, Ryan Kindle, and Finale Doshi-Velez. "Identifying Decision Points for Safe and Interpretable Reinforcement Learning in Hypotension Treatment." *arXiv preprint arXiv:2101.03309*, 2021.