
Policy Learning for Optimal Individualized Dose Intervals

Guanhua Chen

University of Wisconsin, Madison
gchen25@wisc.edu

Xiaomao Li

University of Wisconsin, Madison
gilbert.xm.li@gmail.com

Menggang Yu

University of Wisconsin, Madison
meyu@biostat.wisc.edu

Abstract

We study the problem of learning individualized dose intervals using observational data. There are very few previous works for policy learning with continuous treatment, and all of them focused on recommending an optimal dose rather than an optimal dose interval. In this paper, we propose a new method to estimate such an optimal dose interval, named probability dose interval (PDI). The potential outcomes for doses in the PDI are guaranteed better than a pre-specified threshold with a given probability (e.g., 50%). The associated nonconvex optimization problem can be efficiently solved by the Difference-of-Convex functions (DC) algorithm. We prove that our estimated policy is consistent, and its risk converges to that of the best-in-class policy at a root- n rate. Numerical simulations show the advantage of the proposed method over outcome modeling based benchmarks. We further demonstrate the performance of our method in determining individualized Hemoglobin A1c (HbA1c) control intervals for elderly patients with diabetes.

1 INTRODUCTION

Increasing evidence has shown that personalized prescriptions based on patients' health conditions and medical history will result in more effective treatment (Mancinelli et al., 2000; Collins and Varmus, 2015). Hence, it is desirable to learn personalized treatment policy based on collected data and recommend treatment for future patients based on such a policy. We focus on off-policy learning, which is about evaluating

and optimizing new personalized treatment policies based on observational data.

Despite the abundance of research on policy learning for categorical treatments (Liang and Yu, 2020; Lou et al., 2018; Chen et al., 2017; Kosorok and Laber, 2019), policy learning for continuous treatments has not been studied sufficiently. Meanwhile, many applications in the medical domain involve determining optimal continuous treatment: the optimal dosage of a medical drug, e.g., Warfarin (International Warfarin Pharmacogenetics Consortium, 2009; Rich et al., 2014), and the optimal length of stay in the hospital for hospitalized patients (Borghans et al., 2012). Existing approaches for continuous treatments focused on recommending a fixed optimal dose for each patient. However, for many medical applications, recommending optimal dose interval can be more practical. For example, for chronic disease management metrics (e.g. HbA1c level, blood pressure, and cholesterol), targeted intervals, instead of a particular level, are frequently recommended from major medical organizations such as the American Diabetes Association (American Association of Clinical Endocrinologists, 2019) and the American Heart Association (Flack and Adekola, 2020). Another example is cancer radiation therapy, where radiation received by targeted tumors, not delivered by physicians, can be controlled only at an interval due to the irregular shape of solid tumors and the need to radiate surrounding regions (Scott et al., 2017). Therefore we focus on policy learning for optimal recommendation of an interval instead of a single dose. In addition, this shifted focus will help relax the local approximation needed for optimal dose finding, therefore improving both computational stability and theoretical properties, including a faster convergence rate for the target function over an existing method (Chen et al., 2016).

2 PROBLEM SETUP

2.1 Settings and Assumptions

Denote the dataset with n observations by (X_i, A_i, Y_i) , $i = 1, \dots, n$. Here, $X_i \in \mathcal{X}$ is a d -dimensional vec-

tor representing patient-level covariates. $A_i \in \mathcal{A} \equiv [a_L, a_U]$ is the observed treatment dose whose range has the lower and upper bounds (a_L and a_U). Y_i is the observed outcome from receiving A_i . We adopt the well-known potential outcome notation in causal inference (Imbens and Rubin, 2015) to articulate our research problem. For any dose $a \in \mathcal{A}$, the potential outcome $Y_i(a)$ is the outcome that would have been observed if patient i received treatment a . Note that when $A_i = a$, naturally $Y_i = Y_i(a)$ by the causal consistency assumption (Imbens and Rubin, 2015). An individualized policy or dose assignment rule is a functional $f(X) : \mathcal{X} \rightarrow \mathcal{A}$. For simplicity of presentation, we assume bigger $Y(a)$ is better. The optimal policy is the one that maximizes the policy value $V(f) \equiv E[Y(f(X))]$. To estimate this causal quantity using the observed data, we adopt the following three commonly assumed conditions (Imbens and Rubin, 2015): 1) the stable unit treatment value assumption (SUTVA); 2) the no unmeasured confounding assumption: $Y(a) \perp A|X$; and 3) the positivity assumption which assumes $P(A = a|X) > 0$ for all $a \in \mathcal{A}$. $P(A = a|X)$ is known as the (generalized) propensity score function that models the treatment assignment mechanism (Imbens and Rubin, 2015).

An Individualized Dose Interval (IDI) is desirable for a patient if the doses inside the interval are more likely associated with better outcomes than doses outside the interval. Each IDI is therefore connected with a threshold S that separates good outcomes from bad ones for the patient. In addition, a confidence probability α reflects how strongly the desirable interval is associated with the good outcomes. Specifically, let $f_L(x)$ and $f_U(x)$ be the lower and upper bounds of the IDI for a patient with covariates x such that $a_L \leq f_L(x) \leq f_U(x) \leq a_U$, $\forall x \in \mathcal{X}$. Then, we define a level- α *Probability Dose Interval* (PDI) as follows:

$$\text{PDI}_\alpha(x) \equiv [f_L(x), f_U(x)] : \forall a \in [f_L(x), f_U(x)], \quad (1) \\ P\{Y(a) > S|x\} \geq \alpha.$$

If a clinician follows the PDI recommendation when treating a patient with covariates x , the patient is guaranteed to have the outcome above S with a probability of at least α .

2.2 Related work

Indirect methods: Indirect methods for estimating treatment rules are based on outcome modeling (Zhao et al., 2009; Moodie et al., 2012; Chakraborty and Moodie, 2013). This framework has two steps. The first step models the conditional expectation of the outcome given the treatment and covariates, $E(Y|A, X)$. The models can be parametric such as regression with

Lasso-type regularization (Tibshirani, 1996), or non-parametric such as regression tree (Breiman et al., 1984; Breiman, 2001) and Support Vector Machine (Steinwart and Christmann, 2008; Vapnik, 2013). In the second step, the treatment option which maximizes the conditional expectation in the first step is set as the optimal treatment. For most cases, the second step involves numerical optimization such as grid search due to lack of close-form solutions. Since there are no existing approaches for PDI learning, we propose the following indirect PDI learning approach to serve as a benchmark. In particular, the first step for indirect PDI learning would construct an outcome model to estimate the conditional probability $P(Y > S | A = a, X = x)$, for example, based on a classifier for predicting the binary label $I\{Y > S\}$ using A and X . Subsequently, a grid search for doses is done to identify the dose interval where the predicted conditional probability is greater than the confidence level α . It has been noted in binary treatment settings that solutions from the indirect methods can be unstable due to the need to model $P(Y > S | A = a, X = x)$ for all a and x (Kosorok and Laber, 2019). Furthermore, the grid search step could be computationally challenging.

Direct methods: On the other hand, for optimal policy learning with categorical treatments, direct methods have been developed to learn the optimal policy directly, without using the outcome modeling. These approaches (Moodie et al., 2009; Xu et al., 2015; Lou et al., 2018; Chen et al., 2017; Liang and Yu, 2020) are believed to have better performance when signal-to-noise ratio is low. For learning optimal policy for continuous treatments, Chen et al. (2016) proposed an outcome weighted learning (O-learning) method which is based on a weighted regression. In particular, O-learning maximizes the value function associated with a policy f : i.e. $f_{opt} = \text{argmax}_f V(f)$ with $V(f) = E[Y(f(X))]$. Furthermore, Chen et al. (2016) and other direct methods (Kallus and Zhou, 2018; Bica et al., 2020; Sondhi et al., 2020; Zhu et al., 2020; Schulz and Moodie, 2020; Zhou et al., 2021) use inverse probability weighted (IPW) or augmented IPW estimators of the value function to address confounding issues by reweighting the observational data to mimic randomized trials. However, none of these approaches are directly applicable for the PDI problem.

3 DIRECT POLICY LEARNING FOR PDI

Because the recommended dose interval needs to satisfy probability constraint $P(Y(a) > S | X = x) \geq \alpha$ for $a \in \text{PDI}_\alpha(x)$, maximizing the associated value function for a dose interval becomes less intuitive. Instead,

we propose a loss function such that the estimated optimal policy that minimizes the loss function would (asymptotically) satisfy the constraint.

There are two types of errors that an PDI can incur. When a patient receives a dose inside the PDI, $a \in \text{PDI}_\alpha(x)$, but has an outcome below the threshold, $Y(a) < S$. It presents a false-positive error. On the other hand, when the patient receives a dose outside the interval, $a \notin \text{PDI}_\alpha(x)$, but has an outcome above the threshold, $Y(a) > S$. It presents a false-negative error. The goal is to minimize these two errors. Instead of working with a multi-objective optimization problem, we use α to balance the magnitudes of the two errors and work with the weighted sum, that is, $\alpha I\{Y(a) \leq S, a \in \text{PDI}_\alpha(x)\} + (1 - \alpha)I\{Y(a) > S, a \notin \text{PDI}_\alpha(x)\}$. This weighted sum leads to nice theoretical guarantee of the resulting interval for our proposed method as we demonstrate in Section 4 below.

The indirect methods rely on modeling the outcome Y to predict $I(Y(a) > S)$. Instead, we circumvent modeling of the outcome and directly minimize the above error based on observed data. Specifically, we consider the following risk function for minimization to obtain an estimate for $\text{PDI}_\alpha(x)$ with any given α , and S .

$$E \left[\left\{ \alpha I(Y \leq S) I(A \in \text{PDI}_\alpha(X)) + (1 - \alpha) I(Y > S) I(A \notin \text{PDI}_\alpha(X)) \right\} \frac{1}{P(A|X)} \right] \quad (2)$$

Here $P(A|X)$ is the generalized propensity score account for possible nonrandom treatment assignment in observational studies.

Under mild conditions, the minimizer of the risk function (2) becomes an interval $[f_L(x), f_U(x)]$ where $f_L(x)$ and $f_U(x)$ are boundary functions that minimize the following risk

$$R(f_L, f_U) = E \left[\left\{ \alpha I(Y \leq S) I(A \in [f_L(X), f_U(X)]) + (1 - \alpha) I(Y > S) I(A \notin [f_L(X), f_U(X)]) \right\} \frac{1}{P(A|X)} \right] \quad (3)$$

In cases when Conditions C1 and C2 defined in Section 4 are satisfied, then only the lower bound is needed to separate the desirable doses from the undesirable doses, i.e., $\text{PDI}_\alpha(x) = [f_L(X), a_U]$. Accordingly, we just need to minimize the following to estimate $f_L(x)$.

$$R(f_L) = E \left[\left\{ \alpha I(Y \leq S) I(A \in [f_L(X), a_U]) + (1 - \alpha) I(Y > S) I(A \in [a_L, f_L(X)]) \right\} \frac{1}{P(A|X)} \right] \quad (4)$$

The functional forms for the boundary functions $f_L(x)$ and $f_U(x)$ can be chosen quite flexibly in principle, as in many machine learning approaches. In practice, one may put some smoothness conditions on them for given x to ensure that a small perturbation of x will not lead to a dramatic change in the resulting PDI.

3.1 Non-Convex Surrogate Loss Relaxation

As in the classical classification tasks, it is difficult to directly optimize the empirical risk due to the discontinuity of indicator functions in (3) and (4). Surrogate loss functions have been proposed to replace the indicator functions as long as the surrogate loss functions satisfy certain conditions such as Fisher consistency. Note that unbounded loss functions such as the hinge loss are not good options here as demonstrated in the optimal dose problem (Chen et al., 2016). The reason is due to the sensitivity of the solution of the unbounded surrogate loss to outliers, that is, to observations whose received doses are far from the optimal bounds. Therefore we propose a truncated hinge loss as a surrogate similar to Chen et al. (2016). Such non-convex surrogate losses have also been used in robust SVM (Wu and Liu, 2007).

The risk function (4) of the one-sided $\text{PDI}_\alpha(X)$ after relaxation is as follows. See Figure 1 for a graphical representation.

$$R_\epsilon(f_L) = E \left[\left\{ \alpha I(Y \leq S) \Psi_\epsilon(A, f_L) + (1 - \alpha) I(Y > S) \Psi_\epsilon(f_L, A) \right\} \frac{1}{P(A|X)} \right] \quad (5)$$

where $\Psi_\epsilon(a, b) = \min\{\epsilon^{-1}(a - b)_+, 1\}$. Note that $\Psi_\epsilon(a, b) = \Psi_{\epsilon,1}(a, b) - \Psi_{\epsilon,2}(a, b)$, where $\Psi_{\epsilon,1}(a, b) = \{\epsilon^{-1}(a - b)\}_+$ and $\Psi_{\epsilon,2}(a, b) = \{\epsilon^{-1}(a - b) - 1\}_+$ are convex.


The risk function (3) of the two-sided $\text{PDI}_\alpha(X)$ after relaxation is as follows.

$$R_\epsilon(f_L, f_U) = E \left[\left\{ \alpha I(Y \leq S) \Psi_\epsilon^{in}(f_L, A, f_U) + (1 - \alpha) I(Y > S) \Psi_\epsilon^{out}(f_L, A, f_U) \right\} \frac{1}{P(A|X)} \right] \quad (6)$$

where

$$\Psi_\epsilon^{in}(a, b, c) = \begin{cases} 0, & b < a < c \\ (b - a)/\epsilon, & a < b < a + \epsilon < c \\ 1, & a + \epsilon < b < c - \epsilon, \text{ and} \\ (c - b)/\epsilon, & a < c - \epsilon < b < c \\ 0, & a < c < b \end{cases}$$

$$\Psi_\epsilon^{out}(a, b, c) = 1 - \Psi_\epsilon^{in}(a, b, c).$$


 Figure 1: Loss Functions $\Psi_\epsilon, \Psi_{\epsilon,1}$ and $\Psi_{\epsilon,2}$

3.2 DC Algorithm for One-sided PDI

Assume that f_L belongs to a reproducing kernel Hilbert space (RKHS) \mathcal{H} (Shawe-Taylor et al., 2004). We propose minimizing the following regularized empirical risk of $R_\epsilon(f_L)$ corresponding to (5):

$$\begin{aligned} \hat{f}_{L,n} = \arg \min_{f_L \in \mathcal{H}} \sum_{i=1}^n \left[\left\{ \alpha I(Y_i \leq S) \Psi_\epsilon(A_i, f_L(X_i)) \right. \right. \\ \left. \left. + (1 - \alpha) I(Y_i > S) \Psi_\epsilon(f_L(X_i), A_i) \right\} \frac{1}{P(A_i|X_i)} \right] \\ + \frac{\lambda_n}{2} \|f_L\|_{\mathcal{H}}, \end{aligned} \quad (7)$$

where the regularization term $\frac{\lambda_n}{2} \|f_L\|_{\mathcal{H}}^2$ penalizes the complexity of f_L and $\|\cdot\|_{\mathcal{H}}^2$ is a norm defined in \mathcal{H} .

Optimizing such an objective function is challenging due to the nonconvexity of the loss function. We adopt the Difference-of-Convex-functions (DC) algorithm (Le Thi Hoai and Tao, 1997) to iteratively solve a series of convex minimization problems. For simplicity, assume $f_L(X_i; \mathbf{v}) = \sum_{j=1}^n v_j k(X_i, X_j) + v_0$ where k is the reproducing kernel of \mathcal{H} and denote $\mathbf{v} = (v_1, \dots, v_n, v_0)$ for the unknown parameters. We can write the regularized empirical risk in (7) as

$$\begin{aligned} \mathcal{L}(\mathbf{v}) \equiv \sum_{i=1}^n \left[\left\{ \alpha I(Y_i \leq S) \Psi_\epsilon(A_i, f_L(X_i; \mathbf{v})) \right. \right. \\ \left. \left. + (1 - \alpha) I(Y_i > S) \Psi_\epsilon(f_L(X_i; \mathbf{v}), A_i) \right\} \frac{1}{P(A_i|X_i)} \right] \\ + \frac{\lambda_n}{2} \mathbf{v}^\top \mathbf{K} \mathbf{v} \end{aligned} \quad (8)$$

where \mathbf{K} is a $n \times n$ matrix with $k(X_i, X_j)$ ($i = 1 \dots n; j = 1, \dots, n$) as its elements. Let $\mathcal{L}_1(\mathbf{v})$ and $\mathcal{L}_2(\mathbf{v})$ be similarly defined as $\mathcal{L}(\mathbf{v})$ except with Ψ_ϵ replaced with $\Psi_{\epsilon,1}$ and $\Psi_{\epsilon,2}$ (without the regularization term) respectively in (8). Then $\mathcal{L}(\mathbf{v}) = \mathcal{L}_1(\mathbf{v}) - \mathcal{L}_2(\mathbf{v})$ and the non-convex objective function $\mathcal{L}(\mathbf{v})$ is the difference of two convex functions $\mathcal{L}_1(\mathbf{v})$ and $\mathcal{L}_2(\mathbf{v})$. By replacing the second function $\mathcal{L}_2(\mathbf{v})$ by its first order Taylor expansion expanded around the solution from previous iteration and the fact that a linear function is convex and concave, the DC algorithm solves a convex optimization problem at each iteration.

In summary, the DC algorithm estimates \mathbf{v} by repeatedly solving the following optimization problem

$$\begin{aligned} \mathbf{v}^{t+1} &= \arg \min_{\mathbf{v}} \mathcal{L}_1(\mathbf{v}) - \{\nabla \mathcal{L}_2(\mathbf{v}^t)\}^\top (\mathbf{v} - \mathbf{v}^t) \\ &= \arg \min_{\mathbf{v}} \mathcal{L}_1(\mathbf{v}) - \{\nabla \mathcal{L}_2(\mathbf{v}^t)\}^\top \mathbf{v}, \end{aligned}$$

where \mathbf{v}^{t+1} and \mathbf{v}^t are solutions of the DC-algorithm at $t+1$ -th and t -th iterations, respectively. The optimization problem above is a quadratic programming problem which can be solved by standard packages. The details of the derivations are documented in the supplementary. The convergence property of the DC algorithm has been shown in Lanckriet and Sriperumbudur (2009) such that the solution is guaranteed to converge to a local minimum.

So far, we only discussed the case of one-sided PDIs with a lower bound function. In the case where an upper bound is desired, one could also utilize the algorithm described above with a minor modification. By taking the negative value of the dose, the upper bound

of the original dose will become the lower bound for the negative dose. To find two-sided PDIs, one can use the following strategy to convert it into two separate one-sided PDI problems. Recall that the optimal dose function $f_{opt} = \arg \max_f V(f)$ defined in Section 2.2 and assume that $f_{L,opt}(x) < f_{opt}(x) < f_{U,opt}(x)$ for all $x \in \mathcal{X}$. Then, we can use samples with observed $A < f_{opt}(x)$ to learn the lower bound and samples with observed $A > f_{opt}(x)$ to learn the upper bound. To estimate the unknown function $f_{opt}(x)$, we can use the existing methods for optimal dose finding (Chen et al., 2016; Kallus and Zhou, 2018; Zhu et al., 2020; Zhou et al., 2021). The choice of which PDI (one-sided or two-sided) to use usually depends on practical considerations. When recommending HbA1c level for diabetes patients, because over-control HbA1c is the main concern, learning a one-sided PDI is sufficient (see Section 5.2 for details). In other settings, e.g. optimal personalized dose intervals for the Warfarin drug (an anticoagulant drug), a two-sided PDI may be preferred as overdosing on Warfarin will increase risk of bleeding, and underdosing will increase risk of stroke.

Note that PDIs depend on the threshold S and the probability α . The combinations of these quantities with different values could lead to different interpretations of the resulting PDIs. The pre-determined threshold S that defines the favorable outcomes can also be chosen as dependent on some covariates $z \subset x$, e.g., age or age groups.

3.3 Generalized Propensity Score for Continuous Treatment

To conduct off-policy learning for a continuous treatment, we need to estimate $P(a|x)$ which is the conditional density of dose a given the covariates x (Austin, 2018) and weigh each observation using the inverse of the estimated $P(a|x)$. Multiple approaches are available for estimating $P(a|x)$. One natural approach is regression. In particular, one could assume the density is from parametric or nonparametric distributions (Austin, 2018). This approach usually performs satisfactorily when the distributional assumptions are reasonable. For our simulation settings as well as the data application, the approach led to quite unstable inverse weights. To remedy the problem, we adopt recent methods that directly estimate the weights (instead of the inverse of estimated $P(a|x)$) that could balance the covariates across the range of a continuous treatment (Fong et al., 2018; Kallus and Santacatterina, 2019; Tübbicke, 2020; Huling et al., 2021). In particular, we use a modified version of Kernel Optimal Orthogonality Weighting (Kallus and Santacatterina, 2019) and the independence-inducing weights by Huling et al. (2021) for our numerical analysis due to their stable finite

sample performance.

4 THEORETICAL RESULTS

We include results for one-sided PDIs here and leave those for two-sided PDIs to the supplement. If $[f_L(x), a_U]$ is a one-sided level- α PDI, then it is reasonable to assume $[f_L(x) + \eta(x), a_U]$ is also a one-sided level- α PDI for any $\eta(x) \geq 0$ as it is a narrower interval. We ensure this by making the following assumptions throughout the section:

(C1) For any x , the boundaries a_L and a_U for the dose range satisfy $P\{Y(a_L) > S|x\} \leq \alpha$ and $P\{Y(a_U) > S|x\} \geq \alpha$.

(C2) For any x , there is a unique solution $a_x \in [a_L, a_U]$ such that $P\{Y(a_x) > S|x\} = \alpha$.

The first condition (C1) ensures that our choices of S and α are reasonable in the sense that the level- α is attainable. The second condition (C2) ensures that the limit of our learning algorithm is well defined. Details of the proof can be found in the supplementary file. Given (C1) and (C2), our solution $f_L(x)$ can attain the lowest possible value (i.e., the PDI is as wide as possible).

Note that if $P(Y > S|A = a, X = x)$ is a monotonic function in a for all x , i.e. $P(Y > S|A = a_1, X = X) \leq P(Y > S|A = a_2, X = X)$, $\forall a_1, a_2 \in \text{PDI}_\alpha(X)$ and $a_1 < a_2$, then Conditions (C1) and (C2) will hold with properly chosen α and $S(X)$. However, (C1) and (C2) do not require $P(Y > S|X = x)$ to be monotonic. In other words, the monotonicity is sufficient but not necessary for (C1) and (C2) to hold. We make assumptions that (C1) and (C2) hold to guarantee the existence of $\text{PDI}_\alpha(X)$. Hence, our method can work with multimodal outcomes as long as (C1) and (C2) are satisfied. On the other hand, for many medical applications, one can safely expect $P(Y > S|a, X)$ to be monotonic or unimodal in a . A unimodal $P(Y > S|a, X)$ will require two-sided PDIs.

4.1 Fisher Consistency

Theorem 4.1 below shows that minimizing $R(\cdot)$ defined in (4) will result in a one-sided PDI lower bound, above which the doses will produce outcomes larger than S with a probability of at least α .

Theorem 4.1. *Let $f_{L,opt} = \arg \min_f R(f)$. Then, under the assumptions (C1) and (C2), $f_{L,opt}(x)$ has the following two properties:*

(1) $P\{Y(f_{L,opt}(x)) > S|x\} = \alpha, \forall x$.

(2) For any measurable function $a(\cdot)$ s.t. $f_{L,opt}(x) \leq$

$a(x) \leq a_U, \forall x \in \mathcal{X}$, the potential outcome $Y(a(x))$ satisfies $P(Y(a(x)) > S|x) \geq \alpha$.

Theorem 4.1 shows that for any given value of $X = x$, the length of $[f_{L,opt}(x), a_U]$ is the widest among all valid $\text{PDI}_\alpha(x)$.

4.2 Convergence Rate

Theorem 4.2 below shows that, for one-sided PDI lower bounds, the difference between the risk function $R(\cdot)$ and its non-convex relaxation $R_{\epsilon_n}(\cdot)$ converges to zero with the same rate as $\epsilon_n \rightarrow 0$. That is, the approximation bias goes to 0. Then in Theorem 4.3 we derive the convergence rates of the risk function based on the estimated $f_{L,n}$ and the general results of empirical risk minimization and approximation. This theorem tells us how the empirically optimal policy performs out-of-sample. Here, we assume the Gaussian kernel with bandwidth γ_n (i.e. $\mathbf{K}(x, x') = \exp(-\gamma_n^2 \|x - x'\|^2)$) was used in estimating $\hat{f}_{L,n}$. Since $f_{L,opt}$ may not belong to the RKHS we optimize in, we choose the Gaussian kernel (which is dense in L_2) to make sure $\hat{f}_{L,n}$ can well approximate $f_{L,opt}$.

Theorem 4.2. *For any measurable function $f : \mathcal{X} \mapsto \mathbb{R}$, we have $|R(f) - R_{\epsilon_n}(f)| \leq C\epsilon_n$, where C is a constant.*

Theorem 4.3. *Let $B_{1,\infty}^\delta(\mathbb{R}^d)$ be a Besov space defined as $B_{1,\infty}^\delta(\mathbb{R}^d) = \{f \in L_\infty(\mathbb{R}^d) : \sup_{t>0} (t^{-\delta} w_{(r,L_1)}(f, t)) < \infty\}$, where w is the modulus of continuity. Assume that $f_{L,opt} \in B_{1,\infty}^\delta(\mathbb{R}^d)$. Then for any $\zeta > 0$, $d/(d + \tau) < p < 1$, and $\tau > 0$,*

$$R(\hat{f}_{L,n}) - R(f_{L,opt}) \leq c_1 \frac{\lambda_n}{\gamma_n^d} + c_2 \gamma_n^\delta + c_3 \gamma_n^{-\frac{(1-p)(1+\zeta)d}{2-p}} \lambda_n^{-\frac{p}{2-p}} n^{-\frac{1}{2-p}} + c_4 \frac{\tau^{\frac{1}{2}}}{n^{\frac{1}{2}}} + c_5 \frac{\tau}{n} + c_6 \epsilon_n$$

with probability $1 - 3e^{-\tau}$. Furthermore, if we choose these constants to be

$$\gamma_n = \mathcal{O}\left(\frac{1}{n}\right)^{\frac{1}{2\delta+d}} \quad \lambda_n = \mathcal{O}\left(\frac{1}{n}\right)^{\frac{\delta+d}{2\delta+d}} \quad \epsilon_n = \mathcal{O}\left(n^{-\frac{\delta}{2\delta+d}}\right).$$

We have the following optimal convergence rate: $R(\hat{f}_{L,n}) - R(f_{L,opt}) = \mathcal{O}\left(n^{-\frac{\delta}{2\delta+d}}\right)$.

Theorem 4.3 shows that the O-learning for PDI can potentially achieve a rate of nearly $\mathcal{O}(n^{-\frac{1}{2}})$, when the smoothness parameter δ of the optimal bound function $f_{L,opt}$ is relatively large compared to the covariate dimension d . This rate is faster than the convergence rate of about $\mathcal{O}(n^{-\frac{1}{4}})$ in Chen et al. (2016), which focused on finding the optimal dose. This improvement is due to the fact that more sample points are contained in the neighborhood of the entire optimal

interval than in the neighborhood of the optimal dose. Note that, we have assumed the generalized propensity score $P(a|x)$ can be estimated at the rate of $\mathcal{O}(n^{-1/2})$ or faster. Consequently, the estimation error of the generalized propensity score is negligible. However, if the generalized propensity score can not be estimated at this rate, then convergence rate in Theorem 4.3 will be dominated by the possibly slower convergence rate of the generalized propensity score estimator.

5 EXPERIMENTS

5.1 Simulation

We compare our methods with three indirect methods for learning one-sided $\text{PDI}_\alpha(X)$. Note that existing optimal dose learning methods can not yield a dose interval, hence they are not included in the comparison. For our methods, we implemented linear and Gaussian kernels and denote them as L-O-Learning and G-O-Learning. As the DC-algorithm is an iterative algorithm which relies on initial values, we used covariates-independent lower bounds as such initial values. All indirect methods involve classification of outcomes to find the PDIs. In particular, the indicator $I\{Y > S\}$ was used as the binary outcome for classification so that we could build a model for $\hat{P}(Y > S|a, X) \equiv h(a, X)$. Then we found the PDIs as a range of doses $\mathcal{A}_1(X)$ such that for $\forall a \in \mathcal{A}_1(X)$, $h(a, X) > \alpha$. To find $\mathcal{A}_1(X)$, we evaluated different a value in estimated $h(a, X)$. All a that satisfied $h(a, X) > \alpha$ were included in $\mathcal{A}_1(X)$. In our simulations, we used 200 evenly spaced grids to construct $\mathcal{A}_1(X)$.

We used logistic regression with L_1 penalty, Support Vector Machine (SVM), and Random Forests (RF) as indirect methods. For the penalized logistic regression, we included linear and quadratic terms of the dose, covariates, and pairwise interactions between dose and covariates as predictors. For SVM, we used the Gaussian kernel with the kernel bandwidth fixed using the ‘‘median heuristic’’. For RF, we set the number of trees to be 1000. All methods required parameter tuning and we tuned the L_2 regularization parameters for O-learnings, the L_1 regularization parameter for the logistic regression, the margin trade-off parameter for the SVM, and the percentage of variables that can be used for splitting at each node for the RF. Tuning parameters for all methods were selected using 5-fold cross-validation.

We considered multiple simulation settings with various data generating processes, sample sizes, numbers of covariates as well as signal to noise levels. We set the confidence probability α to be 0.5. Scenario 1 is for linear lower bounds and Scenario 2 for nonlinear lower

bounds. For both scenarios, there is a unique value of a such that $P(Y(a) > S(x)|x) = \alpha$ for a given x .

We assume that the training data contain n i.i.d. observations $(A_i, Y_i, X_i), i = 1 \dots n$. The dose A_i is generated from a truncated normal distribution $\mathcal{N}(\mu_A, 0.5)$ with the truncation limits $[-2, 2]$. X_i is generated from a uniform distribution $\text{Uniform}(-1, 1)^d$. The outcome Y is generated from a normal distribution $\mathcal{N}(\mu_Y, \sigma^2)$. We assume there exist confounders (i.e., X can impact both treatment assignment and the outcome) to mimic the observational study settings. In the supplement, we also include simulations for randomized trial settings.

Scenario 1: Linear bounds

$$\begin{aligned}\mu_A &= 0.3X_1 + 0.3X_2 + 0.3X_3; \\ \mu_Y &= 5e^{10A}/(e^{10\mu_A} + e^{10A}) + D(X); \\ D(X) &= 0.6X_2 + 0.6X_3 + 0.6X_4.\end{aligned}$$

Scenario 2: Non-linear bounds

$$\begin{aligned}\mu_A &= 0.75 \log(|X_1| + 1) - 0.2 \cos(\pi X_2) + 0.2I(X_3 > 0) - 0.4; \\ \mu_Y &= 5e^{10A}/(e^{10C} + e^{10A}) + D(X); \\ D(X) &= 0.4 \sin(\pi X_2) + 0.4I(X_3 > 0) + 0.4|X_4|.\end{aligned}$$

We repeat each setting for 100 times and report the results of all methods by evaluating the empirical risk or the empirical value of $R(\hat{f}_L)$ defined in Equation (4) on corresponding independent testing data sets with sample sizes of 10000. For O-learning methods, we estimate weights that could approximate $1/P(a|x)$ using the approaches mentioned in the Section 3.3. Meanwhile when evaluating the performances of all methods, we plug in the true values of $1/P(a|x)$ for calculating the empirical risks of estimated policies on the testing data sets.

For all methods we compared, we set S to be the fitted value of a polynomial regression of Y with X as predictors (i.e., set $S(x) = \hat{E}[Y|x]$), as there are strong prognostic effects. From Table 1, the L-O-Learning has advantages in the linear bound settings and the G-O-Learning has advantages in the nonlinear settings. In particular, compared to other methods, our methods are more robust to high covariate dimensionality (large d), low sample size (small n), and low signal-to-noise ratio (large σ^2) settings. In terms of computational time, the RF and L-O-Learning/G-O-Learning are comparable, and they are much faster than the SVM.

5.2 Study of HbA1c Control for Diabetic Patients

A fundamental aspect of managing diabetes is achieving tight control of blood sugar levels to reduce the risk of diabetes complications. Current diabetes guidelines recognize that treatment recommendations regarding tight control of HbA1c may not be appropriate for older

patients, especially those with comorbid conditions, due to increased risk for adverse events (American Diabetes Association, 2018). For example, complex patients may be at risk for negative outcomes when treated aggressively (Ismail-Beigi et al., 2011; Riddle and Karl, 2012). These increased risks have even more significance if the benefits of tight control are limited due to shortened life expectancy. If the long-term benefits are not relevant and the short-term risks are significant, there may be little reason to pursue tight control of HbA1c. As a proof-of-concept analysis, we would like to recommend one-sided PDI for the upper bound of HbA1c such that the number hospitalization events are minimized.

We use a de-identified electronic health record (EHR) dataset from a large health care system which contains laboratory records, information on primary care visits, and medication history of diabetes patients between 01/2003 to 12/2011. We use 45 relevant features and 8126 subjects whose observed HbA1c measures that are between 4% and 10% to learn the PDI of HbA1c. Details about the data set can be found in the supplement. For the purpose of notation consistency, we take the negative value of number of hospitalization events as the outcome. Hence a larger value of Y represents a preferred outcome. By this coding, we also set $S = -0.5$ and $\alpha = 0.5$, which leads to the interpretation of PDI such that if a patient controls his/her HbA1c less than the upper bound, they would likely not experience diabetes-related hospitalization (i.e., with the probability of hospitalization less than 0.5) in the next 90-days after HbA1c is measured. Since the observed outcome is a count, any choice of S that is between -1 and 0 will lead to the same result and interpretation as setting $S = -0.5$.

We randomly select 1000 patients as a training data set to learn the policy and use the remaining 7126 patients as a testing data set. We repeat this random splitting 100 times and report the averaged results. We apply all methods compared in our simulations to this real data. The tuning parameters are selected by using the same 5-fold cross-validation strategy.

To demonstrate the result, we plotted the distributions of the observed HbA1c level and the estimated upper bounds of samples in one testing data set in Figure 2. While the majority of patients have observed HbA1c between 5.5% and 7.5%, their estimated upper bounds of the L-O-Learning and G-O-Learning are generally between 6.5% and 8%, which suggests that some patients may be over-controlling their HbA1c levels. In comparison, the PDI recommendations from the comparison methods are highly tilted towards the upper limit, especially the SVM. It is unlikely that physicians will recommend such high HbA1c values to patients.

Table 1: Empirical risk ($R(\hat{f}_L)$) with SD in (.) from comparison methods under different settings

	n	d	L-O-Learning	G-O-Learning	Logistic	SVM	RF
Scenario 1 $\sigma^2 = 2.25$	200	10	0.048 (0.004)	0.049 (0.004)	0.054 (0.023)	0.053 (0.005)	0.053 (0.005)
	200	50	0.055 (0.005)	0.052(0.002)	0.176(0.074)	0.077 (0.041)	0.058 (0.004)
	400	10	0.045 (0.002)	0.047 (0.002)	0.044(0.002)	0.048 (0.002)	0.049 (0.002)
	400	50	0.056 (0.004)	0.051 (0.002)	0.071(0.025)	0.055 (0.003)	0.053 (0.002)
Scenario 1 $\sigma^2 = 9$	200	10	0.122 (0.004)	0.121 (0.003)	0.182(0.049)	0.139 (0.026)	0.134 (0.010)
	200	50	0.132 (0.005)	0.122(0.002)	0.215(0.049)	0.238 (0.032)	0.142 (0.014)
	400	10	0.120 (0.003)	0.121 (0.003)	0.142(0.037)	0.124 (0.007)	0.128 (0.005)
	400	50	0.129 (0.003)	0.122 (0.002)	0.205 (0.049)	0.139 (0.023)	0.127 (0.004)
Scenario 2 $\sigma^2 = 2.25$	200	10	0.054 (0.003)	0.050 (0.002)	0.069 (0.044)	0.055 (0.005)	0.052 (0.005)
	200	50	0.061 (0.005)	0.051 (0.002)	0.187 (0.078)	0.177 (0.081)	0.057 (0.006)
	400	10	0.051 (0.002)	0.049 (0.002)	0.051 (0.010)	0.051 (0.004)	0.049 (0.005)
	400	50	0.061 (0.003)	0.051 (0.002)	0.095(0.050)	0.060 (0.003)	0.051 (0.003)
Scenario 2 $\sigma^2 = 9$	200	10	0.123 (0.003)	0.120 (0.003)	0.177 (0.052)	0.144 (0.035)	0.131 (0.011)
	200	50	0.133 (0.005)	0.121 (0.002)	0.199 (0.053)	0.249 (0.010)	0.145 (0.016)
	400	10	0.121 (0.003)	0.120 (0.002)	0.143 (0.038)	0.123 (0.005)	0.127 (0.009)
	400	50	0.131 (0.004)	0.121 (0.002)	0.198 (0.053)	0.153 (0.040)	0.125 (0.005)

Furthermore, the grid search range of HbA1c levels for the indirect methods is between 4% and 10%, and indirect methods recommended a high proportion of patients to have 10% as the HbA1c levels upper bound (logistic regression: 30% of all patients; SVM: 70%; RF: 50%). If we change the grid search range to be between 4% and 9%, the indirect methods would recommend a high proportion of patients to have 9% as the upper bound. This indicates the policies estimated from indirect methods can have unwanted dependence on the grid search range. Such patterns of the estimated upper bounds of HbA1c levels shown in Figure 2 are similar across all 100 replications.

Furthermore, the upper bounds learned from our methods (especially the G-O-Learning) have much better agreement with the current medical guidance compared to other methods. In particular, the 2018 American Diabetes Association (ADA) Standards of Medical Care in Diabetes advises the following upper HbA1c bounds for diabetes patients (American Diabetes Association, 2018): 6.5% for people who can achieve this goal without experiencing a lot of hypoglycemia episodes or other negative effects of having lower blood glucose levels; 7.5% for all children with diabetes; 8% for people with a history of severe hypoglycemia, people who have had diabetes for many years or have difficulty achieving tighter control, and people with limited remaining life expectancy. In our dataset, 88% of the patients are over 60 years old, 50% over 73 years old, and 20% over 80 years old. Most of these patients have diabetes for many years. Hence, the distributions of upper bounds recommendations from our methods coincide well with the standards given the age distribution.

6 CONCLUSION

We have proposed a framework for learning individualized dose intervals with desired theoretical properties. Our method extends a previous method that can only learn the policy for an optimal dosage (Chen et al., 2016). We proposed a DC algorithm to efficiently solve the underlying non-convex optimization problems with linear or nonlinear policies learning. Note that both direct and indirect methods proposed by us can learn PDI, meanwhile, we advocate using the direct method since it can recommend dose intervals without the need of fitting complex regression models. Learning PDI relies on the function restriction due to Conditions C1 and C2 in Section 4, and such a restriction applies to both direct and indirect methods. If the conditions are not satisfied, the PDI may not be well defined or it will be a union of several discontinuous intervals.

We also believe that S should be chosen to have good interpretation. In our HbA1c example, choosing S between -1 to 0 leads to the same interpretation of the resulting PDI. For the applications where there is no obvious choice of S , we could set it to be a number between $\alpha \times 100\%$ -th percentile of $Y(a_L)|X$ and $\alpha \times 100\%$ -th percentile of $Y(a_U)|X$ to ensure the existence of PDI.

Lastly, learning causally meaningful PDIs using observed data relies on the no-unmeasured confounder assumption. In our HbA1c example, we worked closely with our physician collaborator to include curated variables from EHR and medical claims to ensure this assumption. At the same time, sensitivity analysis may help examine the robustness of the PDI policy. We want to emphasize that before implementing an




Figure 2: Density plot of observed and estimated upper bounds of HbA1c PDIs.

estimated policy for PDI in practice, it is critically important to verify the superiority of the estimated policy over existing clinical practice via randomized trials.

Acknowledgments

We would like to thank all reviewers for their helpful comments that greatly improved this paper. Research reported in this work was partially funded through a Patient-Centered Outcomes Research Institute (PCORI) Award (ME-2018C2-13180). The views in this work are solely the responsibility of the authors and do not necessarily represent the views of the Patient-Centered Outcomes Research Institute (PCORI), its Board of Governors or Methodology Committee.

References

- American Association of Clinical Endocrinologists (2019). Management of common comorbidities of diabetes.
- American Diabetes Association (2018). Standards of medical care in diabetes—2018 abridged for primary care providers. *Clinical diabetes: a publication of the American Diabetes Association*, 36(1):14.
- Austin, P. C. (2018). Assessing the performance of the generalized propensity score for estimating the effect of quantitative or continuous exposures on binary outcomes. *Statistics in Medicine*, 37(11):1874–1894.
- Bica, I., Jordon, J., and van der Schaar, M. (2020). Estimating the effects of continuous-valued interventions using generative adversarial networks. *arXiv preprint arXiv:2002.12326*.
- Borghans, I., Kleefstra, S. M., Kool, R. B., and Westert, G. P. (2012). Is the length of stay in hospital correlated with patient satisfaction? *International Journal for Quality in Health Care*, 24(5):443–451.
- Breiman, L. (2001). Random forests. *Machine learning*, 45(1):5–32.
- Breiman, L., Friedman, J., Stone, C. J., and Olshen, R. A. (1984). *Classification and regression trees*. CRC press.
- Chakraborty, B. and Moodie, E. (2013). *Statistical methods for dynamic treatment regimes*. Springer.
- Chen, G., Zeng, D., and Kosorok, M. R. (2016). Personalized dose finding using outcome weighted learning. *Journal of the American Statistical Association*, 111(516):1509–1521.
- Chen, S., Tian, L., Cai, T., and Yu, M. (2017). A general statistical framework for subgroup identification and comparative treatment scoring. *Biometrics*, 73(4):1199–1209.
- Collins, F. S. and Varmus, H. (2015). A new initiative on precision medicine. *New England journal of medicine*, 372(9):793–795.
- Eberts, M., Steinwart, I., et al. (2013). Optimal regression rates for svms using gaussian kernels. *Electronic Journal of Statistics*, 7:1–42.
- Flack, J. M. and Adekola, B. (2020). Blood pressure and the new acc/aha hypertension guidelines. *Trends in cardiovascular medicine*, 30(3):160–164.
- Fong, C., Hazlett, C., Imai, K., et al. (2018). Covariate balancing propensity score for a continuous treatment: Application to the efficacy of political advertisements. *The Annals of Applied Statistics*, 12(1):156–177.

- Freedman, B. I., Shenoy, R. N., Planer, J. A., Clay, K. D., Shihabi, Z. K., Burkart, J. M., Cardona, C. Y., Andries, L., Peacock, T. P., Sabio, H., et al. (2010). Comparison of glycated albumin and hemoglobin a1c concentrations in diabetic subjects on peritoneal and hemodialysis. *Peritoneal Dialysis International*, 30(1):72–79.
- Huling, J. D., Greifer, N., and Chen, G. (2021). Independence weights for causal inference with continuous exposures. *arXiv preprint arXiv:2107.07086*.
- Imbens, G. W. and Rubin, D. B. (2015). *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press.
- International Warfarin Pharmacogenetics Consortium (2009). Estimation of the warfarin dose with clinical and pharmacogenetic data. *New England Journal of Medicine*, 360(8):753–764.
- Ismail-Beigi, F., Moghissi, E., Tiktin, M., Hirsch, I. B., Inzucchi, S. E., and Genuth, S. (2011). Individualizing glycemic targets in type 2 diabetes mellitus: implications of recent clinical trials. *Annals of internal medicine*, 154(8):554–559.
- Kallus, N. and Santacatterina, M. (2019). Kernel optimal orthogonality weighting: A balancing approach to estimating effects of continuous treatments. *arXiv preprint arXiv:1910.11972*.
- Kallus, N. and Zhou, A. (2018). Policy evaluation and optimization with continuous treatments. In *International Conference on Artificial Intelligence and Statistics*, pages 1243–1251.
- Kosorok, M. R. and Laber, E. B. (2019). Precision medicine. *Annual Review of Statistics and Its Application*, 6:263–286.
- Lanckriet, G. R. and Sriperumbudur, B. K. (2009). On the convergence of the concave-convex procedure. In *Advances in Neural Information Processing Systems*, pages 1759–1767.
- Le Thi Hoai, A. and Tao, P. D. (1997). Solving a class of linearly constrained indefinite quadratic problems by DC algorithms. *Journal of global optimization*, 11(3):253–285.
- Liang, M. and Yu, M. (2020). A semiparametric approach to model effect modification. *Journal of the American Statistical Association*, 0(0):1–13.
- Little, R. R., Rohlfing, C. L., and Sacks, D. B. (2011). Status of hemoglobin a1c measurement and goals for improvement: from chaos to order for improving diabetes care. *Clinical chemistry*, 57(2):205–214.
- Lou, Z., Shao, J., and Yu, M. (2018). Optimal treatment assignment to maximize expected outcome with multiple treatments. *Biometrics*, pages 506–516.
- Luedtke, A. R. and van der Laan, M. J. (2016). Comment on: ‘personalized dose finding using outcome weighted learning’. *Journal of the American Statistical Association*, 111(516):1526–1530.
- Mancinelli, L., Cronin, M., and Sadée, W. (2000). Pharmacogenomics: the promise of personalized medicine. *Aaps Pharmsci*, 2(1):29–41.
- Moodie, E. E., Chakraborty, B., and Kramer, M. S. (2012). Q-learning for estimating optimal dynamic treatment rules from observational data. *Canadian Journal of Statistics*, 40(4):629–645.
- Moodie, E. E., Platt, R. W., and Kramer, M. S. (2009). Estimating response-maximized decision rules with applications to breastfeeding. *Journal of the American Statistical Association*, 104(485):155–165.
- Nitin, S. (2010). Hba1c and factors other than diabetes mellitus affecting it. *Singapore Med J*, 51(8):616–622.
- Radin, M. S. (2014). Pitfalls in hemoglobin a1c measurement: when results may be misleading. *Journal of general internal medicine*, 29(2):388–394.
- Rich, B., Moodie, E. E., and Stephens, D. A. (2014). Simulating sequential multiple assignment randomized trials to generate optimal personalized warfarin dosing strategies. *Clinical trials*, 11(4):435–444.
- Riddle, M. C. and Karl, D. M. (2012). Individualizing targets and tactics for high-risk patients with type 2 diabetes: practical lessons from accord and other cardiovascular trials. *Diabetes Care*, 35(10):2100–2107.
- Schulz, J. and Moodie, E. E. (2020). Doubly robust estimation of optimal dosing strategies. *Journal of the American Statistical Association*, pages 1–13.
- Scott, J. G., Berglund, A., Schell, M. J., Mihaylov, I., Fulp, W. J., Yue, B., Welsh, E., Caudell, J. J., Ahmed, K., Strom, T. S., et al. (2017). A genome-based model for adjusting radiotherapy dose (gard): a retrospective, cohort-based study. *The lancet oncology*, 18(2):202–211.
- Shawe-Taylor, J., Cristianini, N., et al. (2004). *Kernel methods for pattern analysis*. Cambridge university press.
- Sondhi, A., Arbour, D., and Dimmery, D. (2020). Balanced off-policy evaluation in general action spaces. In *International Conference on Artificial Intelligence and Statistics*, pages 2413–2423. PMLR.
- Steinwart, I. and Christmann, A. (2008). *Support vector machines*. Springer Science & Business Media.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288.
- Tübbsicke, S. (2020). Entropy balancing for continuous treatments. *arXiv preprint arXiv:2001.06281*.

- Vapnik, V. (2013). *The nature of statistical learning theory*. Springer science & business media.
- Wu, Y. and Liu, Y. (2007). Robust truncated hinge loss support vector machines. *Journal of the American Statistical Association*, 102(479):974–983.
- Xu, Y., Yu, M., Zhao, Y.-Q., Li, Q., Wang, S., and Shao, J. (2015). Regularized outcome weighted subgroup identification for differential treatment effects. *Biometrics*, 71(3):645–653.
- Zhao, Y., Kosorok, M. R., and Zeng, D. (2009). Reinforcement learning design for cancer clinical trials. *Statistics in medicine*, 28(26):3294–3315.
- Zhou, W., Zhu, R., and Zeng, D. (2021). A parsimonious personalized dose-finding model via dimension reduction. *Biometrika*, 108(3):643–659.
- Zhu, L., Lu, W., Kosorok, M. R., and Song, R. (2020). Kernel assisted learning for personalized dose finding. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '20, page 56–65.

Supplementary Material: Policy Learning for Optimal Individualized Dose Intervals

A DC ALGORITHM FOR ONE-SIDED PDI

Let $f_L(X_i; \mathbf{v}) = \sum_{j=1}^n v_j k(X_i, X_j) + v_0$ and $\mathbf{v} = (v_1, \dots, v_n, v_0)$. As claimed in Section 3.2 in the main paper, the DC algorithm repeatedly updates

$$\mathbf{v}^{t+1} = \arg \min_{\mathbf{v}} (\mathcal{L}_1(\mathbf{v}) - [\nabla \mathcal{L}_2(\mathbf{v}^t)]^\top (\mathbf{v} - \mathbf{v}^t)). \quad (9)$$

We derive the form of the quadratic programming by first showing that

$$\begin{aligned} \nabla \mathcal{L}_2(\mathbf{v}) &= \nabla_{\mathbf{v}} \left(\sum_{i=1}^n \frac{1}{P(A_i|X_i)} \left[(1-\alpha)I(Y_i > S)\Psi_{\epsilon,2}(f_L(X_i), A_i) + \alpha I(Y_i \leq S)\Psi_{2,\epsilon}(A_i, f_L(X_i)) \right] \right) \\ &= \sum_{i=1}^n \frac{1}{P(A_i|X_i)} \left[(1-\alpha)I(Y_i > S)Q_i(\beta) - \alpha I(Y_i \leq S)\tilde{Q}_i(\beta) \right] \nabla_{\mathbf{v}} f_L(X_i) \end{aligned}$$

where $\nabla_{\mathbf{v}} f_L(X_i) = (k(X_i, X_1), \dots, k(X_i, X_n), 1)$, $Q_i(\mathbf{v}) = I(f_L(X_i; \mathbf{v}) - A_i > \epsilon)/\epsilon$ and $\tilde{Q}_i(\mathbf{v}) = I(A_i - f_L(X_i; \mathbf{v}) > \epsilon)/\epsilon$

Following that, we have

$$[\nabla \mathcal{L}_2(\mathbf{v}^t)]^\top \mathbf{v} = \sum_{i=1}^n \frac{1}{P(A_i|X_i)} \left[(1-\alpha)I(Y_i > S)Q_i(\mathbf{v}^t) - \alpha I(Y_i \leq S)\tilde{Q}_i(\mathbf{v}^t) \right] \left\{ \sum_{j=1}^n v_j k(X_i, X_j) + v_0 \right\}$$

By plugging in the above results into the DC procedure,

$$\begin{aligned} \mathbf{v}^{t+1} &= \arg \min_{\mathbf{v}} \mathcal{L}_1(\mathbf{v}) - [\nabla \mathcal{L}_2(\mathbf{v}^t)]^\top (\mathbf{v} - \mathbf{v}^t) = \arg \min_{\mathbf{v}} \mathcal{L}_1(\mathbf{v}) - [\nabla \mathcal{L}_2(\mathbf{v}^t)]^\top \mathbf{v} \\ &= \arg \min_{\mathbf{v}} \sum_{i=1}^n \frac{1}{P(A_i|X_i)} \left((1-\alpha)I(Y_i > S) \left[\Psi_{1,\epsilon}(f_L(X_i; \mathbf{v}), A_i) - Q_i(\mathbf{v}^t) \left(\sum_{j=1}^n v_j k(X_i, X_j) + v_0 \right) \right] \right. \\ &\quad \left. + \alpha I(Y_i \leq S) \left[\Psi_{1,\epsilon}(A_i, f_L(X_i; \mathbf{v})) + \tilde{Q}_i(\mathbf{v}^t) \left(\sum_{j=1}^n v_j k(X_i, X_j) + v_0 \right) \right] \right) + \frac{\lambda_n}{2} \sum_{i=1}^n \sum_{j=1}^n v_i v_j k(X_i, X_j) \quad (10) \end{aligned}$$

To solve the above optimization problem, we now show that it is essentially a quadratic programming problem which can be easily solved using standard optimization packages.

Let $H_i = \frac{(1-\alpha)I(Y_i > S)}{P(A_i|X_i)\lambda_n}$ and $\tilde{H}_i = \frac{\alpha I(Y_i \leq S)}{P(A_i|X_i)\lambda_n}$. Minimizing Equation (10) is equivalent to

$$\begin{aligned} \arg \min_{\mathbf{v}} \quad & \sum_{i=1}^n H_i \xi_i + \sum_{i=1}^n \tilde{H}_i \tilde{\xi}_i + \sum_{i=1}^n (\tilde{H}_i \tilde{Q}_i^t - H_i Q_i^t) \left(\sum_{j=1}^n v_j k(X_i, X_j) + v_0 \right) + \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n v_i v_j k(X_i, X_j) \\ \text{s.t.} \quad & \xi_i - \left[\left(\sum_{j=1}^n v_j k(X_i, X_j) + v_0 \right) - A_i \right] / \epsilon \geq 0, & \xi_i \geq 0 \\ & \tilde{\xi}_i - \left[A_i - \left(\sum_{j=1}^n v_j k(X_i, X_j) + v_0 \right) \right] / \epsilon \geq 0, & \tilde{\xi}_i \geq 0 \end{aligned}$$

Applying the Lagrangian multiplier, the problem above is equivalent to

$$\begin{aligned}
 & \max_{\alpha_i, \tilde{\alpha}_i, \mu_i, \tilde{\mu}_i} \min_{\mathbf{v}, \xi} \sum_{i=1}^n H_i \xi_i + \sum_{i=1}^n \tilde{H}_i \tilde{\xi}_i + \sum_{i=1}^n (\tilde{H}_i \tilde{Q}_i^t - H_i Q_i^t) \left(\sum_{j=1}^n v_j k(X_i, X_j) + v_0 \right) + \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n v_i v_j k(X_i, X_j) \\
 & - \sum_{i=1}^n \alpha_i (\xi_i - [(\sum_{j=1}^n v_j k(X_i, X_j) + v_0) - A_i] / \epsilon) - \sum_{i=1}^n \tilde{\alpha}_i (\tilde{\xi}_i - [A_i - (\sum_{j=1}^n v_j k(X_i, X_j) + v_0)] / \epsilon) - \sum_{i=1}^n \mu_i \xi_i - \sum_{i=1}^n \tilde{\mu}_i \tilde{\xi}_i \\
 \text{s.t.} \quad & \sum_{i=1}^n v_i k(X_i, X_j) + \sum_{i=1}^n (\tilde{H}_i \tilde{Q}_i^t - H_i Q_i^t + \alpha_i / \epsilon - \tilde{\alpha}_i / \epsilon) k(X_i, X_j) = 0 \quad \forall j = 1, \dots, n \\
 & \sum_{i=1}^n (\tilde{H}_i \tilde{Q}_i^t - H_i Q_i^t + \alpha_i / \epsilon - \tilde{\alpha}_i / \epsilon) = 0 \\
 & H_i - \alpha_i - \mu_i = 0 \quad \forall i = 1, \dots, n \\
 & \tilde{H}_i - \tilde{\alpha}_i - \tilde{\mu}_i = 0 \quad \forall i = 1, \dots, n
 \end{aligned} \tag{11}$$

Incorporating the constraint that $\alpha_i \geq 0$, $\tilde{\alpha}_i \geq 0$, $\mu_i \geq 0$, $\tilde{\mu}_i \geq 0$ and denoting $\beta_i = \tilde{H}_i \tilde{Q}_i^t - H_i Q_i^t + \alpha_i / \epsilon - \tilde{\alpha}_i / \epsilon$, we have the final quadratic programming problem as

$$\begin{aligned}
 \min_{\beta} \quad & \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \beta_i k(X_i, X_j) \beta_j + \sum_{i=1}^n \beta_i A_i \\
 \text{s.t.} \quad & \tilde{H}_i \tilde{Q}_i^t - H_i Q_i^t - \tilde{H}_i / \epsilon \leq \beta_i \leq \tilde{H}_i \tilde{Q}_i^t - H_i Q_i^t + H_i / \epsilon \quad \forall i = 1, \dots, n \\
 & \sum_{i=1}^n \beta_i = 0
 \end{aligned} \tag{12}$$

The $t + 1$ step solution of (v_1, \dots, v_n) is thus attained from

$$\sum_{i=1}^n v_i k(X_i, X) = - \sum_{i=1}^n \beta_i k(X_i, X)$$

As for the intercept v_0 , although there exists an analytic solution by applying the KKT conditions, we solve it using a line search for simplicity. We iteratively solve the quadratic programming problem until convergence.

B PROOF OF THEOREMS FOR ONE-SIDED PDI

B.1 Proof of Theorem 4.1

Proof. First, let $f^* \in F \equiv \{f : \forall x \in \mathcal{X}, P(Y(f(x)) > S) = \alpha\}$. For any f' such that $f' \neq f^*$, denote $\mathcal{X}_0 = \{x : x \in \mathcal{X}, \exists f \in F \text{ s.t. } f'(x) = f(x)\}$, $\mathcal{X}_1 = \{x : x \in \mathcal{X}, f'(x) > \max_{f \in F} f(x)\}$, and $\mathcal{X}_2 = \{x : x \in \mathcal{X}, f'(x) < \min_{f \in F} f(x)\}$

Under the ignorability assumption, $P(Y > S | A = f'(x), x) = P(Y(f'(x)) > S | x) > \alpha$ for $x \in \mathcal{X}_1$, while $P(Y > S | A = f'(x), x) = P(Y(f'(x)) > S | x) < \alpha$ for $x \in \mathcal{X}_2$.

Then

$$\begin{aligned}
 & R(f') - R(f^*) \\
 &= E \left[\frac{1}{P(A|X)} \left((1-\alpha)I(Y > S) \left(I(A < f'(X)) - I(A < f^*(X)) \right) \right. \right. \\
 &\quad \left. \left. + \alpha I(Y \leq S) \left(I(A > f'(X)) - I(A > f^*(X)) \right) \right) \right] \\
 &= E \left[\int_{f^*(X)}^{f'(X)} \left((1-\alpha)P(Y > S|A=a, X) \right. \right. \\
 &\quad \left. \left. - \alpha P(Y \leq S|A=a, X) \right) da \middle| X \in \mathcal{X}_1 \right] P(X \in \mathcal{X}_1) \tag{A1}
 \end{aligned}$$

$$\begin{aligned}
 &+ E \left[\int_{f'(X)}^{f^*(X)} \left(- (1-\alpha)P(Y > S|A=a, X) \right. \right. \\
 &\quad \left. \left. + \alpha P(Y \leq S|A=a, X) \right) da \middle| X \in \mathcal{X}_2 \right] P(X \in \mathcal{X}_2) \tag{A2}
 \end{aligned}$$

$$\begin{aligned}
 &+ E \left[\int_{f^*(X)}^{f'(X)} \left((1-\alpha)\alpha - \alpha(1-\alpha) \right) da \middle| X \in \mathcal{X}_0 \right] P(X \in \mathcal{X}_0) \tag{A3}
 \end{aligned}$$

Here

$$\begin{aligned}
 & (A3) = 0 \\
 & (A1) = E \left[\int_{f^*(X)}^{f'(X)} \left(P(Y > S|A=a, X) - \alpha \right) da \middle| X \in \mathcal{X}_1 \right] P(X \in \mathcal{X}_1) \\
 &= E \left[\int_{\max_{f \in \mathcal{F}} f(X)}^{f'(X)} \left(P(Y > S|A=a, X) - \alpha \right) da \middle| X \in \mathcal{X}_1 \right] P(X \in \mathcal{X}_1) \\
 &\geq 0 \\
 &> 0 \quad \text{if } P(X \in \mathcal{X}_1) \neq 0
 \end{aligned}$$

$$\begin{aligned}
 & (A2) = E \left[\int_{f'(X)}^{f^*(X)} \left(- P(Y > S|A=a, X) + \alpha \right) da \middle| X \in \mathcal{X}_2 \right] P(X \in \mathcal{X}_2) \\
 &= E \left[\int_{f'(X)}^{\min_{f \in \mathcal{F}} f(X)} \left(- P(Y > S|A=a, X) + \alpha \right) da \middle| X \in \mathcal{X}_2 \right] P(X \in \mathcal{X}_2) \\
 &\geq 0 \\
 &> 0 \quad \text{if } P(X \in \mathcal{X}_2) \neq 0
 \end{aligned}$$

Therefore, $R(f') - R(f^*) \geq 0$. And $R(f') - R(f^*) = 0$ if and only if $P(X \in \mathcal{X}_1) = P(X \in \mathcal{X}_2) = 0$, i.e. $f' \in F$. Hence any possible optimal left bound function $f_{L,opt}(x) = \arg \min_f R(f)$ belongs to F . In addition, according to the assumptions (C1) and (C2) of partial monotonicity for PDI, we have that $P(Y(f_{L,opt}(x)) > S) = \alpha$ and $P(Y(a'(x)) > S) \geq \alpha$ for $\forall x \in \mathcal{X}$, where $a'(x)$ is an arbitrary measurable function s.t. $f_{L,opt}(x) \leq a'(x) \leq a_U$. Theorem 4.1 is proved.

□

B.2 Proof of Theorem 4.2

Proof. For any measurable function $f : \mathcal{X} \rightarrow \mathbb{R}$, we have

$$\begin{aligned}
 & |R(f) - R_{\epsilon_n}(f)| \\
 &= E \left| \frac{1}{\epsilon_n P(A|X)} \left((1-\alpha) I(Y > S) \max(\epsilon_n - (f(X) - A)_+, 0) \right. \right. \\
 &\quad \left. \left. + \alpha I(Y \leq S) \max(\epsilon_n - (A - f(X))_+, 0) \right) \right| \\
 &= E_X \left| \frac{1}{\epsilon_n} \int_{a \in [a_L, a_U]} \left((1-\alpha) P(Y > S|A = a, X) \max(\epsilon_n - (f(X) - a)_+, 0) \right. \right. \\
 &\quad \left. \left. + \alpha P(Y \leq S|A = a, X) \max(\epsilon_n - (a - f(X))_+, 0) \right) da \right| \\
 &= E_X \left| \frac{(1-\alpha)}{\epsilon_n} \int_{a \in [f(X) - \epsilon_n, f(X)]} P(Y > S|A = a, X) (\epsilon_n - (f(X) - a)) da \right. \\
 &\quad \left. + \frac{\alpha}{\epsilon_n} \int_{a \in [f(X), f(X) + \epsilon_n]} P(Y \leq S|A = a, X) (\epsilon_n - (a - f(X))) da \right| \\
 &= E_X \left| \frac{(1-\alpha)}{\epsilon_n} \int_{z \in [0, \epsilon_n]} P(Y > S|A = f(X) - z, X) (\epsilon_n - z) dz \right. \\
 &\quad \left. + \frac{\alpha}{\epsilon_n} \int_{z \in [0, \epsilon_n]} P(Y \leq S|A = f(X) + z, X) (\epsilon_n - z) dz \right| \\
 &\leq \epsilon_n
 \end{aligned}$$

Therefore Theorem 4.2 follows. \square

B.3 Proof of Theorem 4.3

Proof. According to Theorem 4.2, we have

$$\begin{aligned}
 R(\hat{f}_L) - R(f_{opt}) &\leq R_{\epsilon_n}(\hat{f}_L) - R_{\epsilon_n}(f_{L,opt}) + 2C\epsilon_n \\
 &\leq R_{\epsilon_n}(\hat{f}_L) - R_{\epsilon_n}(f_L^*) + 2C\epsilon_n \\
 &\leq \left(\lambda_n \|\hat{f}_L\|_k^2 + R_{\epsilon_n}(\hat{f}_L) - R_{\epsilon_n}^* \right) + \left(2C\epsilon_n \right) \\
 &= (I) + (II)
 \end{aligned} \tag{13}$$

where f_L^* is the minimizer of $R_{\epsilon_n}(\cdot)$ and $R_{\epsilon_n}^* = R_{\epsilon_n}(f_L^*)$.

To bound (13), we need to first bound (I). As in Chen et al. (2016), we refer to Theorem 7.23 in Steinwart and Christmann (2008). In order to use the oracle inequality in the theorem, there are four conditions to be satisfied.

- (B1) The loss function $L(\cdot)$ has a supremum bound $L(\cdot) \leq B$ for a constant $B > 0$.
- (B2) The loss function $L(\cdot)$ is locally Lipschitz continuous and can be clipped at a constant $M > 0$ such that $f = I(|f| \leq M)f + I(|f| \geq M)M$.
- (B3) The variance bound $E_P(L \circ \tilde{f} - L \circ f_{L,P}^*)^2 \leq V \left(E_P(L \circ \tilde{f} - L \circ f_{L,P}^*) \right)^v$ is satisfied for a constant $v \in [0, 1]$, $V \geq B^{2-v}$ and all $f \in \mathcal{H}$.
- (B4) For fixed $n \geq 1$, there exist constants $p \in (0, 1)$ and $a \geq B$ such that the entropy number $E_{D_X \sim P_X^n} e_i(id : \mathcal{H} \rightarrow L_2(D_X)) \leq ai^{-\frac{1}{2p}}$, $i \leq 1$.

To verify Condition (B1), recall that the relaxed loss function of the lower bound of the one-sided PDI is

$$L_\epsilon(X, A, f_L(X)) = \frac{1}{P(A|X)} \left\{ (1 - \alpha) I(Y > S) \Psi_\epsilon(f_L(X), A) + \alpha I(Y \leq S) \Psi_\epsilon(A, f_L(X)) \right\}$$

Assuming the inverse probability $1/P(A|X)$ is bounded by a constant B , then $L_\epsilon(X, A, f_L(X))$ is naturally bounded by B as well.

To verify Condition (B2), similar to (B1), $L_\epsilon(X, A, f_L(X))$ is Lipschitz continuous with a Lipschitz constant B . Also, f_L can be clipped to have a smaller risk. Suppose that M is an upper bound of the absolute value of a reasonable range of dose, say, $M = \max\{|a_L|, |a_U|\}$, and $\tilde{f}_L = I(|f_L| \leq M) f_L + I(|f_L| \geq M) M$. It naturally follows that $R(\tilde{f}_L) \leq R(f_L)$, since any unreasonably large dose recommendation introduces a larger risk.

Condition (B3) is satisfied with $v = 0$ and $V = 4B^2$, because

$$E[(L_\epsilon \circ \tilde{f}_L - L_\epsilon \circ f_L^*)^2] \leq 2E[(L_\epsilon \circ \tilde{f}_L)^2 + (L_\epsilon \circ f_L^*)^2] \leq 4B^2$$

The Gaussian Kernel used in Sections 5.1 and 5.2 is one type of benign kernels. According to Theorem 7.34 of Steinwart and Christmann (2008), (B4) is satisfied with the constant $a = c_{\epsilon,p} \gamma_n^{-\frac{(1-p)(1+\zeta)d}{2p}}$, where $d/(d+\tau) < p < 1$ and $\zeta > 0$ are two constants.

Since Conditions (B1)-(B4) are satisfied in our case, applying Theorem 7.23 from Steinwart and Christmann (2008) yields

$$(I) \leq 9 * \left\{ \lambda_n \|f_L^0\|_k^2 + R_{\epsilon_n}(f_L^0) - R_{\epsilon_n}^* \right\} + K_0 \left[\frac{1}{\gamma_n^{(1-p)(1+\zeta)d} \lambda_n^p n} \right]^{\frac{1}{2-p}} + 36\sqrt{2}B \left(\frac{\tau}{n}\right)^{\frac{1}{2}} + \frac{15B\tau}{n} \quad (14)$$

with probability $1 - 3e^{-\tau}$.

In order to bound $(\lambda_n \|f_L^0\|_k^2 + R_{\epsilon_n}(f_L^0) - R_{\epsilon_n}^*)$, the approximation error, we refer to Section 2 in Eberts et al. (2013). As (14) holds for any $f_L^0 \in \mathcal{H}_\gamma$, we construct a specific f_L^0 to facilitate the proof. First, similar to Equation (8) in Eberts et al. (2013), we define a function $\mathcal{K}(x) = \sum_{j=1}^r \binom{r}{j} (-1)^{1-j} \frac{1}{j^d} \left(\frac{2}{\gamma^2}\right)^{d/2} \mathcal{K}_{j\gamma/\sqrt{2}}(x)$ where $\mathcal{K}_\gamma = \exp(-\gamma^2 \|x\|_2^2)$ for all $x \in R^d$, and subsequently define f_L^0 via convolution as follows,

$$f_L^0 = \mathcal{K} * f_L^* = \int_{R^d} \mathcal{K}(X - t) f_L^*(t) dt, \quad x \in R^d$$

If we assume $f_L^* \in L_2(R^d) \cap L_\infty(R^d)$, then from Theorem 2.3 in Eberts et al. (2013), it can be shown that $f_L^0 \in \mathcal{H}_\gamma$, where \mathcal{H}_γ is the RKHS with the Gaussian kernel. In addition, Theorem 2.2 provides the upper bound for the access risk of f_L^0 , which can be incorporated to yield the following results.

$$\begin{aligned} & \lambda_n \|f_L^0\|_k^2 + R_{\epsilon_n}(f_L^0) - R_{\epsilon_n}^* \\ &= \lambda_n \|\mathcal{K} * f_L^*\|_k^2 + R_{\epsilon_n}(\mathcal{K} * f_L^*) - R_{\epsilon_n}^* \\ &\leq \lambda_n (\gamma_n \sqrt{\pi})^{-d} (2^r - 1)^2 \|f_L^*\|_{L_2(R^d)}^2 + R_{\epsilon_n}(\mathcal{K} * f_L^*) - R_{\epsilon_n}^* \\ &\leq \lambda_n (\gamma_n \sqrt{\pi})^{-d} (2^r - 1)^2 \|f_L^*\|_{L_2(R^d)}^2 + B \|\mathcal{K} * f_L^* - f_L^*\|_{L_1(P_X)} \\ &\leq \lambda_n (\gamma_n \sqrt{\pi})^{-d} (2^r - 1)^2 \|f_L^*\|_{L_2(R^d)}^2 + BC_{r,1} \|g\|_{L_p(P_X)} \omega_{r,L_1(R^d)}(f_L^*, \gamma_n/2) \end{aligned} \quad (15)$$

Given the fact that $f_{opt} \in B_{1,\infty}^\delta(R^d)$, if we further assume $f_L^* \in B_{1,\infty}^\delta(R^d)$, i.e., $B_{1,\infty}^\delta(R^d) = \{f \in L_\infty(R^d) : \sup_{t>0} (t^{-\delta} \omega_{r,L_1(R^d)})(f, t) < \infty\}$, then $w_{r,L_1}(R^d)(f_L^*, \gamma_n/2) < c_0 \gamma_n^\delta$, where c_0 is a constant. Plugging this into (15), we obtain

$$\lambda_n \|f_L^0\|_k^2 + R_{\epsilon_n}(f_L^0) - R_{\epsilon_n}^* \leq c_1 \lambda_n \gamma_n^{-d} + c_2 \gamma_n^\delta$$

After combining all the parts together,

$$R(\hat{f}_L) - R(f_{opt}) \leq c_1 \frac{\lambda_n}{\gamma_n^d} + c_2 \gamma_n^\delta + c_3 \frac{1}{\gamma_n^{\frac{(1-p)(1+\zeta)d}{2-p}} \lambda_n^{\frac{p}{2-p}} n^{\frac{1}{2-p}}} + c_4 \frac{\tau^{\frac{1}{2}}}{n^{\frac{1}{2}}} + c_5 \frac{\tau}{n} + c_6 \epsilon_n$$

By properly choosing the constant, i.e.

$$\gamma_n \propto \left(\frac{1}{n}\right)^{\frac{1}{2\delta+d}} \quad \lambda_n \propto \left(\frac{1}{n}\right)^{\frac{\delta+d}{2\delta+d}} \quad \epsilon_n = \mathcal{O}\left(n^{-\frac{\delta}{2\delta+d}}\right)$$

we have the convergence rate:

$$R(\hat{f}_L) - R(f_{opt}) = \mathcal{O}\left(n^{-\frac{\delta}{2\delta+d}}\right)$$

□

C THEORETICAL RESULTS FOR TWO-SIDED PDI

To ensure the two-sided PDI exists, we need the following assumptions similar to the one-sided PDI case:

(C3) For any x , the boundaries a_L and a_U for the dose range satisfy $P\{Y(a_L) > S|x\} \leq \alpha$ and $P\{Y(a_U) > S|x\} \leq \alpha$.

(C4) There exists a function $a^*(\cdot) \in [a_L, a_U]$ such that $P\{Y(a^*(x)) > S|x\} \geq \alpha$

(C5) For any x , there are two numbers $a_{1,x} < a_{2,x}$ with $a_{1,x} \in [a_L, a_U]$ and $a_{2,x} \in [a_L, a_U]$ such that $P\{Y(a_{1,x}) > S|x\} = \alpha$ and $P\{Y(a_{2,x}) > S|x\} = \alpha$.

Corollary C.0.1. *Let $[f_{L,opt}, f_{U,opt}] = \arg \min_{f_L \leq f_U} R(f_L, f_U)$. Then, under the assumptions (C3)-(C5), $[f_{L,opt}, f_{U,opt}]$ has the following properties:*

(1) $P\{Y(f_{L,opt}(x)) > S|x\} = \alpha$ and $P\{Y(f_{U,opt}(x)) > S|x\} = \alpha, \forall x$.

(2) For any measurable function $a(\cdot)$ s.t. $f_{L,opt}(x) \leq a(x) \leq f_{U,opt}(x), \forall x \in \mathcal{X}$, the potential outcome $Y(a(x))$ satisfies $P(Y(a(x)) > S|x) \geq \alpha$.

(3) For any $a_1(x) < f_{L,opt}(x), P(Y(a_1(x)) > S|x) < \alpha$, and for any $a_2(x) > f_{U,opt}(x), P(Y(a_2(x)) > S|x) < \alpha$.

Proof. First,

$$F_L \equiv \{f : \forall x \in \mathcal{X}, P(Y(f(x)) > S|x) = \alpha \quad \text{and} \quad \exists a^* < f(x), P(Y(a^*(x)) > S|x) \leq \alpha\}$$

$$F_U \equiv \{f : \forall x \in \mathcal{X}, P(Y(f(x)) > S|x) = \alpha \quad \text{and} \quad \exists a^* > f(x), P(Y(a^*(x)) > S|x) \leq \alpha\}$$

Let $f_L^* \in F_L$ and $f_U^* \in F_U$, then for arbitrary f'_L and f'_U , the sample space of \mathcal{X} can be decomposed as $\mathcal{X} = \mathcal{X}_0 \cup \mathcal{X}_1 \cup \mathcal{X}_2 \cup \mathcal{X}_3 \cup \mathcal{X}_4 \cup \mathcal{X}_5 \cup \mathcal{X}_6$ where

$$\begin{aligned} \mathcal{X}_0 &= \{x : x \in \mathcal{X}, \quad \exists f_L \in F_L \text{ s.t. } f'_L(x) = f_L(x) \quad \text{and} \quad \exists f_U \in F_U \text{ s.t. } f'_U(x) = f_U(x)\} \\ \mathcal{X}_1 &= \{x : x \in \mathcal{X}, \quad \max_{f_L \in F_L} f_L(x) < f'_L(x) \leq f'_U(x) < \min_{f_U \in F_U} f_U(x)\} \\ \mathcal{X}_2 &= \{x : x \in \mathcal{X}, \quad f'_L(x) < \min_{f_L \in F_L} f_L(x) \leq \max_{f_U \in F_U} f_U(x) < f'_U(x)\} \\ \mathcal{X}_3 &= \{x : x \in \mathcal{X}, \quad f'_L(x) \leq \min_{f_L \in F_L} f_L(x) \leq f'_U(x) \leq \max_{f_U \in F_U} f_U(x)\} \\ \mathcal{X}_4 &= \{x : x \in \mathcal{X}, \quad \min_{f_L \in F_L} f_L(x) \leq f'_L(x) \leq \max_{f_U \in F_U} f_U(x) \leq f'_U(x)\} \\ \mathcal{X}_5 &= \{x : x \in \mathcal{X}, \quad f'_L(x) \leq f'_U(x) < \min_{f_L \in F_L} f_L(x) \leq \max_{f_U \in F_U} f_U(x)\} \\ \mathcal{X}_6 &= \{x : x \in \mathcal{X}, \quad \min_{f_L \in F_L} f_L(x) \leq \max_{f_U \in F_U} f_U(x) < f'_L(x) \leq f'_U(x)\} \end{aligned}$$

Then

$$\begin{aligned}
 R(f') - R(f^*) &= E \left[\frac{1}{P(A|X)} \left((1-\alpha)I(Y > S) \left(I(A \notin [f'_L(X), f'_U(X)]) - I(A \notin [f_L^*(X), f_U^*(X)]) \right) \right. \right. \\
 &\quad \left. \left. + \alpha I(Y \leq S) \left(I(A \in [f'_L(X), f'_U(X)]) - I(A \in [f_L^*(X), f_U^*(X)]) \right) \right) \right] \\
 &= E [\dots | X \in \mathcal{X}_1] P(X \in \mathcal{X}_0) & (M0) \\
 &\quad + E [\dots | X \in \mathcal{X}_1] P(X \in \mathcal{X}_1) & (M1) \\
 &\quad + E [\dots | X \in \mathcal{X}_2] P(X \in \mathcal{X}_2) & (M2) \\
 &\quad + E [\dots | X \in \mathcal{X}_3] P(X \in \mathcal{X}_3) & (M3) \\
 &\quad + E [\dots | X \in \mathcal{X}_4] P(X \in \mathcal{X}_4) & (M4) \\
 &\quad + E [\dots | X \in \mathcal{X}_5] P(X \in \mathcal{X}_5) & (M5) \\
 &\quad + E [\dots | X \in \mathcal{X}_6] P(X \in \mathcal{X}_6) & (M6)
 \end{aligned}$$

where $(M3) \geq 0$, $(M3) > 0$ if $P(X \in \mathcal{X}_3) > 0$, and $(M4) \geq 0$, $(M4) > 0$ if $P(X \in \mathcal{X}_4) > 0$, according to the results of Theorem 4.1. These results can be directly applied since it is easily seen that within \mathcal{X}_3 and \mathcal{X}_4 , assumptions for Theorem 4.1 are satisfied. Obviously $(M0) = 0$ and

$$\begin{aligned}
 (M1) &= E \left[\int_{f_L^*(X)}^{f'_L(X)} \left((1-\alpha)P(Y > S|A=a, X) - \alpha P(Y \leq S|A=a, X) \right) da \right. \\
 &\quad \left. + \int_{f'_U(X)}^{f_U^*(X)} \left((1-\alpha)P(Y > S|A=a, X) - \alpha P(Y \leq S|A=a, X) \right) da \middle| X \in \mathcal{X}_1 \right] P(X \in \mathcal{X}_1) \\
 &= E \left[\int_{\max_{f_L \in F_L} f_L(X)}^{f'_L(X)} \left(P(Y > S|A=a, X) - \alpha \right) da \right. \\
 &\quad \left. + \int_{f'_U(X)}^{\min_{f_U \in F_U} f_U(X)} \left(P(Y > S|A=a, X) - \alpha \right) da \middle| X \in \mathcal{X}_1 \right] P(X \in \mathcal{X}_1)
 \end{aligned}$$

Therefore $(M1) \geq 0$ and > 0 if $P(X \in \mathcal{X}_1) \neq 0$. Similarly,

$$\begin{aligned}
 (M2) &= E \left[\int_{f'_L(X)}^{f_L^*(X)} \left(-(1-\alpha)P(Y > S|A=a, X) + \alpha P(Y \leq S|A=a, X) \right) da \right. \\
 &\quad \left. + \int_{f_U^*(X)}^{f'_U(X)} \left(-(1-\alpha)P(Y > S|A=a, X) + \alpha P(Y \leq S|A=a, X) \right) da \middle| X \in \mathcal{X}_2 \right] P(X \in \mathcal{X}_2) \\
 &= E \left[\int_{f'_L(X)}^{\min_{f_L \in F_L} f_L(X)} \left(-P(Y > S|A=a, X) + \alpha \right) da \right. \\
 &\quad \left. + \int_{\max_{f_U \in F_U} f_U(X)}^{f'_U(X)} \left(-P(Y > S|A=a, X) + \alpha \right) da \middle| X \in \mathcal{X}_2 \right] P(X \in \mathcal{X}_2)
 \end{aligned}$$

Therefore $(M2) \geq 0$ and > 0 if $P(X \in \mathcal{X}_2) \neq 0$.

$$\begin{aligned}
 (M5) &= E \left[\int_{f'_L(X)}^{f'_U(X)} \left(-(1-\alpha)P(Y > S|A=a, X) + \alpha P(Y \leq S|A=a, X) \right) da \right. \\
 &\quad \left. + \int_{f'_U(X)}^{f_L^*(X)} \left((1-\alpha)P(Y > S|A=a, X) - \alpha P(Y \leq S|A=a, X) \right) da \middle| X \in \mathcal{X}_5 \right] P(X \in \mathcal{X}_5) \\
 &= E \left[\int_{f'_L(X)}^{f'_U(X)} \left(-P(Y > S|A=a, X) + \alpha \right) da \right. \\
 &\quad \left. + \int_{f'_U(X)}^{f_L^*(X)} \left(P(Y > S|A=a, X) - \alpha \right) da \middle| X \in \mathcal{X}_5 \right] P(X \in \mathcal{X}_5)
 \end{aligned}$$

Therefore $(M5) \geq 0$ and > 0 if $P(X \in \mathcal{X}_2) \neq 0$.

$$\begin{aligned}
 (M6) &= E \left[\int_{f'_L(X)}^{f'_U(X)} \left(- (1 - \alpha)P(Y > S|A = a, X) + \alpha P(Y \leq S|A = a, X) \right) da \right. \\
 &\quad \left. + \int_{f_U^*(X)}^{f_L^*(X)} \left((1 - \alpha)P(Y > S|A = a, X) \right. \right. \\
 &\quad \quad \left. \left. - \alpha P(Y \leq S|A = a, X) \right) da \middle| X \in \mathcal{X}_6 \right] P(X \in \mathcal{X}_6) \\
 &= E \left[\int_{f'_L(X)}^{f'_U(X)} \left(- P(Y > S|A = a, X) + \alpha \right) da \right. \\
 &\quad \left. + \int_{f_U^*(X)}^{f_L^*(X)} \left(P(Y > S|A = a, X) - \alpha \right) da \middle| X \in \mathcal{X}_6 \right] P(X \in \mathcal{X}_6)
 \end{aligned}$$

Therefore $(M6) \geq 0$ and > 0 if $P(X \in \mathcal{X}_6) \neq 0$. Consequently, $R(f'_U, f'_U) - R(f_L^*, f_U^*) \geq 0$, and $R(f'_U, f'_U) - R(f_L^*, f_U^*) = 0$ if and only if $P(X \notin \mathcal{X}_0) = 0$. Hence for any optimal two-sided bound functions $[f_{L,opt}, f_{U,opt}] = \arg \min_{f_L \leq f_U} R(f_L, f_U)$, we have $f_{L,opt} \in F_L$ and $f_{U,opt} \in F_U$. According to the assumptions (C3), (C4) and (C5), we have that $P(Y(f_{L,opt}(x)) > S|x) = P(Y(f_{U,opt}(x)) > S|x) = \alpha$ and $P(Y(a'_x) > S|x) \geq \alpha$ for $\forall x \in \mathcal{X}$, where $a'(x)$ is an arbitrary measurable function s.t. $f_{L,opt}(x) \leq a'(x) \leq f_{U,opt}(x)$. \square

Corollary C.0.2. *For any interval of two measurable functions $[f_L(x), f_U(x)]$, where $f_L(x) \leq f_U(x), \forall x \in \mathcal{X}$, we have $|R(f_L, f_U) - R_{\epsilon_n}(f_L, f_U)| \leq C\epsilon_n$, where C is a constant.*

Proof. For any measurable function $f_L : \mathcal{X} \rightarrow \mathbb{R}$, $f_U : \mathcal{X} \rightarrow \mathbb{R}$ and $f_L(X) \leq f_U(X) \forall x \in \mathcal{X}$, we have

$$\begin{aligned}
 &|R(f) - R_{\epsilon_n}(f)| \\
 &= E_X \left| \frac{(1 - \alpha)}{\epsilon_n} \int_{a \in [f_L(X) - \epsilon_n, f_L(X)]} P(Y > S|A = a, X)(\epsilon_n - (f_L(X) - a)) da \right. \\
 &\quad \left. + \frac{\alpha}{\epsilon_n} \int_{a \in [f_L(X), f_L(X) + \epsilon_n]} P(Y \leq S|A = a, X)(\epsilon_n - (a - f_L(X))) da \right. \\
 &\quad \left. + \frac{\alpha}{\epsilon_n} \int_{a \in [f_U(X) - \epsilon_n, f_U(X)]} P(Y \leq S|A = a, X)(\epsilon_n - (f_U(X) - a)) da \right. \\
 &\quad \left. + \frac{(1 - \alpha)}{\epsilon_n} \int_{a \in [f_U(X), f_U(X) + \epsilon_n]} P(Y > S|A = a, X)(\epsilon_n - (a - f_U(X))) da \right| \\
 &= E_X \left| \frac{(1 - \alpha)}{\epsilon_n} \int_{z \in [0, \epsilon_n]} P(Y > S|A = f_L(X) - z, X)(\epsilon_n - z) dz \right. \\
 &\quad \left. + \frac{\alpha}{\epsilon_n} \int_{z \in [0, \epsilon_n]} P(Y \leq S|A = f_L(X) + z, X)(\epsilon_n - z) dz \right. \\
 &\quad \left. + \frac{\alpha}{\epsilon_n} \int_{z \in [0, \epsilon_n]} P(Y \leq S|A = f_U(X) - z, X)(\epsilon_n - z) dz \right. \\
 &\quad \left. + \frac{(1 - \alpha)}{\epsilon_n} \int_{z \in [0, \epsilon_n]} P(Y > S|A = f_U(X) + z, X)(\epsilon_n - z) dz \right| \\
 &= 2\epsilon_n.
 \end{aligned}$$

\square

We provide a convergence rate by borrowing the results from Theorem 4.3, and replace the convergence of the two-sided PDI estimator with two estimators of one-sided PDIs. In order to do so, extra assumptions have to be made. It would be an interesting future work if better rates under weaker assumptions can be achieved.

Corollary C.0.3. Assume that $f_{L,opt} \in B_{1,\infty}^\delta(\mathbb{R}^d)$ and $f_{U,opt} \in B_{1,\infty}^\delta(\mathbb{R}^d)$ where $B_{1,\infty}^\delta(\mathbb{R}^d) = \{f \in L_\infty(\mathbb{R}^d) : \sup_{t>0} (t^{-\delta} w_{(r,L_1)}(f,t)) < \infty\}$ and w is the modulus of continuity. In addition, assume that there exists a measurable function $f_{M,opt}(x)$, which is known or can be estimated consistently, such that $f_{L,opt}(x) < f_{M,opt}(x) < f_{U,opt}(x)$ for all $x \in \mathcal{X}$. Then for any $\eta > 0$, $d/(d+\tau) < p < 1$, $\tau > 0$, $0 < c < 1$, and properly chosen γ_n, λ_n and ϵ_n , we have

$$R(\hat{f}_{L,n}, \hat{f}_{U,n}) - R(f_{L,opt}, f_{U,opt}) = \mathcal{O}_p\left(n^{-\frac{\delta}{2\delta+d}}\right) \quad (16)$$

with probability $c(1 - 3e^{-\tau})^2$. Here d is the dimension of \mathcal{X} .

Proof. The proof of Corollary C.0.3 consists of two steps. In the first step, we assume that there exists a function or a consistent estimator of a function that separates the dataset into two subsets which we call as left and right subsets, in each of which Conditions (C1) and (C2) are satisfied. In the second step, the integrated excessive risk of the two-sided PDI estimator is bounded by the sum of excessive risks of two independent one-sided PDI estimators.

Suppose there is a function $f_{M,opt}(x)$, such that $f_{L,opt}(x) < f_{M,opt}(x) < f_{U,opt}(x)$ for all $x \in \mathcal{X}$. Such a function is typically unknown in practice, hence we consider the case when $f_{M,opt}(x)$ is estimated from data. One choice of $f_{M,opt}$ is $f_{\max,opt} = \arg \max_f E[Y|A = f(X), X]$, the optimal individual dose rule (IDR) in Chen et al. (2016). Theoretical properties of the estimator \hat{f}_{\max} have been established in Chen et al. (2016) and also discussed in Luedtke and van der Laan (2016). Despite the fact that it is suggested in Luedtke and van der Laan (2016) that the O-learning based estimator might convergence in nearly $\mathcal{O}(n^{-1/2})$, the rate established there was based on the risk (or value function) instead of the rule $f_{\max,opt}$ itself. To facilitate proving the properties of the two-sided PDI, we make the following further assumption for the estimated optimal dose function \hat{f}_{\max} .

$$f_{L,opt}(x) < \hat{f}_{\max}(x) < f_{U,opt}(x), \quad \forall x \in \mathcal{X} \quad (17)$$

In the second step, two subsets are defined where the lower bound function and the upper bound function can be estimated separately using the corresponding subset, assuming there is a $f_{M,opt}$, known or estimated, such that $f_{L,opt}(x) < f_{M,opt}(x) < f_{U,opt}(x)$ for all $x \in \mathcal{X}$. Denote the original dataset as $\mathcal{D} = \{(X_i, A_i, Y_i), 1 \leq i \leq n\}$ and then we can subsequently define the left subset as $\mathcal{D}_{\mathcal{L}} = \{(X_i, A_i, Y_i), \forall i \text{ s.t. } A_i \leq f_{M,opt}(X_i)\}$ and the right subset as $\mathcal{D}_{\mathcal{U}} = \{(X_i, A_i, Y_i), \forall i \text{ s.t. } A_i > f_{M,opt}(X_i)\}$. These two subsets are independent as there is no shared observations.

Now define $\hat{f}'_L = \min_{f_L} \hat{R}_\epsilon(f_L)$ on $\mathcal{D}_{\mathcal{L}}$ and $\hat{f}'_U = \min_{f_U} \hat{R}_\epsilon(f_U)$ on $\mathcal{D}_{\mathcal{U}}$, while $\{\hat{f}_L, \hat{f}_U\} = \min_{f_L, f_U} \hat{R}_\epsilon(f_L, f_U)$ on the dataset of \mathcal{D} . Similarly, define $f'_{L,*} = \min_{f_L} R_\epsilon(f_L)$ on the (sub)population of $\mathcal{D}_{\mathcal{L}}$ and $f'_{U,*} = \min_{f_U} R_\epsilon(f_U)$ on the (sub)population of $\mathcal{D}_{\mathcal{U}}$, while $\{f_{L,*}, f_{U,*}\} = \min_{f_L, f_U} R_\epsilon(f_L, f_U)$ on the original population. To simplify the problem, we assume that $f_{L,*}(x) - f_{U,*}(x) \leq -2\epsilon$, since this inequality will hold when $\epsilon \rightarrow 0$ as $n \rightarrow 0$. By the definitions of $\mathcal{D}_{\mathcal{L}}$, $\mathcal{D}_{\mathcal{U}}$ and \mathcal{D} , we have $\{f'_{L,*}, f'_{U,*}\} = \{f_{L,*}, f_{U,*}\}$, i.e., the population estimates of the lower and upper bounds estimated jointly are the same as the population estimates of the lower and upper bounds estimated individually while assuming the other part is known. In addition, we make the following assumption of the uniform convergence of the empirical risk function.

$$\sup_{f_L, f_U \in \mathcal{H}} |\hat{R}_\epsilon(f_L, f_U) - R_\epsilon(f_L, f_U)| \leq \kappa \sqrt{\frac{\log(n)}{n}} \text{ with probability } c \quad (18)$$

where κ is a constant related to c and the complexity of the RKHS.

We will show how the loss function of the two-sided PDI can be decomposed as the sum of two one-sided PDI losses plus a quantity not related to the dose. Recall that the loss function for the lower bound of the one-sided PDI is

$$L_\epsilon(X, A, f_L(X)) = \frac{1}{P(A|X)} \left\{ (1 - \alpha) I(Y > S) \Psi_\epsilon(f_L(X), A) + \alpha I(Y \leq S) \Psi_\epsilon(A, f_L(X)) \right\}$$

The loss function for the upper bound of the one-sided PDI is

$$L_\epsilon(X, A, f_U(X)) = \frac{1}{P(A|X)} \left\{ (1 - \alpha) I(Y > S) \Psi_\epsilon(A, f_U(X)) + \alpha I(Y \leq S) \Psi_\epsilon(f_U(X), A) \right\}.$$

And the loss function for the two-sided PDI is

$$L_\epsilon(X, A, f_L(X), f_U(X)) = \frac{1}{P(A|X)} \left\{ (1 - \alpha)I(Y > S)\Psi_\epsilon^{out}(f_L(X), A, f_U(X)) + \alpha I(Y \leq S)\Psi_\epsilon^{in}(f_L(X), A, f_U(X)) \right\}.$$

It follows that the difference between the two-sided loss and the sum of the one-sided losses is a constant which does not contain the bound functions, under the assumption that $f_{L,*}(x) - f_{U,*}(x) \leq -2\epsilon$.

$$\begin{aligned} & L_\epsilon(X, A, f_L(X), f_U(X)) - L_\epsilon(X, A, f_L(X)) - L_\epsilon(X, A, f_U(X)) \\ &= \frac{1}{P(A|X)} \left\{ (1 - \alpha)I(Y > S)\Psi_\epsilon(A, f_U(X)) + \alpha I(Y \leq S)(\Psi_\epsilon(f_U(X), A) - 1) \right\} - L_\epsilon(X, A, f_U(X)) \\ &= \frac{-1}{P(A|X)} \left\{ \alpha I(Y \leq S) \right\} \end{aligned}$$

Combined with Corollary C.0.2, we have

$$\begin{aligned} R(\hat{f}_L, \hat{f}_U) - R(f_{L,opt}, f_{U,opt}) &\leq R_{\epsilon_n}(\hat{f}_L, \hat{f}_U) - R_{\epsilon_n}(f_{L,opt}, f_{U,opt}) + 2C\epsilon_n \\ &\leq R_{\epsilon_n}(\hat{f}_L, \hat{f}_U) - R_{\epsilon_n}(f_{L,*}, f_{U,*}) + 2C\epsilon_n \\ &\leq R_{\epsilon_n}(\hat{f}'_L, \hat{f}'_U) - R_{\epsilon_n}(f_{L,*}, f_{U,*}) + 2C\epsilon_n + 2\kappa\sqrt{\frac{\log(n)}{n}} \\ &= R_{\epsilon_n}(\hat{f}'_L) - R_{\epsilon_n}(f'_{L,*}) + R_{\epsilon_n}(\hat{f}'_U) - R_{\epsilon_n}(f'_{U,*}) + 2C\epsilon_n + 2\kappa\sqrt{\frac{\log(n)}{n}} \end{aligned} \quad (19)$$

with probability c as $n \rightarrow \infty$. The probability c comes from the inequality (18). The first inequality of (19) is guaranteed by Theorem C.0.1. The second inequality is from the definition of $f_{L,*}$ and $f_{U,*}$. The third inequality is from the assumption (18). The coefficient is $2 \times \kappa$ as the uniform bound has to be applied twice because the definitions of \hat{f}_L and \hat{f}_U only guarantee $\hat{R}_{\epsilon_n}(\hat{f}_L, \hat{f}_U) < \hat{R}_{\epsilon_n}(\hat{f}'_L, \hat{f}'_U)$ but not $R_{\epsilon_n}(\hat{f}_L, \hat{f}_U) < R_{\epsilon_n}(\hat{f}'_L, \hat{f}'_U)$. The last equality is from the decomposition of the two-sided loss and the equivalence of $\{f'_{L,*}, f'_{U,*}\}$ and $\{f_{L,*}, f_{U,*}\}$.

The theorem follows by applying (14) twice, because \hat{f}'_L and \hat{f}'_U are independently estimated on two non-overlapping datasets \mathcal{D}_L and \mathcal{D}_U . That is, $R_{\epsilon_n}(\hat{f}'_L) - R_{\epsilon_n}(f'_{L,*})$ is independent with $R_{\epsilon_n}(\hat{f}'_U) - R_{\epsilon_n}(f'_{U,*})$. The uniform bound (18) and Theorem 4.3 hold simultaneously with a probability of at least $c(1 - 3e^{-\tau})^2$. Here ϵ_n only shows up in the term $2C\epsilon_n$. Hence by choosing $\epsilon_n = \mathcal{O}(n^{-\frac{\delta}{2\delta+d}})$, it will not impact the convergence rate. The term $2\kappa\sqrt{\frac{\log(n)}{n}}$ corresponds to the convergence rate of $\mathcal{O}(\sqrt{\frac{\log(n)}{n}})$. Hence, (19) is dominated either by the first two risk terms or by $2\kappa\sqrt{\frac{\log(n)}{n}}$. \square

D ADDITIONAL SIMULATIONS WITH NO CONFOUNDING COVARIATES

In this section, we consider scenarios similar to those in the main paper except that there is no confounding. In particular, the treatment A of the training data follows a uniform distribution $\text{Unif}(-2, 2)$ (i.e., independent of X) while other data generating mechanisms are the same as in the main paper. We repeat each setting for 100 times and report results by evaluating the empirical risks on corresponding independent testing data sets with sample sizes of 10000. Similar to the settings with confounding, our methods perform better than the competitive methods.

E MORE BACKGROUND ABOUT THE HBA1C CONTROL STUDY

Hemoglobin is an iron-containing oxygen which transports protein in the red blood cells. HbA1c, the most abundant minor hemoglobin in the human body, is formed when glucose accumulates in red blood cells and binds to the hemoglobin. This process occurs slowly and continuously over the lifespan of red blood cells, which is

Table 2: Empirical risk $R(\hat{f}_L)$ (with SD) under settings without confounding covariates

	n	d	L-O-Learning	G-O-Learning	Logistic	SVM	RF
Scenario 1 $\sigma^2 = 9$	200	10	0.137 (0.008)	0.132 (0.011)	0.154 (0.024)	0.139 (0.011)	0.136 (0.008)
	200	50	0.157 (0.014)	0.133(0.011)	0.185(0.023)	0.181 (0.043)	0.135 (0.012)
	400	10	0.128 (0.006)	0.126 (0.006)	0.132(0.011)	0.129 (0.006)	0.130 (0.009)
	400	50	0.154 (0.006)	0.128 (0.007)	0.174 (0.022)	0.142 (0.008)	0.129 (0.008)
Scenario 2 $\sigma^2 = 9$	200	10	0.136 (0.008)	0.130 (0.011)	0.153 (0.022)	0.138 (0.011)	0.133 (0.012)
	200	50	0.159 (0.013)	0.134 (0.010)	0.194 (0.024)	0.226 (0.037)	0.139 (0.017)
	400	10	0.129 (0.005)	0.123 (0.006)	0.132 (0.012)	0.127 (0.005)	0.124 (0.005)
	400	50	0.156 (0.007)	0.126 (0.008)	0.180 (0.026)	0.144 (0.008)	0.127 (0.004)

120 days on average. This makes HbA1c an ideal biomarker of long-term glycemic control. Patients who are susceptible to high blood glucose are usually recommended to have their HbA1c levels regularly measured and recorded.

HbA1c test, since became commercially available in 1978, has been one of the standard tools for monitoring diabetes progression. American Diabetes Association (ADA) recommended using A1c measurement in 1988. The Diabetes Control and Complications Trial (DCCT) demonstrated its importance as a predictor of diabetes-related outcomes, and the ADA started recommending specific A1c targets in 1994 (Little et al., 2011). In 2010, ADA added $\text{HbA1c} \geq 6.5\%$ as a criterion for diabetes diagnosis. The prevalence of HbA1c test can be partially explained by its clinic convenience. As HbA1c is unaffected by acute perturbations in glucose levels, there is no need for fasting or timed samples. For people without diabetes, the normal range for the HbA1c level is between 4% and 5.6%. HbA1c levels between 5.7% and 6.4% imply a higher chance of getting diabetes. Levels of 6.5% or higher indicate the presence of diabetes.

Despite the popularity of HbA1c tests, HbA1c may not reflect the real progression of diabetes given various conditions that patients are predisposed to. Interfering factors may lead to false results of HbA1c tests (Radin, 2014). In some cases, the direction of bias is predictable while in other cases not. For instance, HbA1c can be falsely elevated by iron deficiency anemias, vitamin B-12 anemias, folate deficiency anemias, asplenia and other conditions associated with decreased red cell turnover (Nitin, 2010). HbA1c measures may be lower for patients with conditions that shorten the life of the red blood cell or increase its turnover rate (Freedman et al., 2010), including acute and chronic blood loss, hemolytic anemia, splenomegaly and end-stage renal disease. While in some other cases, the direction of bias is not always easily predictable. For example, the complex interplay of glycemic control and other treatments, such as erythropoietin therapy and the treatment of uremia may influence HbA1c level in a more case-by-case way (Radin, 2014).

Due to the susceptibility of HbA1c measures to various medical and physiological factors, HbA1c measures reflect glycemia differently for different patients. Therefore, the recommendation of a single HbA1c control level for all the patients may not always be desirable. A lower glycemic level is not always better, and over-controlling glycemia can lead to faint or compromised life quality. Especially for patients who are older, with severe complications, or in the phase of early post-surgery rehabilitation, HbA1c is not supposed to be controlled for them as low as for younger and healthier patients.

As a response to the progress in medical research, the idea of glucose control has been evolving over time. In the past, an HbA1c of 7% was considered the golden rule of health for everyone. In recent years, however, the importance of a patient-centered approach to managing HbA1c levels has been recognized, which may better correspond to the patient’s needs for diabetes management and their personal conditions and preferences. To fill this gap, we introduce the HbA1c dataset to study the relationship between HbA1c control and the events of diabetic patients. Specifically, the sample is a subset of a type II diabetes dataset which is drawn from a large Midwestern multi-specialty physician and patient group. Electronic Health Record (EHR) datasets are merged with Medicare claim data to obtain laboratory records, information on primary care visits, and medication history. All patients have been tracked for a specific time period from the first quarter of 2003 to the fourth quarter of 2011. The full observational records are not available for the majority of patients.

Patients identified with diabetes must have at least 1 inpatient or SNF Medicare claim or at least two carrier claims for diabetes, which should be less than 2 years apart. The diagnosis is defined by one or more of the

following categories: the ICD-9 codes, Diabetes mellitus (250.xx), Polyneuropathy in diabetes (357.2), Diabetes retinopathy (362.0x), and Diabetic cataract (366.41). There are several types of patients excluded from the dataset, including those who are without continuous coverage with Medicare Parts A & B for a baseline year or for at least one subsequent quarter, without Medicare railroad benefits or were not enrolled in a Medicare HMO. Each patient needs to have at least 5 quarter time periods in one of the 3 insurers from 2003 to 2011 including 4 baseline quarters and a measurement period of at least 1 subsequent quarter. Also, the patients should have A1c records over the 5 quarter period which are available in the health system. In the sample of analysis, 51.3% of patients are female and 91.8% of the patients are white. The ages of participants at the cohort start time range from 18 to 102 years old, with an average of 65.7 and a standard deviation of 13.7.

We regard the HbA1c measurement as the treatment dose in this study. Observations are centered around the availability of HbA1c measurements. Each observation is associated with a unique HbA1c measurement from an individual patient following the baseline period. The outcome variable is the number of negative events, which include hospitalizations and Emergency Department (ED) visits that cover 90 days after the HbA1c measurement. In addition, we also include a list of short-term covariates that we average within 90 days prior to the HbA1c measurement as well as a list of long-term covariates that are averaged from up to 900 days prior to the HbA1c measurement. After removing subjects with missing baseline covariates, 8126 patients remain. An illustration of the data structure can be found in Figure 3. Each observation is generated from an outcome period (90 days after the HbA1c measure), a short-term covariate period (90 days before the HbA1c measure) and a long-term covariate period (including the baseline information).

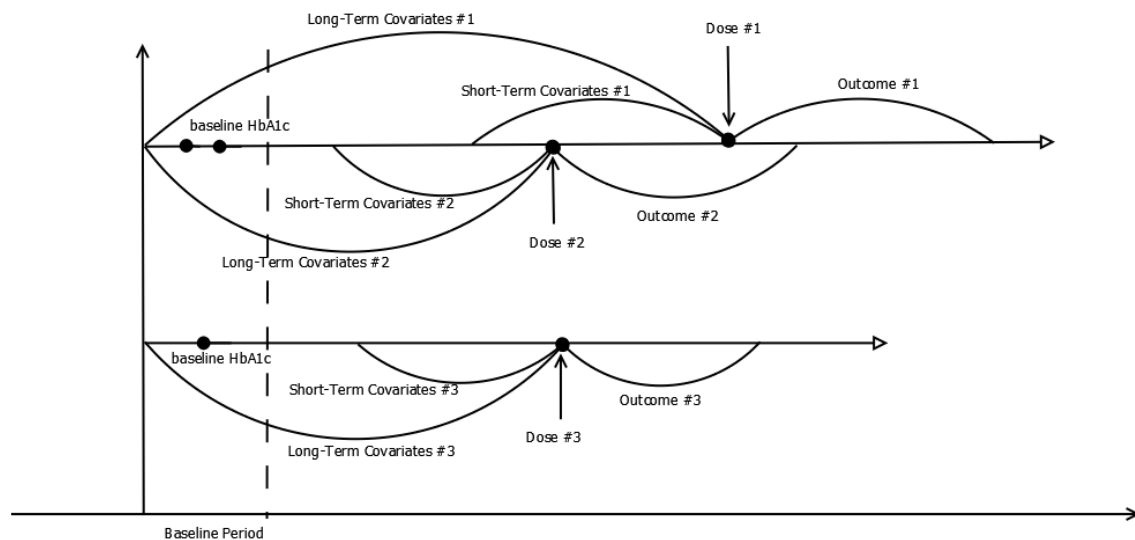


Figure 3: Illustration of Data Structure

The covariates in the data set fall in the following categories.

1. Demographics: age, gender, race
2. Long-term measurements: height, the indicator of disability at baseline, length of enrollment, average of HbA1c in the past 91-900 days, the indicator of Medicaid, number of measurements recorded in past the 91-900 days, Sulfonylureas intake, Insulin intake, and other medications intake
3. Short-term measurements: number of HbA1c measurements in the past 1-90 days; average of HbA1c in the past 1-90 days, number of hospitalization, number of ED-visits; low-density lipoprotein cholesterol, systolic blood pressure, diastolic blood pressure, BMI, weight, cardiovascular disease, ischemic heart disease, congestive heart failure, hepatocellular carcinoma, hypoglycemia and injury, infections; and 18 other comorbidities.