

Stage-Aware Learning for Dynamic Treatments

Hanwen Ye

*Department of Statistics
University of California
Irvine, CA 92617, USA*

HANWENY@UCI.EDU

Wenzhuo Zhou

*Department of Statistics
University of California
Irvine, CA 92617, USA*

WENZHUZ3@UCI.EDU

Ruoqing Zhu

*Department of Statistics
University of Illinois
Urbana-Champaign, IL 61820, USA*

RQZHU@ILLINOIS.EDU

Annie Qu

*Department of Statistics
University of California
Irvine, CA 92617, USA*

AQU2@UCI.EDU

Abstract

Recent advances in dynamic treatment regimes (DTRs) provide convenient optimal treatment searching algorithms, which are tailored to individuals' specific needs and able to maximize their expected clinical benefits. However, existing algorithms could suffer from insufficient sample size under optimal treatments, especially for chronic diseases involving long sequences of decision-making. To address the aforementioned practical challenges, we propose a novel individualized learning method that estimates the DTR with a focus on prioritizing alignment percentages between the observed and underlying optimal regime across decision stages. By relaxing the restriction that observed samples must be fully aligned with the optimal regime at all treatment stages, our approach substantially improves the sample efficiency of inverse probability weighted based approaches. In particular, the proposed learning scheme builds a more general framework which includes the popular outcome weighted learning framework as a special case of ours. Moreover, we introduce the notion of stage importance scores along with an attention-based deep learning architecture to explicitly account for heterogeneity among decision stages. We establish the theoretical properties of the proposed method, including the Fisher consistency and finite-sample performance bound. The proposed approach is evaluated through extensive simulations and a real-world application to the UC COVID-19 dataset.

Keywords: Attention mechanism; Efficient learning; Health analytics; Individualized treatment; Recommender systems

1 Introduction

There has been great interest and demand for individualized modeling and personalized prediction, with applications ranging from medicine to education programs and marketing. For instance, the outbreak of COVID-19 in recent years has highlighted the growing demand

for developing an effective dynamic treatment which can be tailored to individual patients. Dynamic treatment regime (DTR) (Tsiatis et al., 2019), as an emerging individualized treatment approach, has thus found much attention in the medical field. In contrast to traditional, one-size-fits-all medical treatments, DTR continuously adapts a patient’s treatment plan based on their personal response to previous treatments and changes in their medical conditions. Since different treatment regimes could lead to contrasting clinical outcomes, the main goal of current medical care is to find the optimal DTR that would yield the largest expected long-term benefits for each patient (Rubin, 1974; Robins, 1986).

However, searching for the optimal regime in practice is not an intuitive task, commonly due to limited clinical data, complex heterogeneity among patients, and a combinatorial number of possible treatments involved in a long sequence of decisions. Although extensive studies have been conducted to establish the optimal DTR searching methodology, the developed algorithms could still suffer from these empirical challenges. Current search algorithms typically can be categorized into two main frameworks: indirect and direct value-search. The indirect methods, such as Q-learning (Watkins and Dayan, 1992; Nahum-Shani et al., 2012), A-learning (Murphy, 2003; Blatt et al., 2004; Shi et al., 2018), and tree-based methods (Laber and Zhao, 2015; Tao et al., 2018) primarily model the conditional distribution of a clinical outcome given the patients’ past covariates and treatment information, and choose the treatment that maximizes the modeled outcome as optimal. However, as the selection process follows backward induction, an optimal regime might not be recovered if one of the outcome models is not correctly specified (Schulte et al., 2014). The situation only gets worse in a chronic disease setting which involves a long sequence of decision stages. Though one can formulate the outcome model via a semi- or non-parametric approach (Ernst et al., 2005; Geurts et al., 2006; Zhao et al., 2009) to allow more model flexibility and mitigate the risk of misspecification, the fitted models oftentimes are hard to interpret and thus less appealing for decision-makers to apply.

On the other hand, value search methods include outcome weighted learning (Zhao et al., 2012, 2015), residual weighted learning (Zhou et al., 2017), robust estimators (Zhang et al., 2012a,b, 2013; Shi et al., 2016; Schulz and Moodie, 2021), distributional learning (Mo et al., 2021), and angle-based learning (Qi et al., 2020; Xue et al., 2022). These methods directly posit a model class to the treatment regime and maximize the expected outcome under a fitted regime. The optimal regime is defined as the maximizer over the pre-specified function class. However, depending on the choice of density correction technique, the value search methods could have vastly different empirical performances. For instance, outcome weighted learning maximizes the inverse-probability weighted estimator (IPWE) of the expected outcome, which is known to be unstable when there is an insufficient number of optimal instances observed in the collected sample (Zhang et al., 2012a,b). Unfortunately, such an empirical dilemma can be commonly found in real-world applications where the optimal treatment for a new disease is barely studied, or the optimal treatment cannot be assigned to a vast patient population due to various reasons. To improve the stability of the IPWE, doubly-robust estimators (Zhang et al., 2013) augment a mean-zero term to capture information from patients who receive non-optimal treatments. However, the additional augmented term introduces an extra amount of computational burden, and there still is no guarantee of either the propensity or the outcome model being correctly modeled, especially when the regime involves a large number of decision steps.

In this paper, we propose a novel DTR estimation method, namely Stage Aware Learning (SAL), which aims to enhance the alignment between the observed treatments and optimal regime, while maximizing the expected outcome. The key idea is that treatment variations between the observed and the optimal regime are now allowed to increase data efficiency. Thus, instead of only seeking optimal regimes among patients assigned with optimal treatments at all decision stages, our approach includes patients treated under all strategies and places particular emphasis on those whose observed treatments align more closely with the optimal ones. In addition, to better capture the difference in treatment effectiveness at varying stages, we introduce the notion of stage importance scores and further propose the Stage Weighted Learning (SWL) method based on SAL.

The proposed SAL has several unique advantages. First, all empirical clinical data are incorporated into the value search process no matter how many stages of the observed treatments are aligned with the optimal regime. Second, compared to the existing value search methods, the proposed method increases the stability of the IPWE-based approach and reduces the computational burden due to the augmentation process of the robust estimators. Third, our method indicates a more general framework that includes the outcome weighted learning as a special case and can be adapted to various empirical settings. In addition, SAL implies that heterogeneity among decision stages can be explicitly included in the estimation procedure which further facilitates estimating the optimal regime.

The main contributions of our paper are as follows. First, our work bridges the gap between DTR searching algorithms and empirical challenges. Specifically, we overcome the difficulties of applying IPWEs in empirical settings where few optimal instances have been observed. Second, to the best of our knowledge, this is the first work combining a number of IPWEs under different treatment matching scenarios in value search algorithms to estimate optimal regimes. The flexibility of combining and weighing each IPWEs increases the generalizability of DTR estimation methods to any empirical setting. Third, our work builds the theoretical connection between multi-stage DTR search problems and multi-label classification problems to simplify the optimal regime estimation procedure.

The remainder of the paper is structured as follows. In Section 2, we introduce the notations and background of outcome weighted learning. In Section 3, we propose a novel k -IPWE estimator and introduce the SAL and SWL methods to account for stage heterogeneity. Section 4 presents Fisher consistencies and finite-sample performance error bound of the proposed SWL method, and Section 5 explains the computation algorithm and implementation details. In Section 6, extensive simulation results are presented to illustrate the empirical performance advantages of our proposed methods. In Section 7, we apply the proposed methods to the COVID-19 data from UC hospitals. Lastly, we conclude with discussions in Section 8.

2 Background and Related Works

In this section, we introduce the necessary notations and assumptions used in the paper, and formulate the multi-stage DTR estimation procedure under the Inverse Probability Weighed Estimators (IPWE) (Horvitz and Thompson, 1952; Robins et al., 1994) framework. In addition, we will briefly discuss Outcome Weighted Learning (OWL) (Zhao et al., 2012, 2015), which is one of the most representative learning schemes based on IPWE.

2.1 Notation and Preliminary

Consider a balanced multistage decision setting where all patients in the study have a total number of T stages (visits). For each patient at their j^{th} clinic visit, where $j = 1, \dots, T$, a set of time-varying variables $X_j \in \mathcal{X}_j$ are recorded to collect individual health status. Consequently, a new treatment assignment A_j is delivered based upon the patient's longitudinal historical information from their first visit to the j^{th} visit, denoted by $H_j = \mathbf{S}_j(X_1, A_1, X_2, A_2, \dots, A_{j-1}, X_j) \in \mathcal{H}_j$, where \mathbf{S}_j is a deterministic summary function. In this paper, we consider the binary treatment settings, i.e., $A_j \in \mathcal{A} = \{1, -1\}$. In addition, we assume that patients' baseline information $H_1 = \mathbf{S}(X_1)$ is not confounded by treatment assignments (Schulte et al., 2014). After the final visit, a clinical outcome R , also known as the total reward, is obtained to reflect the benefits of the allocated treatment assignment. More specifically, the total reward R is a sum of the pseudo immediate rewards $R = \sum_{j=1}^T r_j$, where the immediate rewards $\{r_j\}_{j=1}^T$ are typically unobservable.

A dynamic treatment regime (DTR) is a sequence of decision rules $\mathcal{D} = \{D_j : \mathcal{H}_j \mapsto \mathcal{A}\}_{j=1}^T$ that map patients' historical information onto treatment space. The decision rules \mathcal{D} could also be represented by a composite function of the real-valued functions $\mathbf{f} = \{f_j \in \mathcal{F} : \mathcal{H}_j \mapsto \mathbb{R}\}_{j=1}^T$, where a realization of the decision that follows the treatment regime at the j^{th} visit is $d_j = D_j(H_j) = \text{sign}(f_j(H_j))$. Now, assume that the full trajectory of an observation sequence $\{X_1, a_1, X_2, a_2, \dots, X_T, a_T, R\}$ follows a data distribution P . Our goal is to seek the optimal treatment regime \mathcal{D}^* which yields the largest expected rewards among all regimes:

$$\mathcal{D}^* \in \underset{\mathcal{D}}{\text{argmax}} \mathbb{E}^{\mathcal{D}}\{\mathbf{R}\}. \quad (1)$$

Note that the expectation operator $\mathbb{E}^{\mathcal{D}}$ in equation (1) is taken with respect to an unknown restricted distribution $\{X_1, A_1 = d_1, X_2, A_2 = d_2, \dots, X_T, A_T = d_T, R\} \sim P^{\mathcal{D}}$, which describes the probability distribution when the treatments are assigned according to the regime \mathcal{D} . By convention, we call the corresponding \mathcal{D} the target regime. Since the historical information and potential outcome under $P^{\mathcal{D}}$ are unobservable, to infer the decision rules from the observed data while avoiding confounding issues between the assignments and expected rewards, we adopt the well-known assumptions of the *stable unit treatment value assumption* (STUVA) (Rubin, 1980) and *no unmeasured confounding* (Robins, 1986). Additionally, we suppose that the probabilities of assigning either treatment at every stage are positive, also called the *strong positivity* assumption.

2.2 Inverse Probability Weighted Estimators

Under the afore-mentioned assumptions, it can be shown that the expected total reward under the target regime \mathcal{D} is estimable by leveraging the inverse probability weighting (Qian and Murphy, 2011):

$$\mathcal{D}^* = \underset{\mathcal{D}}{\text{argmax}} \mathbb{E}^{\mathcal{D}}\{R\} = \underset{\mathcal{D}}{\text{argmax}} \mathbb{E} \left\{ \frac{R \cdot \prod_{j=1}^T \mathbb{I}(A_j = D_j(H_j))}{\prod_{j=1}^T \pi_j(a_j | H_j)} \right\}, \quad (2)$$

where π_j is the propensity score function and the density corrected expected reward is referred to as the Inverse Probability Weighting Estimator (IPWE). Provided that the propensity

functions are correctly specified, the optimal treatment regime \mathcal{D}^* is the maximizer of the IPWE. Accordingly, learning schemes that directly estimate the optimal regime by maximizing the IPWE estimator could be categorized as IPWE-based approaches.

In the existing literature, Outcome Weighted Learning (OWL) (Zhao et al., 2012) is one of the popular IPWE-based approaches to finding the optimal regime from a weighted classification perspective without the need to model the outcome. Although it was first designed to solve a single-stage decision task, the OWL framework was later extended to the multi-stage setting via the backwards-OWL (BOWL) and simultaneous-OWL (SOWL) procedures (Zhao et al., 2015). BOWL applies backward induction to break down multi-stages and searches for optimal decisions through the single-stage OWL framework. However, BOWL only utilizes patients whose observed treatments are exactly aligned with the decision rules at each stage during the backward process. As a result, the number of qualified patients could decrease exponentially when the number of stages increases. To avoid the situation where no patient is left as qualified before the backward process is completed, SOWL was subsequently proposed to transform the multi-stage estimation problem into a single-stage estimation problem. It simultaneously estimates the optimal regime at all stages by modeling the product of the indicator functions in equation (2) through a convex surrogate function $\psi(z_1, \dots, z_T) = \min(z_1 - 1, \dots, z_T - 1, 0)$ as follows:

$$V^{OW}(f_1, \dots, f_T) \doteq \mathbb{E} \left\{ \frac{R \cdot \min\{(A_1 f_1 - 1), \dots, (A_T f_T - 1), 0\}}{\prod_{j=1}^T \pi_j(a_j | H_j)} \right\}. \quad (3)$$

However, this surrogate function is non-smooth, and it is difficult to formulate the variable constraints under the quadratic programming when T is large. Besides, apart from the optimization difficulties, there is one more fundamental problem associated with the IPWEs that remains unresolved. Due to the density ratio corrections, a reward is only involved when an individual's treatments are fully matched with the target regime. Such a full-matching requirement is reflected in aligning treatments with the target regime at each backward stage from BOWL, and the surrogate function involves all decision stages from SOWL. Clearly, the strict full-matching requirement may lead to difficulties on both the optimization and estimation sides. We call this phenomenon the curse of full-matching.

3 Methodology

Our paper is motivated by the curse of full-matching under the IPWE framework. In this section, we propose a new DTR learning method, namely Stage-Aware Learning (SAL) which allows treatment mismatches with the target regime as a resolution. We will show that the optimized regime under SAL will achieve a more adaptable empirical performance compared to the IPWE-based regimes when the sample size is small and the number of decision stages is large. Furthermore, to additionally account for the heterogeneity of treatment effects at different decision stages, we introduce the Stage Weighted Learning (SWL) method upon SAL. In particular, we propose a finely-designed 0-1 metric called the stage importance score where the scores favor treatment alignments on stages with large treatment effects and hence facilitate the learning process.

3.1 The Curse of Full-matching

To fully understand the dilemma of full-matching, we consider a conditional expectation form of the expected total reward under the target regime \mathcal{D} , $\mathbb{E}^{\mathcal{D}}[R]$ in equation (2) as,

$$\mathbb{E}^{\mathcal{D}}\{R\} = \mathbb{E} \left\{ \underbrace{\frac{R}{\prod_{j=1}^T \pi_j(a_j|H_j)}}_{\text{Full-matching expected reward}} \left| \underbrace{\prod_{j=1}^T \mathbb{I}(A_j = D_j(H_j)) = 1}_{\text{Target regime assignment rate}} \right. \right\} \cdot P \left(\prod_{j=1}^T \mathbb{I}(A_j = D_j(H_j)) = 1 \right). \quad (4)$$

In equation (4), the expected reward to be maximized is determined by two major factors: first, the expected reward among the population whose observed treatments are fully matched with the target regime \mathcal{D} , i.e., *full-matching expected reward*; and second, the probability of assigning treatments conforming with the target regime at all stages, i.e., *target regime assignment rate*. At the population level, the probability of assigning any arbitrary dynamic regime is ensured to be non-zero due to the *strong positivity* assumption. However, if the optimal regime assignment rate is small among the patient population, it is highly possible that none of the sampled patients may fully follow the optimal treatments at all stages, and thus the optimal regime is infeasible in practice.

We demonstrate this phenomenon with the concrete illustrative example presented in Figure 1. The optimal treatment regime ($A_1 = 1, A_2 = 1$) yields the highest expected total reward of value 100, but with an extremely small assignment rate equaling 0.5%. In such a scenario, on average, only 1 person out of 200 sampled patients could fulfill the full-matching requirement. Hence, it might require many more samples to be collected before IPWE-based approaches could estimate the optimal regime and achieve near-asymptotic properties. Furthermore, note that the above example consists of only two decision stages. In practice, a growing number of stages tends to make estimation even harder since the target regime assignment rate will decrease exponentially. For instance, even if the assignment rate of the optimal decision at each stage is 80%, after 10 stages, the full-matching rate of the optimal regime at all stages can easily drop to 10% (0.8^{10}).

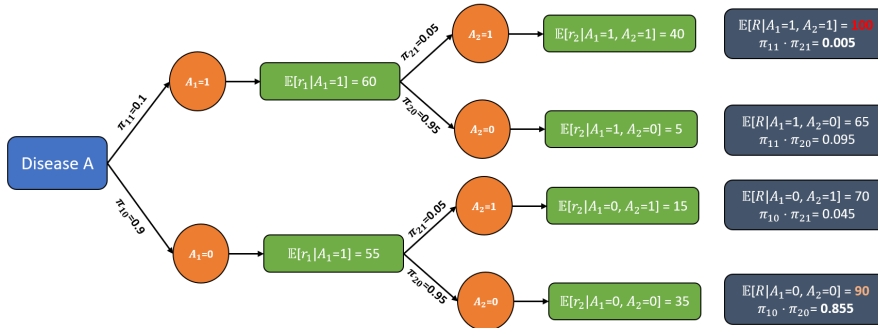


Figure 1: An example to illustrate the curse of full-matching. We consider a sequential randomized trial with a static treatment regime setting (Tsiatis et al., 2019). The assignment rate π_{ta} specifies the probability of allocating treatment a at stage t . The expected total rewards and assignment rates under corresponding regimes are presented in the rightmost column.

The above example implies that the IPWE-based approaches intrinsically require a full-matching condition to achieve competitive performance. Unfortunately, this condition is likely to be too restrictive to satisfy in practice, especially when the sample size is small and the number of stages is large. This motivates us to propose a novel framework which can break the curse of full-matching.

3.2 The k -partially Matching Estimator: k -IPWE

The empirical performance of IPWE-based approaches is largely restricted by the full-matching assignment rate of the optimal regime among the patient populations. Though we might not control how treatments are administrated according to the optimal regime at every decision stage, in the following, we propose to relax the strict full-matching requirement by allowing decision discrepancies between the assigned treatments and the target regime \mathcal{D} at $T - k$ number of stages, where $0 \leq k \leq T$. In other words, the optimal regime is allowed to partially match the treatment sequence at exactly k number of arbitrary decision stages. Hence, the probability of patients receiving optimal decisions at k number of stages could be much larger than the probability of patients receiving optimal decisions at all stages. We call this relaxed version of the full-matching requirement the *k -partially matching requirement*.

To formalize the notation, we let random variable K denote the number of correct alignments between treatments $\{A_j\}_{j=1}^T$ and decisions from an arbitrary target regime \mathcal{D} ,

$$K \doteq |\mathbf{A} \cap \mathcal{D}| = \sum_{j=1}^T \mathbb{I}(A_j = d_j). \quad (5)$$

If we specify K to be a realized value k within the range $\{0, \dots, T\}$, we are constraining our optimal regime search to the patient population who are k -partially matched with the optimal regime. A more extreme k value (e.g., $k = 0$ or $k = T$) corresponds to a more restricted alignment requirement. Under the IPWE framework, only patients with $K = T$ are included in the estimation procedure.

When $K = k$, a new restricted unknown distribution is induced, i.e., $(X_1, A_1 = \tilde{d}_1, \dots, X_T, A_T = \tilde{d}_T, R) \sim P^{\mathcal{D}^{(k)}}$, where $\tilde{d}_j = (-1)^{\mathbb{I}(j \in \mathcal{K})+1} \cdot d_j$ indicates that at any stage $0 \leq j \leq T$, the decision \tilde{d}_j is the same as the target regime d_j only if j is among the indexes of k arbitrary matching stages \mathcal{K} (i.e., $\tilde{d}_j = d_j$ if $j \in \mathcal{K}$). Subsequently, $P^{\mathcal{D}^{(k)}}$ is a distribution of an observation sequence where its k out of T assignments $\{A_j\}_{j \in \mathcal{K}}$ are followed by the regime \mathcal{D} . Based on the derivation of k -matching potential outcomes provided in Appendix A, we can similarly adopt the density ratio correction and obtain an IPWE for the expected rewards evaluated under the new measure $P^{\mathcal{D}^{(k)}}$. We denote the new estimator $\mathbb{E}^{\mathcal{D}^{(k)}}[R]$ as k -IPWE and present the results in proposition 1:

Proposition 1. *Under the SUTVA and no unmeasured confounding assumptions, the expected total reward under the target regime \mathcal{D} with k number of matching stages equals*

$$\mathbb{E}^{\mathcal{D}^{(k)}}[R] = \mathbb{E} \left\{ \frac{R \cdot \mathbb{I}(|\mathbf{A} \cap \mathcal{D}| = k)}{\prod_{j=1}^T \pi_j(a_j | H_j)} \right\}, \quad (6)$$

and the corresponding maximizing regime $\tilde{\mathcal{D}}^{(k)}$ is defined as

$$\tilde{\mathcal{D}}_{(k)} = \operatorname{argmax}_{\mathcal{D}} \mathbb{E}^{\mathcal{D}(k)} \{R\}. \quad (7)$$

The regime $\tilde{\mathcal{D}}_{(k)}$ maximizing the k -IPWE would yield the largest expected reward if patients were treated by \mathcal{D} at k number of stages. In addition, note that we do not require the k matching stages to be the same for each individual. As long as there is an exact k number of treatment matchings between the assignments and target regime \mathcal{D} , those patients' rewards are involved in the maximization process. As a result, the k -IPWE provides a superior level of flexibility.

However, the performance of the proposed k -IPWEs still depends on the pre-selected k value. The purpose of designing the K treatment matching number is to increase the conditional probability of the constrained population receiving optimal treatments. When the value k is poorly selected, the probability of patients being k -partially matched could be small, and we could encounter a similar aforementioned empirical dilemma. For instance, in a population with a 99% full-matching assignment rate, specifying the random variable K with values other than T leads to a small k -partially matching rate. In other words, the curse of full-matching can be effectively minimized only if the pre-specified k has the highest k -partially matching probability, i.e., $k = \operatorname{argmax}_{k \in \{0, \dots, T\}} P(K = k)$. Finding such a k value is doable but computationally cumbersome, and yet not every patient will participate in the optimization process due to the population conditional constraints. To reduce the uncertainty of selecting k values and potentially include all individuals from the sample, we further construct a learning method based on k -IPWE which can incorporate all scenarios of k -partially matching through applying a weighting scale on the matching number K .

3.3 Stage-Aware Learning Method (SAL)

In the following, we propose a novel learning method to combine all levels of k -IPWEs into one single estimation task. This would address the dependencies of k -IPWEs at the pre-selected k -values. Since the new estimator accounts for treatment and regime matching status at any number of stages, we name the new learning method as the Stage-Aware Learning (SAL).

Specifically, we weight each k -IPWE by k/T , proportional to the number of matching stages k , and estimate the optimal regime simultaneously by maximizing the following SAL value function derived in Appendix B,

$$V^{SA}(\mathcal{D}) = \sum_{k=0}^T \frac{k}{T} \cdot \mathbb{E}^{\mathcal{D}(k)} \{R\} = \mathbb{E} \left\{ \frac{R \cdot \frac{1}{T} \sum_{j=1}^T \mathbb{I}(A_j = D_j(H_j))}{\prod_{j=1}^T \pi_j(a_j|H_j)} \right\}, \quad (8)$$

We denote the estimated maximizing regime $\tilde{\mathcal{D}} = \operatorname{argmax}_{\mathcal{D}} V^{SA}(\mathcal{D})$. In our choice, the applied weights increasing with the k value indicate that, we prioritize the regime which has closer alignment with the optimal decisions, and meanwhile, achieves higher expected rewards during estimation. Structurally, the weighting component of the final SAL value function resembles the formulation of the Hamming loss (Tsoumakas and Katakis, 2007), which has the following three unique advantages.

First of all, compared to the IPWE value function (2), the SAL value function replaces the product of indicator functions with the correct treatment alignment percentage. Therefore,

instead of excluding patients completely if one of their observed treatments is not aligned with the optimal decision, SAL still considers those patients but discounts their rewards based on the degree of alignment between observed and optimal treatments across all of the decision stages. That is, even if the decision sequence is long, the new learning process is able to utilize all available patients' outcomes and maximize the alignment percentages over those with high rewards. As a result, all patients and their treatment strategies will be participated into the optimal regime searching procedure.

Secondly, the SAL learning scheme is now analogous to a multi-label classification framework. Instead of employing a surrogate function involving all decision stages simultaneously in outcome weighted learning (3), regime can be optimized at each individual stage to match with the optimal decisions. This leads to computational convenience as the optimization of the Hamming loss function has been well established and the multi-stage learning task can be segmented into several single-stage sub-tasks.

Furthermore, our proposed SAL suggests a more general DTR learning framework by allowing a probability distribution on the matching number K . Suppose the density function of K is proportional to a scale function $\phi(\cdot)$. After taking the iterated expectation of k -IPWEs under all possible choices of k values, we obtain a new estimator of the expected rewards aimed to maximized for any arbitrary target regime \mathcal{D} ,

$$\mathbb{E}_K \left\{ \mathbb{E}^{\mathcal{D}^{(k)}} \{R\} \right\} = \sum_{k=0}^T \phi(k) \cdot \mathbb{E}^{\mathcal{D}^{(k)}} \{R\} = \mathbb{E} \left\{ \frac{R \cdot \phi(|\mathbf{A} \cap \mathcal{D}|)}{\prod_{j=1}^T \pi_j(a_j | H_j)} \right\}. \quad (9)$$

The new estimator results in a generic form where the total rewards are weighted by the density of K . In return, the maximization procedure not only finds the regime that yields the largest expected total rewards, but also identifies the one which matches with the observed treatment sequence closest to the underlying distribution of the treatment matching number. Additionally, compared to the previous regime estimator $\tilde{\mathcal{D}}_{(k)}$ based on the k -IPWE in equation (7), the general estimator no longer constrains the rewards under a sub-population of patients. Instead, it takes every patient's reward into account as long as the density is positive for all matching number K , i.e., $\phi(k) > 0 \forall k \in \{1, \dots, T\}$.

Importantly, the induced general framework allows a flexible specification of the density scale function $\phi(\cdot)$, which can be used to summarize our prior knowledge of how well the observed treatments match with the optimal regime given the k matching number. For instance, the proposed SAL method adopts a linear scale function (i.e., $\phi(k) = k$) under our assumption that patients are more likely to receive a larger number of treatment assignments which coincide with the optimal regime. More choices of the $\phi(\cdot)$ function will be left in Remark 1. In particular, we demonstrate that the outcome weighted learning IPWE-based method is a special case of ours under our general framework.

Remark 1. *The outcome weighted learning IPWE-based framework can be recovered when a degenerated density function ($\phi(k) = \mathbb{I}(k = T)$) is specified. It assumes the probability of the assigned treatments fully aligning with the target regime at all decision stages is equal to 1. Consequently, only patients qualifying the full-matching requirement can be counted in the regime estimation process, and the resulting maximizing regime \mathcal{D} is equivalent to the optimal treatment regime \mathcal{D}^* , which enjoys all of the pre-established theoretical results in (Qian and Murphy, 2011; Zhao et al., 2015). Similarly, when K follows a degenerated distribution*

$\phi(k) = \mathbb{I}(k = j)$ at other stages j ($0 \leq j < T$), we can solve the general framework under the k -IPWE and obtain the maximizing regime $\tilde{\mathcal{D}} = \tilde{\mathcal{D}}_{(j)}$. It assumes there is most likely only one type of matching scenario between the assigned treatments and the optimal regime in the empirical data. Furthermore, several matching scenarios can be combined together. For instance, when the distribution density is specified as $\phi(k) = \frac{1}{2} \cdot \mathbb{I}(k = T) + \frac{1}{2} \cdot \mathbb{I}(k = T - 1)$, the maximizing regime $\tilde{\mathcal{D}}$ will optimize both IPWE and $(T - 1)$ -IPWE at the same time without preference due to the equal $1/2$ weights. As a result, it would yield the largest total rewards if the patient has followed all, or all but one, decisions suggested by \mathcal{D} .

In summary, we propose a novel SAL method, along with a more general DTR estimation framework, to improve data efficiency and empirical adaptivity of IPWE-based methods in this subsection. By combining each level of k -IPWEs, we attend to all matching scenarios and relax the selection of the K value. Through imposing higher weights to k -IPWEs with larger k value, SAL possesses the empirical-driven interpretation of searching samples with high rewards and large optimal treatment matching percentages. In the next subsection, we will further propose a weighted learning scheme based on SAL to incorporate stage heterogeneity and facilitate the learning process.

3.4 Stage Weighted Learning Method (SWL)

The SAL method transforms the DTR estimation into a multi-label classification problem. It simultaneously maximizes the total rewards and matching percentages between observed and optimal decisions. Individually, the total rewards component captures the combined treatment effects across all decision stages, while the matching percentage component promotes overall optimal treatment alignment. However, due to the uniform stage weights ($1/T$) outlined in equation (8), the SAL learning scheme faces an additional empirical challenge when attempting to pinpoint the treatment effects at each individual decision stage. That is, when dealing with samples that share the same matching percentage but possess different total rewards, it requires extra number of matching scenarios to segment the difference in the expected rewards to a single stage. To enhance the stage heterogeneity and facilitate learning process of treatment effects at each stage, we propose the Stage Weighted Learning (SWL) method corresponding to stage importance scores.

The stage importance scores are real-valued metrics designed to compare and quantify the relative importance of treatment effects among all decision stages. We denote the importance score at the j^{th} stage to be ω_j , where $\omega_j \in [0, 1]$ and $\sum_{j=1}^T \omega_j = 1$ for any j in $\{1, \dots, T\}$. Intuitively, a higher stage importance score indicates a more substantial treatment effect contributing to the total rewards. Formally, it affects patients' stage-wise rewards as

$$r_{ij} = \mathbf{r}_j(H_{ij}, A_{ij}) = \tilde{\mathbf{r}}_j(\omega_j, H_{ij}, A_{ij}), \quad (10)$$

where $\mathbf{r}_j : \mathcal{H} \times \mathcal{A} \mapsto \mathbb{R}^+$ is an original immediate reward function at j^{th} decision stage and we further assume there exists a counterpart reward function $\tilde{\mathbf{r}}_j : \mathbb{R}_{[0,1]} \times \mathcal{H} \times \mathcal{A} \mapsto \mathbb{R}^+$, which takes the stage importance score as an additional parameter and is able to disintegrate the decision stage-level heterogeneity effect from its original reward. To illustrate, consider a typical reward function consisted of two components: main effect α and treatment effect β , i.e., $r_{ij} = \alpha_{ij} + \beta_{ij}$. An exemplary reward counterpart indexed by the stage importance score

could be represented as, $r_{ij} = \alpha_{ij} + \omega_j \cdot \frac{\beta_{ij}}{\omega_j} = \alpha_{ij} + \omega_j \cdot \tilde{\beta}_{ij}$. Noticeably, with a direct association with the treatment effect, the stage importance scores measure the intrinsic heterogeneity existed at each decision stage and are invariant to individual patients. Correspondingly, we utilize the importance scores to include stage heterogeneity and hereby introduce the stage weighted learning method as

$$V^{SW}(\mathcal{D}) \doteq \mathbb{E} \left\{ \frac{R \cdot \sum_{j=1}^T \omega_j \cdot \mathbb{I}(A_j = D_j(H_j))}{\prod_{j=1}^T \pi_j(a_j | H_j)} \right\}. \quad (11)$$

Compared to SAL, the SWL scheme rescales the reward by a weighted average of the treatment alignments. Treatment mismatch on stages with large importance scores would lead to a substantial loss in total expected rewards; whereas the expected rewards only have minor fluctuations at stages with negligible importance scores, regardless of treatment assignments. In other words, the nonidentical stage importance scores impose additional stage heterogeneity on the learning scheme, and push the method to prioritize improving treatment alignment accuracy at important stages with large importance scores. In essence, maximizing the aforementioned value function is equivalent to solving a weighted multi-label classification, where the importance scores could be viewed as the label weights.

However, the estimation of the importance scores is non-trivial, mainly because the immediate rewards used to evaluate the treatment effects at each stage are unobservable. To solve this challenge, we estimate the immediate reward counterpart functions and the importance scores scalars from a semi-parametric point of view. Since the expected total reward is the summation of expected immediate rewards, we minimize the L^2 loss between total rewards and constructed surrogate rewards, i.e.,

$$\{\hat{\omega}_1, \dots, \hat{\omega}_T\} = \underset{(\omega_1, \dots, \omega_T) \in \mathbb{R}_{[0,1]}^{|T|}, (\tilde{\mathbf{r}}_1, \dots, \tilde{\mathbf{r}}_T) \in \mathcal{R}^{|T|}}{\operatorname{argmin}} \mathbb{E} \left\{ \left(R - \sum_{j=1}^T \tilde{\mathbf{r}}_j(\omega_j, H_j, A_j) \right)^2 \right\}. \quad (12)$$

Considering the importance scores represent the relative significance of treatment effects at each decision stage, we leverage the attention mechanism and propose the attention-based recurrent neural network architecture in Figure 2, which allows us to estimate the unobservable immediate rewards and approximate the importance scores.

The idea behind the attention mechanism (Bahdanau et al., 2014) is to scale an input sequence by relevance to the predicting outcomes, with a more relevant part being assigned a higher weight. These weights focus the attention of the prediction model on the most relevant part of the input sequence to improve model performance. In our scenario, we view the stage importance scores as the attention weights so that stages with higher contributions and more relevance to the final total rewards possess larger attention weights. Therefore, the weights not only explicitly impose stage heterogeneity on the proposed neural network, but also direct the network to pay more attention to the stages with large treatment effects when predicting the total rewards. The final importance scores are estimated empirically based on the loss function presented in (13) and the attention framework in Figure 2, i.e.,

$$\{\hat{\omega}_1, \dots, \hat{\omega}_T\} = \underset{(\omega_1, \dots, \omega_T) \in \mathbb{R}^{|T|}}{\operatorname{argmin}} \frac{1}{n} \sum_{i=1}^n \left(R_i - \sum_{j=1}^T \tilde{\mathbf{r}}_j(\omega_j \cdot (\mathbf{H}_j(X_{ij}, A_{i,j-1} | H_{i,j-1}), A_{ij})) \right)^2. \quad (13)$$

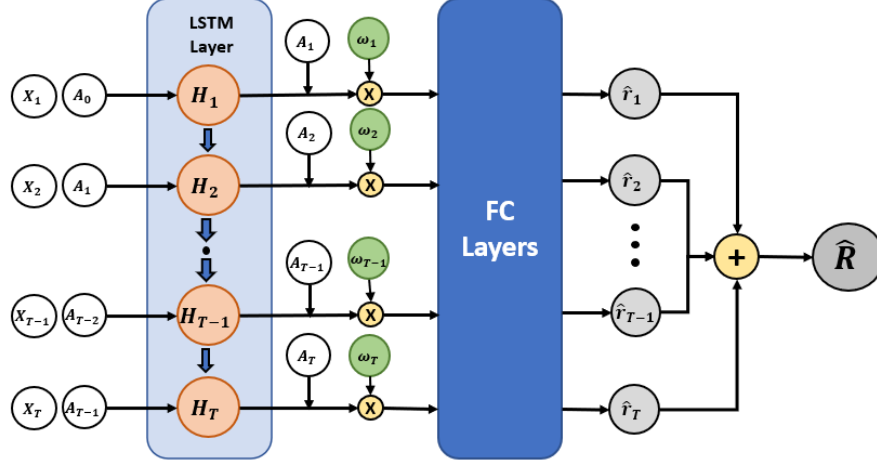


Figure 2: Architecture of stage importance scores searching network. The stage importance scores are treated as the attention weights applied on the patients’ historical information by the LSTM layer (Hochreiter and Schmidhuber, 1997), and are later estimated by minimizing the MSE between the observed and surrogate total rewards after the fully-connected (FC) layers transformation.

Note that we do not limit the parametric form of the reward counterpart function. For illustration purposes, we represent the reward function with a fully-connected (FC) network due to its flexible capability of function approximation (DeVore et al., 2021). In addition, we adopt a long-short term memory (LSTM) network (Hochreiter and Schmidhuber, 1997) to capture the unobserved patients’ historical information H_t at stage t using up-to-date patients’ covariate information $X_{1:t}$ and past treatments $A_{1:(t-1)}$. With estimated importance scores, we can search for the optimal treatment regime under the SWL value function (11), via maximizing the objective function with a smooth convex surrogate function ψ (Bartlett et al., 2006):

$$\hat{\mathcal{D}}_{\psi}^{SW} = (\hat{f}_{\psi 1}, \dots, \hat{f}_{\psi T})^{SW} = \underset{\{f_1, \dots, f_T\} \in \mathcal{F}^{|T|}}{\operatorname{argmax}} \frac{1}{n} \sum_{i=1}^n \frac{R_i \sum_{j=1}^T \hat{\omega}_j \psi(a_{ij} \cdot f_j(H_{ij}))}{\prod_{j=1}^T \pi_j(a_{ij} | H_{ij})}. \quad (14)$$

In particular, we can employ the logistic function as a surrogate to the indicators, i.e., $\psi(x; \lambda) = (e^{-\lambda x} + 1)^{-1}$. The hyper-parameter λ controls the growth rate where a larger λ value makes the surrogate converge to the 0-1 indicator function faster.

To conclude, the proposed SWL method inherits the property of treatment mismatching from SAL and further adopts stage heterogeneity via the empirical stage importance scores estimated from the attention-based recurrent neural network. The resulting SWL scheme is able to estimate the sequential DTR under the weighted multi-label classification framework. In the next section, we will present the theoretical results of the proposed SWL method.

4 Theoretical Results

In this section, we present the theoretical results of the proposed SWL method. In Theorem 1, we show the Fisher consistency of the SWL surrogate estimator $\hat{\mathcal{D}}_{\psi}^{SW}$ compared to the

SWL optimal regime $\tilde{\mathcal{D}}^{SW}$. Next, we demonstrate the SWL Fisher consistency of the optimal regime \mathcal{D}^* in Theorem 2, and establish the finite-sample performance error bound with a flexible metric entropy in Theorem 3. To the best of our knowledge, Theorem 2 is the first theoretical result to fully discuss the gap in Fisher consistency between multi-stage DTR methods and the multi-label classification framework. The proof of the theorems and additional Lemmas are provided in Supplemental Materials.

We first introduce our notation for the technical developments. We let $\mathcal{G}^{|T|}$ be a generic product function space; and denote $\mathbf{g}^* = (g_1^*, \dots, g_T^*) \in \mathcal{G}^{|T|}$ and $\mathbf{g}_\psi^* = (g_{\psi_1}^*, \dots, g_{\psi_T}^*) \in \mathcal{G}^{|T|}$ as the optimal treatment regimes with respect to the SWL value function V^{SW} (11) and its surrogate counterpart V_ψ^{SW} , respectively. Moreover, we define a parametric product function space $\mathcal{F}^{|T|}$, where we search the maximizer $\mathbf{f}^* = (f_1^*, \dots, f_T^*) \in \mathcal{F}^{|T|}$ and $\mathbf{f}_\psi^* = (f_{\psi_1}^*, \dots, f_{\psi_T}^*) \in \mathcal{F}^{|T|}$ for the function approximations for \mathbf{g}^* and \mathbf{g}_ψ^* , respectively. Given the observed data, we further define $\hat{\mathbf{f}}_n = (\hat{f}_{\psi_1}, \dots, \hat{f}_{\psi_T}) \in \mathcal{F}^{|T|}$ as the empirical maximizer of the SWL objective function \hat{V}_ψ^{SW} (14). Before establishing the theoretical results, we present the following necessary regularity conditions.

Assumption 1. (*Finite Reward*) *The total reward is positive and upper-bounded by a finite constant M , i.e., $0 \leq \|R\|_\infty \leq M < \infty$.*

Assumption 2. (*Positivity*) *The propensity score $\pi_j(a_j|H_j)$ is lower-bounded by a positive real number, c_0 , s.t. $0 < c_0 \leq \pi_j < 1$.*

Assumption 3. (*No Approximation Error*) *Suppose for any parameterized functional space $\mathcal{F}^{|T|}$, the approximation error ϵ_{app} satisfies the following:*

$$\epsilon_{app} := \sup_{\mathbf{g} \in \mathcal{G}^{|T|}} \inf_{\mathbf{f} \in \mathcal{F}^{|T|}} \|\mathbf{g} - \mathbf{f}\|_\infty = 0.$$

Assumption 1 is a standard assumption, which requires the total reward to be positive and bounded. Assumption 2 indicates that the probability of assigning any treatment to arbitrary stages is positive in the observational studies. Assumption 3 defines an approximation error ϵ_{app} due to the difference between the parameterized space $\mathcal{F}^{|T|}$ and the generic space $\mathcal{G}^{|T|}$. By setting ϵ_{app} to zero, the assumption states that for any function sequence \mathbf{g} in the generic function space $\mathcal{G}^{|T|}$, we can find a function sequence \mathbf{f} in the parameterized space $\mathcal{F}^{|T|}$ such that $\mathbf{f} = \mathbf{g}$. Equivalently, it can be shown that the optimal regime \mathbf{g}^* belongs to $\mathcal{F}^{|T|}$ and $\mathbf{g}^* = \mathbf{f}^*$.

4.1 Surrogate Fisher consistency

The adoption of surrogate functions eases the optimization procedure. In this subsection, we establish the Fisher consistency between the value function V^{SW} and its surrogate form V_ψ^{SW} . Specifically, we show that the optimal surrogate treatment decision $\text{sign}(f_{\psi_j}^*)$ is aligned with the optimal decision $\text{sign}(f_j^*)$ at each stage. The obtained result is presented in Theorem 1.

Theorem 1. *Let $\psi(a, f; \lambda) : \mathcal{A} \times \mathcal{F} \times \Lambda \mapsto \mathbb{R}$ be a surrogate function with tuning parameters λ that satisfies $\psi(a, f; \lambda) = \psi(-a, -f; \lambda)$ and $\text{sign}(\psi(1, f; \lambda) - \psi(-1, f; \lambda)) = \text{sign}(f)$. Then,*

for all $t = 1, \dots, T$ and $H_t \in \mathcal{H}_t$,

$$\text{sign}(f_{\psi_t}^*(H_t)) = \text{sign}(f_t^*(H_t)) = \underset{a_t \in \{-1, 1\}}{\text{argmax}} \mathbb{E} \left\{ r_t + \sum_{j=t+1}^T r_j \mid A_t = a_t, H_t \right\}. \quad (15)$$

Theorem 1 guarantees the same treatment decisions could be obtained from the surrogate estimators and the maximizing regime under the SWL scheme. This validates the usage of smooth surrogate functions to approximate the indicator functions at each decision stage which cannot be solved under the IPWE framework. In addition, according to Lemma 4, the surrogate is only required to be an even function and produce the same sign as the treatment effect. In fact, a wide class of surrogate functions, such as indicators, logistic functions, and binary cross-entropy, satisfies such requirements and increases the optimization flexibility of the proposed SWL method.

In the following subsection, we also demonstrate the Fisher consistency between SWL and the optimal DTR $\mathcal{D}^* = \{\mathcal{D}_j^*\}_{j=1}^T$. The optimal decision d_t^* on arbitrary stage t from equation (1) has the form of

$$d_t^* = \mathcal{D}_t^*(H_t) = \underset{a_t \in \{-1, 1\}}{\text{argmax}} \mathbb{E} \left\{ r_t + \sum_{j=t+1}^T r_j \mid A_t = a_t, H_t, \{A_j\}_{j=t+1}^T = \{d_j^*\}_{j=t+1}^T \right\}. \quad (16)$$

Compared to SWL maximizing regime \mathbf{f}^* in equation (15), the optimal DTR \mathcal{D}^* also aims to maximize expected future reward, but further assumes every future treatment step matches with the optimal decision. Hence, depending on the underlying behavioral distributions of rewards and treatment assignments, SWL could recover \mathcal{D}^* asymptotically only under much more stringent conditions.

4.2 Fisher consistency: optimal treatment dominance

To fill in the gap between SWL and optimal DTR, we investigate the boundary condition where the SWL estimators might produce different treatment decisions from the optimal regime. This condition could be quantified via the relationships between the treatment effects and expected future rewards, and is defined as

Condition 1. (*Optimal Treatment Dominance*) Any decision Stage t , where $1 \leq t \leq T$, is said to be dominated by the optimal treatment if

$$\begin{aligned} \Psi_t(H_t; \{d_j^*\}_{j=t+1}^T) \cdot P(\{A_j\}_{j=t+1}^T = \{d_j^*\}_{j=t+1}^T) \geq \\ \mathbb{E} \left\{ \Psi_t(H_t; \{A_j\}_{j=t+1}^T) \mid \{A_j\}_{j=t+1}^T \neq \{d_j^*\}_{j=t+1}^T \right\} \cdot P(\{A_j\}_{j=t+1}^T \neq \{d_j^*\}_{j=t+1}^T), \end{aligned} \quad (17)$$

where $U_t(H_t; \{A_j\}_{j=t+1}^T) = \mathbb{E}\{\sum_{j=t+1}^T r_j \mid H_t, \{A_j\}_{j=t+1}^T\}$ is the expected future reward resulting from preassigned future actions after the t^{th} decision stage, and $\Psi_t(H_t; \{A_j\}_{j=t+1}^T) = |\mathbb{E}\{r_t + U_t(H_t; \{A_j\}_{j=t+1}^T) \mid A_t = 1, H_t\} - \mathbb{E}\{r_t + U_t(H_t; \{A_j\}_{j=t+1}^T) \mid A_t = -1, H_t\}|$ represents the absolute difference in the expected total rewards due to treatment at Stage t , i.e., the

treatment effect. Overall, Condition 1 indicates optimal treatment dominance if the treatment effect of the optimal decisions, weighted by its matching probability, is larger than any other combination of non-optimal decisions. Then, if Condition 1 is satisfied at every decision stage, Fisher consistency between SWL and optimal DTR can be reached,

Theorem 2. *For all stages $t = 1, \dots, T$ and $H_t \in \mathcal{H}_t$, $\text{sign}(f_{\psi_t}^*) = \text{sign}(f_t^*) = d_t^*$ if and only if Stage t is dominated by the optimal treatment effect.*

As Theorem 2 shows, Fisher consistency depends on the dominance of the optimal treatments. For a better understanding of Condition 1, we begin with an extreme scenario where every individual receives the optimal decisions. Obviously, optimal decisions in this case are dominant as there is no other regime assigned, and Condition 1 is satisfied at every stage, i.e., $P\left(\{A_j\}_{j=t+1}^T \neq \{d_j^*\}_{j=t+1}^T\right) = 0$. Furthermore, since the expected future reward $\mathbb{E}\left\{\sum_{j=t+1}^T r_j \mid H_t\right\}$ is equal to the expected future optimal reward $\mathbb{E}\left\{\sum_{j=t+1}^T r_j \mid H_t, \{A_j\}_{j=t+1}^T = \{d_j^*\}_{j=t+1}^T\right\}$, it is straightforward to show that the SWL maximizing regime \mathbf{f}^* is the same as the optimal DTR \mathbf{d}^* .

In a more general sense, Condition 1 considers the treatment effects and the optimal regime assignment rate at the same time when estimating the treatment regime. Though the optimal regime could be hardly assigned to the patients, i.e., $P\left(\{A_j\}_{j=t+1}^T = \{d_j^*\}_{j=t+1}^T\right)$ is small, the optimal decision d_t^* is still preferred if its treatment effect is much larger and therefore dominates other regimes, i.e., $\Psi_t(H_t; \{d_j^*\}_{j=t+1}^T) > \mathbb{E}\left\{\Psi_t(H_t; \{A_j\}_{j=t+1}^T) \mid \{A_j\}_{j=t+1}^T \neq \{d_j^*\}_{j=t+1}^T\right\}$. On the other hand, if the treatment effects of all possible regimes are similar, the SWL estimators choose the decision which is more often applied to the population. Accordingly, due to Condition 1, our proposed SWL is able to recover the optimal regime and adaptively make treatment changes based on the treatment effect and the regime-assigned rate from the collected empirical data.

4.3 Finite-sample performance error bound

Theorems 1 and 2 establish the consistency properties of the proposed SWL method. To investigate finite sample performance of the proposed approach, we also establish the performance error bound and investigate the convergence rate. In the following, we require measuring the function space complexity for the parameterized functional space \mathcal{F} .

Assumption 4. *(Capacity of Function Space) Let $\mathcal{F} = \{f \in \mathcal{F} : \|f\| \leq 1\}$ and $H_1, \dots, H_n \in \mathcal{H}$. There exist constants $C > 0$ and $0 < \alpha < 1$ such that for any $u > 0$, the following condition on metric entropy is satisfied:*

$$\log \mathcal{N}_2(u, \mathcal{F}, H_{1:n}) \leq C \left(\frac{1}{u}\right)^{2\alpha}. \quad (18)$$

Assumption 4 characterizes the functional space complexity with the logarithmic minimum number of balls with radius u required to cover a unit ball in \mathcal{F} , and is satisfied under various functional spaces such as the reproducing kernel Hilbert space (RKHS) and Sobolev

space (Van de Geer, 2000; Steinwart and Christmann, 2008). Consequently, the performance bound between $V^{SW}(\mathbf{f}^*)$ and $\hat{V}_\psi^{SW}(\hat{\mathbf{f}}_n)$ is provided in the following theorem.

Theorem 3. *Under Assumptions 1-4, there exist constants $C_1 > 0$ and $0 < \alpha < 1$ such that for any $\delta \in (0, 1)$, w.p. at least $1 - \delta$, the performance error is upper-bounded by:*

$$\left| V^{SW}(\mathbf{f}^*) - \hat{V}_\psi^{SW}(\hat{\mathbf{f}}_n) \right| \leq \underbrace{\frac{M}{c_0^T} \sum_{j=1}^T \omega_j \epsilon_{n,j}}_{\text{Surrogate error}} + \underbrace{\frac{6(\alpha + 1)}{\alpha} \left[\alpha C_1 \sqrt{\frac{T}{n}} \left(\frac{\lambda M}{4c_0^T} \right)^\alpha \right]^{\frac{1}{\alpha+1}} + \frac{9M}{c_0^T} \sqrt{\frac{\log 2/\delta}{2n}}}_{\text{Empirical estimation error}}, \quad (19)$$

where $\epsilon_{n,j} = \sup_{A_j, H_j} |\mathbb{I}(A_j f_j(H_j) > 0) - \psi(A_j f_j(H_j); \lambda_n)|$.

In Theorem 3, the finite-sample performance bound can be broken down into two separate bounds: the surrogate error bound between V^{SW} and V_ψ^{SW} and the empirical estimation error bound of V_ψ^{SW} . As a result, the SWL performance error convergence rate could be obtained at $\mathcal{O}(\epsilon_{n,j} + n^{-1/(2\alpha+2)})$. In particular, the first term depends on the choice of the surrogate function. When the surrogate is well-selected as a logistic function, e.g., $\psi(x; \lambda_n) = (e^{-\lambda_n x} + 1)^{-1}$ with a rate hyper-parameter λ_n , the surrogate error $\epsilon_{n,j}$ vanishes to zero at the rate of $\mathcal{O}(e^{-n})$, which is much faster than the second term; and therefore the performance error bound could be reduced to $\mathcal{O}(n^{-1/(2\alpha+2)})$.

Furthermore, based on the $\mathcal{O}(n^{-1/(2\alpha+2)})$ convergence rate, Theorem 3 provides the finite-sample upper bound which validates the estimation risk and demonstrates how SWL converges under different parametric space settings. For instance, when the historical information \mathcal{H} is an open Euclidean ball in \mathbf{R}^d and the functional space is specified as the Sobolev space $\mathbb{W}^k(\mathcal{H})$ where $k > d/2$, one can choose $\alpha = d/2k$ to obtain an error upper bound of rate at $\mathcal{O}(n^{-d/2(d+2k)})$. In addition, when the functional space is finite, the upper error bound could reach the best rate at $\mathcal{O}(n^{-1/2})$ as $\alpha \rightarrow 0$, which achieves the optimal rate provided in the literature (Zhao et al., 2015, 2019).

To summarize, our finite-sample performance error bound recovers the best-performing convergence rate found in the existing literature and meanwhile provides a non-asymptotic explanation of the proposed multi-stage DTR method in empirical settings. With different choices of metric entropy, the performance error bound can be flexibly adapted to various functional spaces and is not limited to the RKHS discussed in Zhao et al. (2015). In addition, our metric condition from Assumption 4 only needs to be satisfied under a more relaxed empirical L_2 -norm on the collected samples $H_{1:n}$, compared to the supremum norm assumption in Zhao et al. (2019).

5 Implementation and Algorithm

In this section, we provide our main algorithm for stage importance scores searching and the optimal SWL regime estimation. The goal is to find a set of optimal parameters that minimizes the MSE of rewards (13) and maximizes the objective function of SWL (14).

The algorithm starts with finding the stage importance weights by constructing the neural network as specified in Figure 2, which could be summarised into two major steps. First, we model the deterministic summary function \mathbf{S} via LSTMs and estimate patients' historical

information \mathbf{H}_j at each stage. Second, we use the estimators of stage importance scores to scale \mathbf{H}_j and apply fully-connected (FC) layers on the weighted historical information to estimate the total rewards. Once the surrogate total rewards are estimated, the MSE loss between the observed and surrogate total rewards can be computed, and the parameters in the neural networks are updated from the back-propagation process with stochastic gradient descent (SGD)-based optimizers (Robbins and Monro, 1951).

For estimating the optimal regime, there are still two missing pieces need to be filled in according to the SWL objective function (14): the function representations of the target regime and the propensity scores of each observed treatment. In this proposed algorithm, we model the treatment rules $\{f_j\}_{j=1}^T$ with a FC-network. Since the treatment rule could be linear or non-linear, we adjust the activation functions applied on each layer within the network accordingly. To estimate the propensity scores, we apply the logistic regression for each individual at each stage j , i.e., $\{\hat{\pi}_{ij}\}_{i=1}^n$. Finally, we combine every component and managed to present the entire workflow in Algorithm 1.

Algorithm 1 Stage Weighted Learning

1. **Initialize** stage weights $\{\omega_j\}_{j=1}^T$; the LSTMs parameterized by θ_L ; the stage-weight FC-network parameterized by θ_s ; the treatment FC-network parameterized by θ_f ; learning rate λ ; maximum iterations T_{max} ; and a stopping error criterion ϵ_s
 2. **Input** all observed sequence $\{(X_{i1}, A_{i1}, X_{i2}, A_{i2}, \dots, X_{iT}, A_{iT}, R_i)\}_{i=1}^n$
 3. **For** $k \leftarrow 1$ to T_{max} **do**
 4. Compute gradient w.r.t. θ_l, θ_s and $\{\omega_j\}_{j=1}^T$ as
 5.
$$\mathcal{L}_1^k = \frac{1}{n} \sum_{i=1}^n \left[R_i - \sum_{j=1}^T FC_{\theta_s}(\omega_j \cdot LSTM_{\theta_l}(X_{ij}, A_{i,j-1})) \right]^2$$
 6. Update parameters of interests $(\{\omega_j\}_{j=1}^T, \theta_l, \theta_s)^k \leftarrow (\{\omega_j\}_{j=1}^T, \theta_l, \theta_s)^{k-1} - \lambda \cdot \nabla \mathcal{L}_1^k$
 7. Stop if $|\mathcal{L}_1^k - \mathcal{L}_1^{k-1}| \leq \epsilon$
 8. **Normalize** $\hat{\omega}_j = \exp(|\hat{\omega}_j|) / \sum_{j=1}^T \exp(|\hat{\omega}_j|)$ and finalize $\hat{H}_{ij} = LSTM_{\theta_l^k}(X_{ij}, A_{i,j-1})$
 9. **Estimate** $\{(\hat{\pi}_{ij})_{j=1}^T\}_{i=1}^n$ via logistic regressions on $A_{ij} \sim 1 + \hat{H}_{ij}$
 10. **For** $k \leftarrow 1$ to T_{max} **do**
 11. Compute gradient w.r.t. θ_f as
 12.
$$\mathcal{L}_2^k = -\frac{1}{n} \sum_{i=1}^n \left[\frac{R_i}{\prod_{j=1}^T \hat{\pi}_{ij}} \cdot \sum_{j=1}^T \hat{\omega}_j \cdot \left(\exp(-A_{ij} \cdot FC_{\theta_f^k}(\hat{H}_{ij})) + 1 \right)^{-1} \right]$$
 13. Update parameters of interests $\theta_f^k \leftarrow \theta_f^{k-1} - \lambda \cdot \nabla \mathcal{L}_2^k$
 14. Stop if $|\mathcal{L}_2^k - \mathcal{L}_2^{k-1}| \leq \epsilon$
 15. **Return** estimated treatment regime network $\hat{\theta}_f = \theta_f^k$
-

Algorithm 1 demonstrates the end-to-end procedure of the SWL optimal regime estimation, which adopts the standard SGD optimization method, but other techniques could be also considered to improve the convergence and estimation results. For instance, Adam optimizer (Kingma and Ba, 2014) could be a suitable alternative to improve the convergence

performance of SGD on highly-complex and non-convex objective functions. In addition, instead of setting learning rates to be constant presented in the algorithm, utilizing cosine annealing warm-restart schedule (Loshchilov and Hutter, 2016) and different initialization seeds (Diamond et al., 2016) could improve the optimization to achieve a better local convergence. Furthermore, we can also tune the hyper-parameters, such as the learning rate and the number of network hidden layers, by conducting a d -fold cross-validation on the dataset. The detailed cross-validation procedure is described as follows. The dataset is first randomly partitioned into d evenly-sized subsets, and then the neural network is trained on each of the $(d - 1)$ subsets and tested on one remaining subset. After averaging d testing loss, the set of hyper-parameters with minimal testing loss is selected as optimal ones based on the empirical dataset. Once the two-step algorithm is converged and the maximal value of the objective function is reached, we obtain the optimal empirical SWL regime.

6 Simulation

In this section, we present simulation studies to showcase the empirical advantages of our proposed methods over popular multi-stage DTR frameworks; e.g., QLearning (Zhao et al., 2009), BOWL (Zhao et al., 2015), residual weighted learning (RWL) (Zhou et al., 2017) and the augmented IPWE-based approach (AIPW) (Zhang et al., 2013), where the later two methods are robust estimators. Specifically, we investigate the effects of the sample size, number of stages, optimal regime assignment rate and regime function complexity on the model performance. Furthermore, we show the advantages of incorporating stage importance scores when stage heterogeneity exists in the decision stages.

The general simulation setting is described as follows. First of all, a total number of 20 features $\{X_{i1k}\}_{k=1}^{20}$ are independently generated from a standard normal distribution $N(0, 1)$ at baseline ($t = 1$), and progress according to the treatment assigned at the previous decision stage, i.e.,

$$X_{i,t+1,k} = \begin{cases} 0.8 \cdot X_{itk} + 0.6 \cdot \epsilon & \text{if } A_{it} = 1 \\ 0.6 \cdot X_{itk} + 0.8 \cdot \epsilon & \text{if } A_{it} = 0 \end{cases} \quad (20)$$

where Gaussian random noise $\epsilon \sim N(0, 1)$. The variance of the generated covariates is kept constant and the treatment heterogeneity is incorporated.

We design the optimal treatment regime function at each decision stage under linear and non-linear settings. Under the linear setting, the optimal regime function is a linear combination of the covariates without interactions; whereas under the nonlinear setting, we select functions g_t from a basis of functions $\{X, X^2, X^3, \arctan X, \text{sign}(X)\}$, and interaction terms are included among the transformed covariates to increase the function complexity. Accordingly, we formalize the optimal treatment regime generation procedure, i.e.,

$$f_t^*(\mathbf{X}_t) = \begin{cases} \sum_{j \in \mathcal{J}} \beta_{tj} \cdot X_{.,t,j} & \text{Linear setting} \\ \sum_{j \in \mathcal{J}} \beta_{tj} \cdot \prod_{s \in \mathcal{S}_j} g_{ts}(X_{.,t,s}) & \text{Non-linear setting.} \end{cases} \quad (21)$$

Here, $\beta_{tj} \sim N(0, 1)$, \mathcal{J} is the randomly selected covariate index with cardinality $|\mathcal{J}| \sim \text{Unif}(1, 20)$, \mathcal{S}_j is a random index set with cardinality $|\mathcal{S}_j| \sim \text{Unif}(1, 3)$ specifying the interaction terms for the j^{th} transformed covariate, and each g_{ts} is a nonlinear function randomly sampled from the pre-specified functional basis. Consequently, the optimal decision

functions and the number of covariates contributing to the optimal treatment rule vary at each decision time, and can be expanded to long decision sequences.

We further define a linear immediate reward function after obtaining the optimal decision d_t^* from the optimal regime at each step as

$$r_t(\mathbf{X}_t, \mathbf{A}_t) = \omega_t \left[\left(\sum_{j \in \mathcal{J}_r} \beta_{tj}^r \cdot X_{\cdot,t,j} \right) + A_t \cdot d_t^* \right] + \epsilon_r, \quad (22)$$

where $|\mathcal{J}_r| \sim \text{Unif}(5, 20)$, $\beta_{tj}^r \sim N(0, 1)$ and $\epsilon_r \sim N(0, 1)$. Notice that the reward function consists of three main components: the base reward from patients' covariates, the treatment effect, and the stage importance scores. To specify the values of the importance scores $\{\omega_t\}_{t=1}^T$, we randomly sample from a Dirichlet distribution with important stages having a weight parameter of 100 and non-important stages having weight parameter 1, i.e., $\{\omega_t\}_{t=1}^T \sim \text{Dirichlet}(\{1, 100\})$. Correspondingly, important stages have larger importance scores and hence more substantial treatment effects. Based on the immediate rewards, the total rewards can be computed via $R_i = R(\mathbf{X}_i, \mathbf{A}_i) = \mathbb{E} \left\{ \sum_{t=1}^T r_t(X_{it}, A_{it}) \right\}$ and the value function $V(\mathbf{X}, \mathbf{A}) = \frac{1}{n} \sum_{i=1}^n R_i$ is used to evaluate the performance of the assigned regime.

In general, for each specification of listed parameters under the general setting, we repeat experiments 50 times for data generation. All methods are trained using 80% of the simulated training data, and evaluated on the 20% testing set via value functions and the matching accuracy between the estimated and optimal treatment regimes. Due to page restrictions, we provide partial experiment results in Supplemental Materials.

6.1 Effects of sample size and number of decision stages

Sample sizes and the number of decision stages are the two important factors that affect the curse of full matching as introduced in section 3.1. To comprehensively examine the effects of these two factors on our proposed algorithm, we first conduct simulations on sample sizes $n = 500, 1000, 5000$, and the number of decision stages $T = 5, 8, 10$ where treatments A_t are randomly matched with the nonlinear optimal decisions at 50% chance and the number of important stages is set to 0. The results of model performance are summarized in Table 1.

According to Table 1, we notice that SWL outperforms all other competing methods with respect to the estimated total rewards. Given a fixed number of stages, every single model deteriorates as expected when the sample size decreases as indicated by the smaller values of estimated total rewards, but the improvement margin of SWL compared to the best-performing competing method increases. In particular, when $T = 5$, the SWL improves the estimated total rewards nearly seven times as much, from 8.51% to 56.63%, as the sample size decreases from 5000 to 200. In addition, the difference between our model performance and other competing methods enlarges with an increasing number of decision stages. For instance, when $n = 5000$, the improvement rates increase from 8.51% to 64.42% as T grows from 5 to 10. This implies that the proposed method has a more efficient utilization of the observed information and the SWL has more advantages when the sample size is small and the number of treatment stages is large.

T	n	QLearning	AIPW	RWL	BOWL	SWL	Oracle	Imp-rate (To Best)
5	5000	0.107 (0.042)	0.247 (0.024)	0.106 (0.053)	0.343 (0.018)	0.372 (0.026)	0.601 (0.004)	8.51%
	1000	0.102 (0.044)	0.147 (0.033)	0.099 (0.063)	0.239 (0.032)	0.268 (0.045)	0.601 (0.014)	11.87%
	500	0.093 (0.047)	0.105 (0.039)	0.114 (0.063)	0.185 (0.051)	0.247 (0.047)	0.598 (0.020)	33.44%
	200	0.064 (0.069)	0.074 (0.052)	0.140 (0.066)	0.116 (0.070)	0.220 (0.069)	0.600 (0.030)	56.63%
8	5000	0.043 (0.027)	0.091 (0.015)	0.055 (0.025)	0.143 (0.014)	0.191 (0.016)	0.375 (0.005)	33.10%
	1000	0.030 (0.028)	0.052 (0.016)	0.054 (0.024)	0.080 (0.021)	0.145 (0.022)	0.375 (0.009)	80.85%
	500	0.033 (0.029)	0.035 (0.025)	0.066 (0.027)	0.058 (0.03)	0.140 (0.028)	0.378 (0.015)	110.88%
	200	0.014 (0.038)	0.008 (0.039)	0.053 (0.035)	0.022 (0.037)	0.125 (0.04)	0.370 (0.024)	134.96%
10	5000	0.021 (0.024)	0.055 (0.01)	0.043 (0.017)	0.086 (0.012)	0.141 (0.014)	0.300 (0.006)	64.42%
	1000	0.016 (0.021)	0.031 (0.011)	0.041 (0.017)	0.044 (0.015)	0.112 (0.021)	0.300 (0.007)	151.80%
	500	0.015 (0.025)	0.018 (0.018)	0.039 (0.021)	0.030 (0.019)	0.109 (0.021)	0.299 (0.012)	179.49%
	200	0.003 (0.027)	0.006 (0.024)	0.039 (0.027)	0.014 (0.031)	0.106 (0.021)	0.298 (0.019)	175.13%

Table 1: Estimated total rewards when the optimal regime is nonlinear, assigned treatment $A_t \sim \text{Bernouli}(0.5) \cdot d_t^*$ and no important stage. Standard errors are listed next to the estimated means. The Oracle stands for the best estimated total rewards if all treatments are assigned optimally. The improvement rate compares SWL against the best performer of competing methods.

6.2 Full-matching rates between assigned and optimal treatments

We illustrate the curse of full matching under various sample sizes and numbers of stages. However, the empirical dilemma could compromise the convergence of the DTR methods at the same time. To describe a more straightforward association between full-matching rates and model performance while minimizing the effects from non-convergent results, we set $n = 5000$, $T = 10$, and directly adjust the matching probabilities between assigned and optimal treatment. The full-matching rates are $0.5^{10} \approx 0.001$, $0.7^{10} \approx 0.03$, and $0.8^{10} \approx 0.1$.

Matching Probability ($P(\sum_{t=1}^{10} \mathbb{I}(A_t = d_t^*) = 10)$)	QLearning	AIPW	RWL	BOWL	SWL	Oracle	Imp-rate (To BOWL)
Scenario 1 (0.100)	0.044 (0.022)	0.102 (0.012)	0.098 (0.014)	0.196 (0.008)	0.207 (0.007)	0.300 (0.006)	5.61%
Scenario 2 (0.030)	0.042 (0.023)	0.094 (0.013)	0.090 (0.015)	0.170 (0.010)	0.203 (0.009)	0.300 (0.006)	19.41%
Scenario 3 (0.001)	0.041 (0.022)	0.090 (0.010)	0.086 (0.016)	0.144 (0.009)	0.201 (0.012)	0.300 (0.006)	39.58%

Table 2: Estimated total rewards of listed models when $n = 5000$, $T = 10$, and the treatments are matched with the nonlinear optimal decisions based on the pre-specified full-matching rates. Standard errors are provided in parentheses. Improvement rates compare SWL against BOWL.

Based on the results presented in Table 2, we observe that the proposed SWL method outperforms the rest of the competing methods, and the improvement rates against BOWL increase as the full-matching probability decreases. This result is expected since the IPWE-based approaches (BOWL and RWL) rely heavily on the number of full-matching treatments from the empirical samples in order to achieve convergence. In contrast, the proposed SWL model is more adaptive to all matching scenarios by design and thus is more robust against small full-matching rates. In addition, the relatively stable performance of the AIPW model

suggests that the augmentation term in the AIPW estimator might lessen the curse of full-matching compared to the IPWEs. However, under the scenario where the propensity mechanism can be easily found, the additional potential outcome model increases the computational burden of AIPW and therefore further deteriorates its performance compared to BOWL and SWL. Finally, we remark that the difference between the performance of BOWL and RWL indicates that the adoption of an unsmooth surrogate function involving all stages in (3) used in RWL induces more empirical optimization difficulty when the decision sequence is long; whereas the proposed smooth SWL objective function (14) maximizes regimes at each stage separately, and hence eases the optimization process.

6.3 Optimal regime function complexities

In this numerical study, we are interested in analyzing the sensitivity of our proposed method to the functional complexity of the underlying optimal regime. Specifically, we also consider the linear treatment regimes and include the homogeneous decision rule setting where the optimal rules are the same at all time points, i.e. $f_j^* = f_1^*$ for all $2 \leq j \leq T$. Note that the homogeneous rules can be oftentimes encountered in a high-frequency treatment session as the optimal regime is unlikely to update in a short period of time. We combine the results of four settings when $T = 10$ in Figure 3.

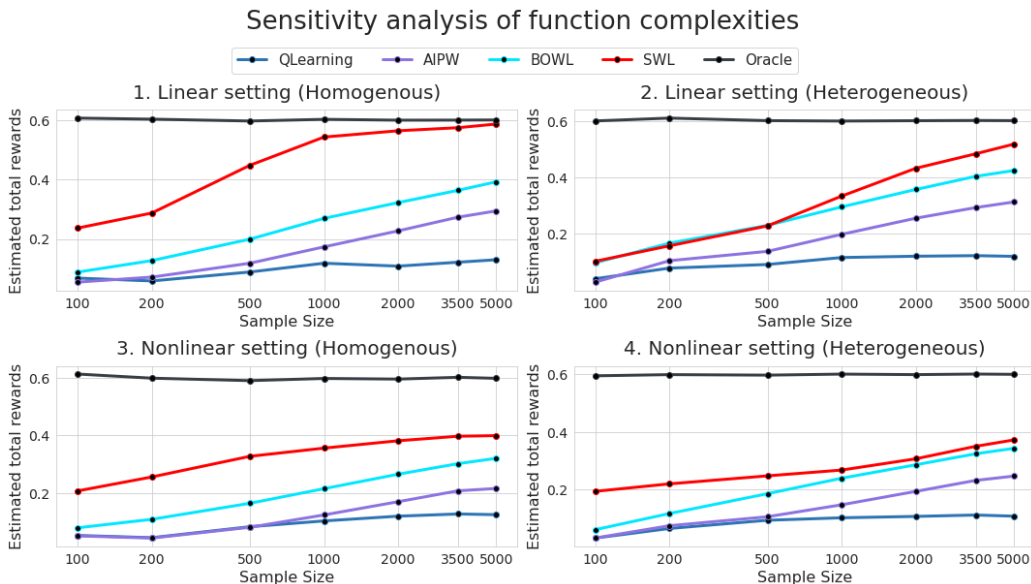


Figure 3: Sensitivity plots of estimated total rewards under four function settings against sample sizes. The number of decision stages is set to 10 for this example.

As algorithms are converging, a decrease in the performance improvement rate/slope is expected. According to this criterion, we observe that the proposed SWL method converges to Oracle faster than all other methods. For instance, under the linear homogeneous treatment rule setting, SWL achieves an averaged 95% matching accuracy, compared to 72% for BOWL when the sample size reaches 1000. Though the presented empirical results can be affected by the implementation and choices of hyper-parameters, based on the similar

performance increasing rate between SWL and BOWL when the sample size is larger than 2000, we can confirm with the developed theoretical results that our SWL reaches the same state-of-the-art asymptotic convergence rate as BOWL. In addition, from the deteriorated performance in the nonlinear heterogeneous setting compared to the linear homogeneous setting, we verify that the increasing level of functional complexity raises the optimization difficulties and hinders the model convergence rate with limited empirical examples.

6.4 Stage heterogeneity: number of important stages

In this subsection, we illustrate the advantages of incorporating stage heterogeneity with the stage importance scores. We adjust the level of heterogeneity by changing the number of important stages, where fewer important stages induce stronger heterogeneity. In addition, to maximize the heterogeneity among stages, we consider the linear homogeneous decision rule setting where the optimal rules are linear and stay the same at all time points, i.e. $f_j^* = f_1^*$ for all $2 \leq j \leq T$. Note that the homogeneous rules can be frequently encountered in a high-frequency treatment session, as the optimal regime is unlikely to be updated in a short period of time. Figure 4 provides the obtained results when $n = 500$ and $T = 10$.

As shown, while the proposed SAL with linear prior still outperforms the rest of the competing methods under the first three scenarios, the importance scores can further improve the performance of SAL with a greater margin when the heterogeneity among the stages gets stronger. Moreover, the stability of the SAL can be improved with the stage weights when strong stage heterogeneity exists. We conclude that the proposed stage importance scores are able to explicitly incorporate stage heterogeneity into the SAL framework and can be used for regime searching on stages which contribute to improving treatment effects.

Important-Stage counts	QLearning	AIPW	RWL	BOWL	SAL-linear	SWL	Oracle	Imp-rate (To SAL)
8	0.063 (0.098)	0.103 (0.071)	0.001 (0.052)	0.165 (0.073)	0.402 (0.31)	0.542 (0.269)	0.896 (0.037)	34.84%
5	0.098 (0.098)	0.2 (0.084)	-0.01 (0.047)	0.253 (0.087)	0.484 (0.283)	0.746 (0.157)	0.889 (0.025)	54.22%
3	0.114 (0.142)	0.304 (0.133)	0.007 (0.056)	0.361 (0.133)	0.522 (0.226)	0.811 (0.107)	0.883 (0.018)	55.28%
1	0.161 (0.253)	0.424 (0.169)	-0.007 (0.065)	0.594 (0.133)	0.452 (0.134)	0.782 (0.109)	0.828 (0.022)	72.88%

Table 3: The estimated total rewards of listed models under three scenarios of numbers of important stages when $n = 500$, $T = 10$, and optimal treatment rule is linear and homogeneous. Improvement rate compares SWL against SAL with linear prior.

7 Data Analysis

In this section, we apply the proposed method to UC COVID Research Data Sets (UC CORDS) (University of California Health, July 15, 2020), which combines timely COVID-related testing and hospitalization healthcare data from six University of California schools and systems. As of December 2022, UC CORDS include a total number of 108,914 COVID patients, where 31,520 of them had been hospitalized and 2,333 of them had been admitted to the ICU. Aiming to facilitate hospital management by reducing inpatients' length of stay at hospitals and further prevent them from developing more severe symptoms, we are interested in selecting effective treatments tailored to individual patients.

Linear setting (Homogeneous): value function v.s. number of important stages

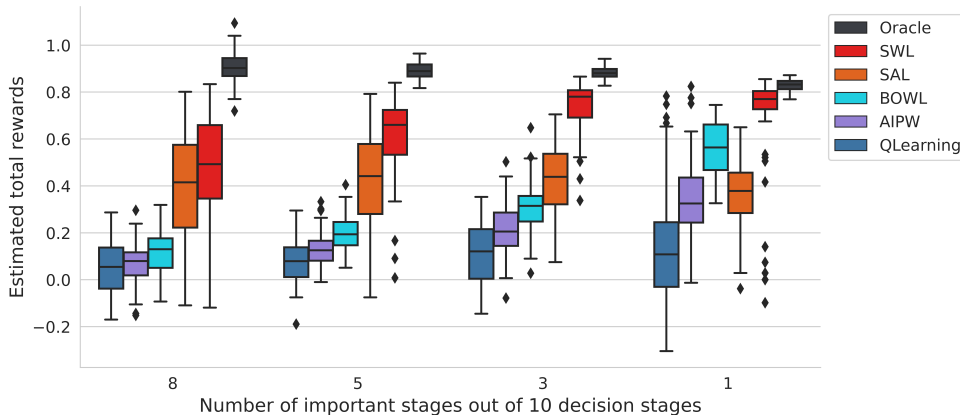


Figure 4: Boxplots of the estimated total rewards of listed methods versus the number of important stages when $n = 500$, $T = 10$, and the optimal treatment rule is linear and homogeneous.

One of the first few FDA-approved drugs that have been found effective against covid was Dexamethasone (Ahmed and Hassan, 2020). However, the precise treatment plan using Dexamethasone still remains unclear. As suggested by Waterer and Rello (2020), clinicians need to consider individual risks especially among elderly patients with age over 65 years old and patients with comorbidities, such as diabetes and cardiovascular diseases. In fact, according to the UC CORDS medical records, elderly patients spend 3 more days on average in both hospitals and ICUs compared to younger patients. Thus, our goal is to apply DTR methods to provide an optimal individualized treatment decision for Dexamethasone (i.e., whether the patient should take the drug or not) with the incorporation of heterogeneity among patients at a decision stage.

In this application, we list two emerging technical challenges. First, the number of decision stages involved can be large during the average 7 days of inpatient stay, and unlike the infinite horizon DTR method (Ertefaie and Strawderman, 2018; Zhou et al., 2022), a finite number of treatment stages is considered in this application. As a result, an efficient DTR estimation method should be robust against the curse of full-matching, e.g., long decision sequences and few patients receiving optimal treatments. Second, according to Lee et al. (2021) who found an early administration of Dexamethasone can reduce hospital stay, we speculate that stage heterogeneity exists in evaluating the treatment effect, and therefore examine methods which can incorporate timing effects on the estimation procedure. Taking the above two challenges into account, the proposed SAL framework and SWL are able to fulfill these needs for this real-data application.

We first pre-process UC CORDS data following the procedure elaborated in Supplemental Materials. Then we fit the proposed models and competing methods to search for the optimal DTR of Dexamethasone which can reduce the number of inpatient or ICU day stays for admitted patients. In this application, patients receiving a total number of 5, 8, and 10 treatment decisions during their stay in the hospital are included. We randomly select 80% of the data as a training set and repeat the process 20 times to obtain a Monte-Carlo sample of the model performance scores. All methods are evaluated under the empirical

value outcome according to Zhao et al. (2015), i.e.,

$$\hat{V}^d = \frac{\mathbb{E}_n \left\{ R \cdot \prod_{j=1}^T \mathbb{I}(a_j = d_j) / \prod_{j=1}^T \hat{\pi}_j(a_j | H_j) \right\}}{\mathbb{E}_n \left\{ \prod_{j=1}^T \mathbb{I}(a_j = d_j) / \prod_{j=1}^T \hat{\pi}_j(a_j | H_j) \right\}}. \quad (23)$$

To show the model performance, we list the estimated outcomes in terms of hospitalization days in Table 4 and Figure 5, where a smaller value indicates better model performance.

Stay Type	Number of Stage	QLearning	AIPW	RWL	BOWL	SAL-linear	SWL
ICU	5 (n=623)	18.217 (17.920)	15.251 (11.720)	10.455 (10.240)	11.720 (9.412)	3.305 (7.832)	5.855 (4.666)
	8 (n=345)	44.200 (16.167)	41.300 (17.254)	7.173 (6.237)	10.340 (8.295)	4.913 (3.375)	6.749 (7.636)
Inpatient	5 (n=3256)	10.108 (3.291)	10.003 (6.660)	9.493 (1.507)	11.525 (2.212)	9.019 (3.051)	7.790 (2.861)
	8 (n=1876)	24.950 (11.200)	19.815 (13.451)	9.589 (2.790)	7.012 (0.389)	6.704 (3.541)	5.664 (3.263)
	10 (n=1419)	27.712 (10.713)	27.800 (10.884)	11.412 (4.031)	15.129 (12.910)	5.582 (2.318)	4.817 (1.958)

Table 4: Estimated number of hospitalization days obtained from DTR methods under two different stay types: Inpatient and ICU. The sample size is listed next to the number of stages.

Both Table 4 and Figure 5 show that the SAL framework achieves the overall best performance in terms of reducing the length of hospital stay for patients following the suggested DTR. Specifically, the proposed method reduces almost 70% of the ICU duration from 10 days to 3 days, compared to the RWL method when a total of 5 decision stages is involved. Apart from the averaged performance, our methods attain leading model stability compared to other competing methods. In particular, the SWL can further improve the model stability of SAL under a larger sample case, where the stage importance scores can be better-captured by the attention-based neural network.




Figure 5: Boxplots of the estimated number of inpatient days by the number of important stages.

Based on our analyses, we can summarize that the Q-Learning has difficulties to estimate the outcome when a large number of decision stages are involved and the underlying reward mechanism is complicated, especially in the case of COVID where the association of recovery

time and Dexamethasone treatment still remains unclear. Meanwhile, the AIPW directly estimates the DTR and achieves slightly better results than the Q-Learning. But with the augmentation term, the AIPW might suffer similar challenges, such as outcome estimation and increasing computation burden. Note that RWL and BOWL are both based on the IPWE framework and require a sufficient number of patients to receive optimal treatments at all stages. Thus, as the number of stages increases and the number of involved patients decreases, both methods show deteriorating performance and larger variance. In comparison, our methods are able to combine the efficiency of IPWE-based methods and meanwhile improve model stability by taking into account heterogeneity-matching schemes between the observed and underlying optimal regimes. Our real-world application to the UC CORDS dataset demonstrates the superior empirical performance of the proposed SAL framework.

8 Discussion

In this paper, we propose a novel individualized learning framework for estimating the optimal dynamic treatment regime. The proposed framework utilizes the matching status between the observed and underlying optimal regime at any stage and substantially improves the sample efficiency of the inverse-probability-weighted approaches. With the stage importance scores, the proposed framework incorporates stage heterogeneity and therefore more accurately captures the differences in treatment effectiveness at various stages. In theory, we establish the Fisher consistency and a finite-sample performance error bound, which achieves the best convergence rate in the literature and provides a non-asymptotic explanation.

There are future improvements and extensions for our work. For example, to estimate the stage importance scores, we construct an attention-based neural network in the current work. As a result, the estimated importance scores remain fixed at all stages and provide a general interpretation of the importance of a stage among the patient population. However, the treatment stage heterogeneity could vary among the patients. As for future exploration, we can incorporate the multi-head attention architecture (Vaswani et al., 2017) which provides individualized stage importance scores at the patient level. In addition, we can provide a data-driven procedure to estimate the prior distributions for the SAL framework, and devise a smooth surrogate function for the non-convex component of the general k -IPWE.

Acknowledgments and Disclosure of Funding

All acknowledgements go at the end of the paper before appendices and references. Moreover, you are required to declare funding (financial activities supporting the submitted work) and competing interests (related financial activities outside the submitted work). More information about this disclosure can be found on the JMLR website.

References

Mukhtar H Ahmed and Arez Hassan. Dexamethasone for the treatment of coronavirus disease (covid-19): a review. *SN Comprehensive Clinical Medicine*, 2(12):2637–2646, 2020.

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- Peter L Bartlett, Michael I Jordan, and Jon D McAuliffe. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156, 2006.
- Doron Blatt, Susan A Murphy, and Ji Zhu. A-learning for approximate planning. *Ann Arbor*, 1001:48109–2122, 2004.
- Ronald DeVore, Boris Hanin, and Guergana Petrova. Neural network approximation. *Acta Numerica*, 30:327–444, 2021.
- Steven Diamond, Reza Takapoui, and Stephen Boyd. A general system for heuristic solution of convex problems over nonconvex sets. *arXiv preprint arXiv:1601.07277*, 2016.
- Damien Ernst, Pierre Geurts, and Louis Wehenkel. Tree-based batch mode reinforcement learning. *Journal of Machine Learning Research*, 6:503–556, 2005.
- Ashkan Ertefaie and Robert L Strawderman. Constructing dynamic treatment regimes over indefinite time horizons. *Biometrika*, 105(4):963–977, 2018.
- Pierre Geurts, Damien Ernst, and Louis Wehenkel. Extremely randomized trees. *Machine Learning*, 63(1):3–42, 2006.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8), 1997.
- Daniel G Horvitz and Donovan J Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*, 47(260):663–685, 1952.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Eric B Laber and Ying-Qi Zhao. Tree-based methods for individualized treatment regimes. *Biometrika*, 102(3):501–514, 2015.
- Hyun Woo Lee, Jimyung Park, Jung-Kyu Lee, Tae Yeon Park, and Eun Young Heo. The effect of the timing of dexamethasone administration in patients with covid-19 pneumonia. *Tuberculosis and Respiratory Diseases*, 84(3):217, 2021.
- Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016.
- Weibin Mo, Zhengling Qi, and Yufeng Liu. Learning optimal distributionally robust individualized treatment rules. *Journal of the American Statistical Association*, 116(534):659–674, 2021.
- Susan A Murphy. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(2):331–355, 2003.
- Inbal Nahum-Shani, Min Qian, Daniel Almirall, William E Pelham, Beth Gnagy, Gregory A Fabiano, James G Waxmonsky, Jihneeh Yu, and Susan A Murphy. Q-learning: a data analysis method for constructing adaptive interventions. *Psychological Methods*, 17(4):478, 2012.
- Zhengling Qi, Dacheng Liu, Haoda Fu, and Yufeng Liu. Multi-armed angle-based direct learning for estimating optimal individualized treatment rules with various outcomes. *Journal of the American Statistical Association*, 115(530):678–691, 2020.

- Min Qian and Susan A Murphy. Performance guarantees for individualized treatment rules. *Annals of Statistics*, 39(2):1180, 2011.
- Herbert Robbins and Sutton Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, 22(3):400–407, 1951.
- James Robins. A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Mathematical Modelling*, 7(9-12):1393–1512, 1986.
- James M Robins, Andrea Rotnitzky, and Lue Ping Zhao. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association*, 89(427):846–866, 1994.
- Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688, 1974.
- Donald B Rubin. Randomization analysis of experimental data: The Fisher randomization test comment. *Journal of the American Statistical Association*, 75(371):591–593, 1980.
- Phillip J Schulte, Anastasios A Tsiatis, Eric B Laber, and Marie Davidian. Q-and A-learning methods for estimating optimal dynamic treatment regimes. *Statistical Science: A review journal of the Institute of Mathematical Statistics*, 29(4):640, 2014.
- Juliana Schulz and Erica EM Moodie. Doubly robust estimation of optimal dosing strategies. *Journal of the American Statistical Association*, 116(533):256–268, 2021.
- Chengchun Shi, Rui Song, and Wenbin Lu. Robust learning for optimal treatment decision with np-dimensionality. *Electronic Journal of Statistics*, 10:2894, 2016.
- Chengchun Shi, Alin Fan, Rui Song, and Wenbin Lu. High-dimensional A-learning for optimal dynamic treatment regimes. *Annals of Statistics*, 46(3):925, 2018.
- Ingo Steinwart and Andreas Christmann. *Support Vector Machines*. Springer Science, 2008.
- Yebin Tao, Lu Wang, and Daniel Almirall. Tree-based reinforcement learning for estimating optimal dynamic treatment regimes. *The Annals of Applied Statistics*, 12(3):1914, 2018.
- Anastasios A Tsiatis, Marie Davidian, Shannon T Holloway, and Eric B Laber. *Dynamic Treatment Regimes: Statistical Methods for Precision Medicine*. Chapman and Hall/CRC, 2019.
- Grigorios Tsoumakas and Ioannis Katakis. Multi-label classification: An overview. *International Journal of Data Warehousing and Mining (IJDWM)*, 3(3):1–13, 2007.
- University of California Health. University of california health creates centralized data set to accelerate COVID-19 research, July 15, 2020. URL <https://www.universityofcalifornia.edu/press-room/university-california-health-creates-centralized-data-set-accelerate-covid-19-research>.
- Sara Van de Geer. *Empirical Processes in M-estimation*, volume 6. Cambridge Press, 2000.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 2017.

- Grant W Waterer and Jordi Rello. Steroids and COVID-19: We need a precision approach, not one size fits all. *Infectious Diseases and Therapy*, 9(4):701–705, 2020.
- Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine Learning*, 8(3):279–292, 1992.
- Fei Xue, Yanqing Zhang, Wenzhuo Zhou, Haoda Fu, and Annie Qu. Multicategory angle-based learning for estimating optimal dynamic treatment regimes with censored data. *Journal of the American Statistical Association*, 117(539):1438–1451, 2022.
- Baqun Zhang, Anastasios A Tsiatis, Marie Davidian, Min Zhang, and Eric Laber. Estimating optimal treatment regimes from a classification perspective. *Stat*, 1(1):103–114, 2012a.
- Baqun Zhang, Anastasios A Tsiatis, Eric B Laber, and Marie Davidian. A robust method for estimating optimal treatment regimes. *Biometrics*, 68(4):1010–1018, 2012b.
- Baqun Zhang, Anastasios A Tsiatis, Eric B Laber, and Marie Davidian. Robust estimation of optimal dynamic treatment regimes for sequential treatment decisions. *Biometrika*, 100(3):681–694, 2013.
- Ying-Qi Zhao, Donglin Zeng, Eric B Laber, and Michael R Kosorok. New statistical learning methods for estimating optimal dynamic treatment regimes. *Journal of the American Statistical Association*, 110(510):583–598, 2015.
- Ying-Qi Zhao, Eric B Laber, Yang Ning, Sumona Saha, and Bruce E Sands. Efficient augmentation and relaxation learning for individualized treatment rules using observational data. *The Journal of Machine Learning Research*, 20(1):1821–1843, 2019.
- Yingqi Zhao, Donglin Zeng, A John Rush, and Michael R Kosorok. Estimating individualized treatment rules using outcome weighted learning. *Journal of the American Statistical Association*, 107(499):1106–1118, 2012.
- Yufan Zhao, Michael R Kosorok, and Donglin Zeng. Reinforcement learning design for cancer clinical trials. *Statistics in Medicine*, 28(26):3294–3315, 2009.
- Wenzhuo Zhou, Ruoqing Zhu, and Annie Qu. Estimating optimal infinite horizon dynamic treatment regimes via pt-learning. *Journal of the American Statistical Association to appear*, 2022.
- Xin Zhou, Nicole Mayer-Hamblett, Umer Khan, and Michael R Kosorok. Residual weighted learning for estimating individualized treatment rules. *Journal of the American Statistical Association*, 112(517):169–187, 2017.