

Gaussian Processes for Individualized Continuous Treatment Rule Estimation

Pavel Shvechikov*

Evgeniy Riabenko*

Abstract

Individualized treatment rule (ITR) recommends treatment on the basis of individual patient characteristics and the previous history of applied treatments and their outcomes. Despite the fact there are many ways to estimate ITR with binary treatment, algorithms for continuous treatment have only just started to emerge. We propose a novel approach to continuous ITR estimation based on explicit modelling of uncertainty in the subject's outcome as well as direct estimation of the mean outcome using gaussian process regression. Our method incorporates two intuitively appealing properties - it is more inclined to give a treatment with the outcome of higher expected value and lower variance. Experiments show that this direct incorporation of the uncertainty into ITR estimation process allows to select better treatment than standard indirect approach that just models the average. Compared to the competitors (including OWL), the proposed method shows improved performance in terms of value function maximization, has better interpretability, and could be easier generalized to multiple interdependent continuous treatments setting.

Key Words: Dose Finding, Individualized Treatment Rule, Gaussian Process Regression, Continuous Treatment, Personalized Medicine

1. Introduction

Many clinician studies conclude that reactions to a fixed treatment could vary significantly among individuals in a population (Ishigooka et al., 2000; Havlir et al., 2011; Bolondi et al., 2012). Such heterogeneity of treatment effects (Kravitz et al., 2004) raises a question about the applicability of this "one size fits all" approach to clinical practice and leads naturally to personalized medicine paradigm – the idea of healthcare improvement as a result of tailoring treatments on the basis of patient individual characteristics.

1.1 Treatment rules

One of the crucial concepts in personalized medicine research is *individualized treatment rule* (ITR) (aka individualized treatment *regime*) which is usually defined as a mapping from measured patient characteristics to recommended treatment (Laber and Qian, 2017). An obvious goal of ITR construction is a maximization of patients' response to treatment averaged over all individuals in a population. This goal is sometimes referred to as expected reward in a population or *value function*. ITR estimation is a procedure which uses data from either observational or randomized clinician studies to fulfill this goal.

Personalized medicine often comes in one of the two forms (Lipkovich et al., 2017): either identification of the subgroups of patients who benefit from a given treatment, or identification of the optimal treatment for a specific patient. The latter form (which we also adhere to in this work) could be further subdivided by the approach to ITR estimation:

1. *Indirect approach* usually consists of a two-stage procedure. First, modeling of either conditional mean outcomes or contrasts between mean outcomes takes place.

*National Research University Higher School of Economics, 20 Myasnitskaya Ulitsa, Moscow 101000 Russia. The reported study was funded by RFBR, according to the research project No. 16-07-01192 A.

Second, the ITR is derived as the rule assigning the treatments that maximize the model of the conditional mean outcome.

2. *Direct approach* approximates expected reward in a population with its sample estimate as a function of ITR. Maximization of this expected reward estimate with regard to parameters of ITR yields the desired treatment rule.

Specific ITR estimation methods differ a lot depending on whether the treatment is discrete or continuous. In the latter case ITR is also called *individualized dose rule* (IDR) and is of a particular practical importance because of its usage in therapies balancing drug dose between insufficient (not causing any effect) and excessive (causing severe adverse effects). One example of such dose administering therapies is thrombosis prevention using a drug called warfarin, which is shown to be more effective when prescribed on the basis of environmental and genetic patient factors (International Warfarin Pharmacogenetics Consortium et al., 2009). Another illustrative example is a treatment of inflammatory bowel disease with thiopurines (Teml et al., 2007). These examples motivate a strong need in developing reliable and efficient methods for ITR estimation.

There is a substantial amount of literature devoted to ITR estimation dealing with discrete (mostly binary) treatment using an indirect approach (Qian and Murphy, 2011; Moodie et al., 2014; Zhang and Zhang, 2015; Tao and Wang, 2017) as well as a direct one (Zhao et al., 2012; Zhou et al., 2015). However, despite its importance, as far as we know, there are very few methods dealing with IDR estimation.

IDR estimation poses additional challenges because the set of admissible treatments has a cardinality of the continuum. There are infinitely many treatments to choose from, only a few of which might be actually present in the data. These difficulties yield trivial generalization of discrete methods to continuous treatment cases either impossible or intractable.

One of the recent works devoted to IDR estimation Laber and Zhao (2015) follows a direct approach with the focus on policy interpretability, which is achieved through the use of decision trees with splits guided by complicated purity measure. This split criterion is based on the usual value function from discrete ITR with additional smoothing. The proposed smoothing is analogous to introducing a stochastic policy with probability mass concentrated around the deterministic policy assignments observed in the training data. An interesting peculiarity of the work is that only treatments observed in the training data are considered as a split candidates.

Another recent approach to IDR estimation (Chen et al., 2016) is devoted to developing an extension of O-learning algorithm (Zhao et al., 2012). It resembles in spirit the work of Laber and Zhao (2015) because of chosen *direct* approach to IDR estimation and the presence of smoothing in value function estimate. However, unlike in Laber and Zhao (2015), the smoothing is used only as an intermediate step required for optimization purposes, so that the annealing of the kernel bandwidth to zero gives precisely the same value function as is used for discrete ITR estimation.

Specifically, Chen et al. (2016) introduced a regression based approach inspired by the Zhao et al. (2012) and SVM; the proposed algorithm comprises global optimization task tackled by DC procedure which is known to depend on initialization and could guarantee only local convergence.

All of the aforementioned algorithms follow the direct approach to IDR estimation consisting of maximizing mean conditional reward over individuals in a population (aka value function), but do not deal with variance of the patients' potential outcome, which might be crucial for a real world application.

Consider a simple example when patient is told to choose between two mutually exclu-

sive options:

1. Take a drug A with 0.99 chance of recovery (reward = 100) and 0.01 probability to acquire physical disability (reward = -1000).
2. Take a drug B with 0.85 chance of recovery (reward = 100) and 0.15 probability not to experience any effect (reward = 0).

Mean reward in the first case is 89 while in the second one it is 85, suggesting that rational patient should always choose the first option. However, it is known that people are not rational and tend to averse losses (Kahneman and Tversky, 1979). This common behavior was also shown to be true in a health domain (Attema et al., 2013). Therefore, it seems very natural to explicitly incorporate reasoning about mean and variance of a possible outcome into an ITR. Nevertheless, such reasoning is impossible when using *direct* approach.

1.2 Indirect methods

Our work aims to develop an *indirect* approach to ITR estimation. *Indirect* methods are usually criticized for potential overfitting during the first stage (fitting regression model to rewards) and for possible model misspecification.

In defense of *indirect* methods, we argue that the concerns about misspecification could be reduced to a great extent through the use of a flexible model that could fit the data well (Taylor et al., 2015).

However, the more powerful model is, the more it is susceptible to overfitting. To cope with the overfitting problem on the level of model design, we have chosen to follow Bayesian paradigm, which allow to do natural regularization through the averaging over model parameters and is known to be free of overfitting problem (Bishop, 2006). Nevertheless, any Bayesian model becomes prone to overfitting as soon as optimization with regards to hyperparameters is involved (so-called type II maximum likelihood procedure). Careful treatment of overfitting requires either very small number of hyperparameters or fully Bayesian inference with Markov chain Monte Carlo methods (Neal, 1997; Gramacy, 2007).

Another argument in favor of *indirect* methods is their capacity to estimate the uncertainty of the expected outcome. Such estimation is impossible when following *direct* approach but could be invaluable in practice, providing the means for balancing exploration and exploitation.

General task of ITR estimation is similar to the classic reinforcement learning (RL) setting (Sutton and Barto, 1998): patient characteristics, their treatments and responses are direct analogs of state, action and reward concepts from RL literature. Moreover, *direct* and *indirect* approaches to ITR estimation in RL community are called value-based and policy-based approaches. Recently, it was shown (O'Donoghue et al., 2016; Schulman et al., 2017) that these two approaches share much more in common than it was believed before. To some extent this finding reconciles two different ITR estimation approaches and gives rise to an unexplored idea of their combination.

The contributions of this work are summarized as follows.

1. We propose a new *indirect* approach to IDR estimation that benefits from separation of uncertainty due to finite sample estimation and irreducible variance of patient outcomes.
2. We show how to reduce the hard task of continuous inversion to simple univariate global optimization problem which is well-known in active learning and Bayesian optimization literature.

3. We evaluate the proposed method on various simulation studies, confirming its superior performance for randomized dose-finding clinical trials.

2. Proposed method

We assume that train data comes from randomized dose trial and every training record consists of three major components:

1. patient covariates $\mathbf{c} \in \mathbb{R}^p$;
2. treatment doze $a \in [0, 1]$, selected at random within safe bounded dose range, which w.l.o.g. is assumed to be equal to $[0, 1]$ interval;
3. patient outcome $r \in [r_{lo}, r_{hi}]$, $r_{lo}, r_{hi} \in \mathbb{R}$.

Note that we do not restrict the sign of outcomes but do require them to be bounded. We also introduce a shorthand notation $\mathbf{x}_i = [\mathbf{c}_i^\top, a_i]^\top$ for a concatenation of covariate vector and a treatment value.

As was mentioned previously, *indirect* approach to IDR estimation requires building a model of the outcome value given patient covariates and treatment doze. One of the models meeting all criteria of being sufficiently flexible, robust to overfitting due to Bayesian nature, and able to make probabilistic inference is Gaussian process (GP) regression model (Rasmussen and Williams, 2005), which is the basis of this work.

2.1 Gaussian process framework

We proceed with minimal introduction to GP regression. For more thorough discussion one could refer to Rasmussen and Williams (2005).

Gaussian process is a collection of random variables any finite number of which have joint Gaussian distribution. A GP is completely specified with its mean and covariance functions which we will denote as $m(\mathbf{x})$ and $k(\mathbf{x}, \mathbf{x}')$ respectively. For clarity and without loss of generality (Rasmussen and Williams, 2005) we will consider Gaussian process f with zero mean $m(\mathbf{x}) = 0$ and denote $f(\mathbf{x}) \sim \mathcal{GP}(0, k(\mathbf{x}, \mathbf{x}'))$.

We assume that we were given a training dataset $\{(\mathbf{x}_i, y_i) \mid i = 1, \dots, n\}$ that consists of pairs of object description \mathbf{x}_i and noisy observation of function realization, i.e. $y_i = f(\mathbf{x}_i) + \varepsilon$, where the noise $\varepsilon \sim \mathcal{N}(0, \sigma_n^2)$.

Having specified a joint Gaussian prior over the test output f_* and noisy train outputs \mathbf{y} , the inference is carried out with conditioning on observed train outputs. This procedure allows to compute the mean and the variance of f_* posterior distribution given the test point \mathbf{x}_* using the following equations:

$$\begin{aligned} f_* &\sim \mathcal{N}(\bar{f}_*, \mathbb{V}\text{ar}[f_*]) \\ \bar{f}_* &= \mathbf{k}_*^\top (K + \sigma_n^2 I)^{-1} \mathbf{y} \\ \mathbb{V}\text{ar}[f_*] &= k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^\top (K + \sigma_n^2 I)^{-1} \mathbf{k}_*, \end{aligned} \tag{1}$$

where $\mathbf{k}_*^\top = (k(\mathbf{x}_*, \mathbf{x}_1), \dots, k(\mathbf{x}_*, \mathbf{x}_n))$ and covariance matrix K has entries $K_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$. To reduce clutter we will subsequently denote $(K + \sigma_n^2 I)^{-1}$ as Λ .

One of the main constituents of GP flexibility is covariance function. The choice of covariance function for particular application can be guided by Bayesian model selection (Rasmussen and Williams, 2005), however, for illustration purposes we have chosen probably the most commonly used squared exponential function with separate length-scales. We

also extend this covariance function with the noise (aka nugget) term:

$$k(\mathbf{x}_p, \mathbf{x}_q) = \sigma_f^2 \exp \left\{ - \sum_{i=1}^p \frac{(x_{pi} - x_{qi})^2}{2\theta_i} \right\} + \sigma_n^2 \delta_{pq},$$

where θ_i is called length-scale, x_{pi} is a p -th feature of i -th object and δ_{pq} is the Kronecker delta function.

Note that all of the subsequent derivations could be replicated with other popular choices of covariance function (i.e. the Matérn class) and are not unique to squared exponential covariance function.

2.2 Method formulation

Motivation of the proposed method is as follows. To solve the prescription problem for a new patient with characteristics \mathbf{c}_* , doctor, given the history of previously seen patients $\{(\mathbf{c}_i, a_i, r_i)\}_{i=1}^n$, have to simultaneously solve two conflicting tasks:

1. Doctor has to be **confident** about potential dose effect and this confidence should be based on previous cases. The more similar are patient characteristics \mathbf{c}_* to the ones from previous practice \mathbf{c}_i , the more doctor is sure that giving the same a_i to patient with \mathbf{c}_* will result in r_i .
2. Doctor has to prescribe **the best** treatment possible. The bigger is r_i , the more doctor wants a_* to be similar to a_i .

These two intuitive observations coupled with the notion of loss aversion in health domain (Attema et al., 2013) motivate prescription of the dose which maximizes probabilistic lower surface for the potential outcome. That is, in terms of mean and variance of posterior distribution over f_*

$$a_* = \arg \max_a \bar{f}_* - s \sqrt{\text{Var}[f_*]}, \quad (2)$$

where s is a scaling coefficient controlling the penalty for the uncertainty in the outcome.

We argue that this simple heuristic, which hereinafter will be referred to as LCSL (lower confidence surface learning):

1. is a strong IDR competitive with the previously proposed approaches;
2. provides the means to account for loss aversion and uncertainty;
3. provides the means to perform adaptive trial design;
4. allows to explicitly address an exploration-exploitation tradeoff.

2.3 Modelling correlation with exponential power family

Formula (2) does not gives us a method to perform maximization over treatment domain. So we need to design a constructive way to carry out such an inversion.

Lets rewrite each of the terms in (2) separately assuming that covariance function is squared exponential with separate length-scales. For the sake of brevity we omit length-scales in following derivations

$$k(\mathbf{x}_p, \mathbf{x}_q) = \sigma_f^2 \exp \left\{ - \frac{\|\mathbf{x}_p - \mathbf{x}_q\|_2^2}{2} \right\} = \sigma_f^2 \exp \left\{ - \frac{\|\mathbf{c}_p - \mathbf{c}_q\|_2^2}{2} \right\} \exp \left\{ - \frac{\|a_p - a_q\|_2^2}{2} \right\}$$

Lets denote

$$K(\mathbf{c}_*) = \sigma_f^2 \cdot \text{diag} \left(\exp \left\{ - \frac{\|\mathbf{c}_* - \mathbf{c}_1\|_2^2}{2} \right\}, \dots, \exp \left\{ - \frac{\|\mathbf{c}_* - \mathbf{c}_n\|_2^2}{2} \right\} \right)$$

Then one could write \bar{f}_* in a form of a linear combination of exponential terms where α_i depends only on the train data and known covariates \mathbf{c}_* of test object:

$$\begin{aligned}\bar{f}_* &= \left(\exp \left\{ -\frac{\|a - a_1\|^2}{2} \right\}, \dots, \exp \left\{ -\frac{\|a - a_n\|^2}{2} \right\} \right) K(\mathbf{c}_*)^\top \Lambda \mathbf{y} \\ &= \sum_i \alpha_i \exp \left\{ -\frac{\|a - a_i\|^2}{2} \right\},\end{aligned}$$

where $\alpha_i = [K(\mathbf{c}_*)^\top \Lambda \mathbf{y}]_i$.

The same transformation could be done for $\text{Var}[f_*]$:

$$\begin{aligned}\text{Var}[f_*] &= k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^\top \Lambda \mathbf{k}_* \\ &= k(\mathbf{x}_*, \mathbf{x}_*) - \sum_{i,j} \gamma_{ij} \exp \left\{ -\frac{\|a - a_i\|^2 + \|a - a_j\|^2}{2} \right\}\end{aligned}$$

Here $\gamma_{ij} = \sigma_f^4 \Lambda_{ij} \exp \left\{ -\frac{\|\mathbf{c}_* - \mathbf{c}_i\|^2 + \|\mathbf{c}_* - \mathbf{c}_j\|^2}{2} \right\}$ and $\forall x_* k(\mathbf{x}_*, \mathbf{x}_*) = \sigma_f^2$, so neither β_{ij} nor $k(x_*, x_*)$ depends on unknown a .

To sum up, our optimization task could be simplistically formulated as follows

$$\max_a \sum_i \alpha_i \exp \left\{ -\frac{(a - a_i)^2}{2} \right\} - s \sqrt{k(\mathbf{x}_*, \mathbf{x}_*) - \sum_{i,j} \gamma_{ij} \exp \left\{ -\frac{(a - a_i)^2 + (a - a_j)^2}{2} \right\}} \quad (3)$$

2.4 Global optimization

As is easily seen, the objective (3) is a sum of unnormalized exponential functions with coefficients α_i, γ_{ij} of arbitrary sign. Thus, to allow for explicit maximization of this non-convex objective, one need to employ global optimization methods.

Even though the task (3) is a very hard one, it is surely not a new one – problems of this type are inevitable in Bayesian optimization (Brochu et al., 2010), arising during maximization of the so-called acquisition function.

There are many methods to solve global optimization problem of the form (3), i.e. DIRECT (Jones et al., 1993), which is derivative free and deterministic. In this work we adhere to a simpler yet efficient two-step approach to global optimization that is often implemented in Bayesian optimization frameworks (Bayesian Optimization, 2014; Spearmint, 2014). First, seeds are generated – either uniformly at random, in a form of fixed step grid or as a Sobol’ sequence (Sobol’, 1967). Second, function is evaluated at each seed and (optionally) local optimization method is started with each seed as an initial point. Frequent choice of such local optimization algorithm is L-BFGS-B (Zhu et al., 1997).

2.5 Interpretability of the model

Flexibility of any model usually comes with the price of reduced interpretability which is often of particular interest for practical applications. However, as is evident from formula (1), mean predictions of GP regression are additive with respect to covariance terms \mathbf{k}_* . This means that, given test object \mathbf{x}_* , one could argue which objects from the train dataset influence mean prediction the most based on the components of covariance vector \mathbf{k}_* and coefficients $[\Lambda \mathbf{y}]_i$. The same reasoning holds for variance interpretation with the only difference of quadratic dependence replacing the linear one.

Additionally, the learned values of length-scales could serve as indicators of feature importances, giving rise to the so-called automatic relevance determination (ARD) effect.

These simple model interpretations coupled with the recent work devoted to the trade-off between accuracy and interpretability (Plate, 1999) along with possible interpretability increase due to additive GP modelling (Duvenaud et al., 2011) allow to get even closer to interpretability level of linear and tree-based models.

2.6 Applicability to observational dose trials

The requirement of training data being obtained from randomized study is by no means necessary. Under usual assumptions in casual inference literature (Hernán and Robins, 2017) the proposed method could be seen as a model of counterfactual mean outcome. These assumptions are as follows:

1. Conditional exchangeability (aka strong ignorability) – potential outcomes R^a are conditionally independent of A given covariates C , i.e. $\forall a, R^a \perp A | C$.
2. Positivity – there is nonzero probability of assigning any of the treatment levels to any of possible objects, i.e. $\forall C: p(C) > 0, \forall a p(a | C) > 0$.
3. Consistency – each observed outcome is precisely the one that would have been observed under given treatment level, i.e. $R = R^A$.

In case these assumptions hold true, structural model $\mathbb{E}[R^a | C = c]$ of counterfactual outcome R^a is equal to conditional mean model $\mathbb{E}[R | C = c, A = a]$. This imply that the proposed method falls under umbrella of *outcome regression* causal inference models.

3. Numerical results

All experiments, unless stated otherwise, were conducted as follows. Train and test sample generation was repeated 50 times for each combination of scenario, train sample size, algorithm and variance penalty factor. Test sample size was fixed at the value of 1000. Rewards in train dataset were scaled to fit into $[0, 1]$. For each train dataset we fit a Gaussian process regression that was allowed to internally restart optimization of marginal likelihood 10 times. The model maximizing marginal likelihood across these 10 restarts is chosen for subsequent evaluation on test data. Approximate model performance is obtained by plugging predicted optimal dose for test objects into the true underlying state-action value function $Q(C, A)$. Method quality is reported in terms of $\hat{\mathcal{V}}(f)$, which is defined as approximate model performance averaged over all data generations.

We have empirically observed that the choice of a particular global optimization method out of those described in section 2.4 does not change performance much. Thus, we report all of the results using the simplest one – fixed step grid evaluation with number of grid point fixed at 50 for all of the simulation scenarios and settings.

In this section we refer to our method as L_{CSL}. X , where X specifies the value of variance penalty (3) by the equation $s = \Phi^{-1}(X/100)$, where Φ is normal CDF.

All our experiments are built on top of Gaussian process framework GPy (2012).

3.1 Simulation study

In this section we first present results of two simulation experiments which show that for simple scenarios our method works remarkably well.

We denote \mathbf{C} for patient covariates matrix, i.e. $\mathbf{C} = [\mathbf{c}_1, \dots, \mathbf{c}_n]^\top$, where \mathbf{c}_i is a column vector of i -th patient characteristics; A for a column vector of treatments, i.e. $A = [a_1, \dots, a_n]^\top$; C_i for an i -th column of \mathbf{C} matrix.

Our first scenario simulates optimal treatment dependence on single covariate in a form of parabola (figure 1, left):

$$R \sim \mathcal{N}(Q_1(\mathbf{C}, A), 0.01), \quad A \sim \mathcal{U}(0, 1), \quad \mathbf{c}_i \sim \mathcal{U}(0, 1)$$

$$Q_1(\mathbf{C}, A) = -100 \cdot (f_1^{opt}(\mathbf{C}) - A)^2$$

$$f_1^{opt}(\mathbf{C}) = 4 \cdot (C_1 - 0.5)^2$$

Our second scenario is a more complex one – it simulates optimal treatment assignment as a highly nonlinear function with local periodicity and abrupt change in the middle of covariate space (figure 1, right) :

$$R \sim \mathcal{N}(Q_{2s}(\mathbf{C}, A), 0.01), \quad A \sim \mathcal{U}(0, 1), \quad \mathbf{c}_i \sim \mathcal{U}(0, 1)$$

$$Q_2(\mathbf{C}, A) = -100 \cdot (f_2^{opt}(\mathbf{C}) - A)^2$$

$$f^{runge}(x) = \cos(3\pi(4x - 0.3)) / (1 + 25(4x - 0.55)^2)$$

$$f^{step}(x) = 0.1x \cdot (10 + \sin(20x) + \sin(50x)) - 1.3$$

$$f_2^{opt}(\mathbf{C}) = (f^{runge}(C_1) + f^{step}(C_1) \cdot \mathbb{I}(C_1 > 0.5)) / 1.5 - 0.7$$

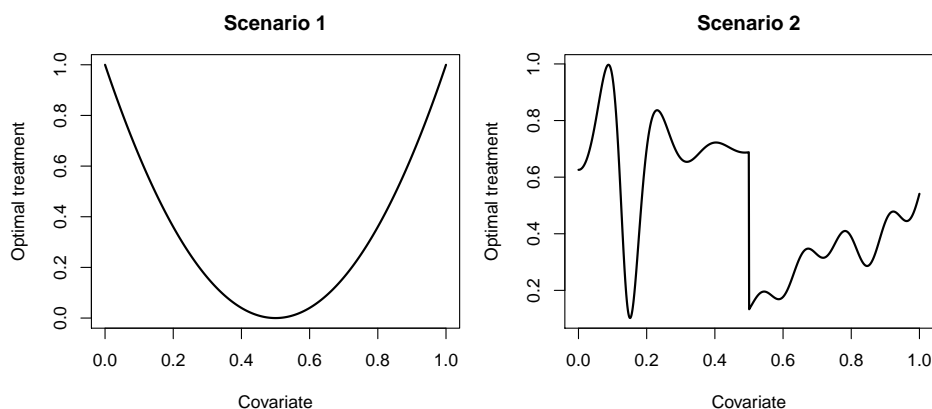


Figure 1: Illustration of $f^{opt}(\mathbf{C})$ from each of the proposed scenarios 1 and 2.

Because true $f^{opt}(\mathbf{C})$ is nonlinear in both scenarios, we don't include linear OWL (L-OWL) from Chen et al. (2016) into comparison. Kernel OWL (K-OWL) results are obtained by reusing the code from the supplementary material of Chen et al. (2016).

Simulation results (table 1, 2) show that the proposed method works extremely well in both scenarios, showing unimprovable results for the first scenario.

Table 1: Mean and std of $\hat{\nu}(f)$, Scenario 1

Sample size	LCSL.95	K-OWL
50	-0.01 (0.00)	-2.30 (1.37)
100	0.00 (0.00)	-0.93 (0.50)
200	0.00 (0.00)	-0.42 (0.18)
400	0.00 (0.00)	-0.23 (0.08)
800	0.00 (0.00)	-0.17 (0.06)

Table 2: Mean $\hat{\mathcal{V}}(f)$, Scenario 2

Sample size	LCSL.95	K-OWL
50	-3.78 (1.39)	-3.92 (0.63)
100	-2.19 (1.31)	-3.63 (0.37)
200	-0.40 (0.26)	-3.37 (0.23)
400	-0.22 (0.21)	-3.21 (0.20)
800	-0.10 (0.06)	-3.06 (0.19)

We proceed with results of LCSL on several more complex and higher dimensional simulations first introduced in Chen et al. (2016). In particular, we evaluate and compare our method on three scenarios, first two of which simulate randomized trial and the last one simulates observational trial.

Scenario 3, randomized trial:

$$R \sim \mathcal{N}(Q_3(\mathbf{C}, A), 1), \quad A \sim \mathcal{U}(0, 2), \quad \mathbf{c}_i \sim \mathcal{U}(-1, 1)^{30}$$

$$Q_3(\mathbf{C}, A) = 8 + 4C_1 - 2C_2 - 2C_3 - 25 \cdot (f_3^{opt}(\mathbf{C}) - A)^2$$

$$f_3^{opt}(\mathbf{C}) = 1 + 0.5C_1 + 0.5C_2$$

Scenario 4, randomized trial:

$$R \sim \mathcal{N}(Q_4(\mathbf{C}, A), 1), \quad A \sim \mathcal{U}(0, 2), \quad \mathbf{c}_i \sim \mathcal{U}(-1, 1)^{10}$$

$$Q_4(\mathbf{C}, A) = 8 + 4 \cos(2\pi C_2) - 2C_4 - 8C_5^3 - 15 \cdot |f_4^{opt}(\mathbf{C}) - A|$$

$$f_4^{opt}(\mathbf{C}) = 0.6 \cdot \mathbb{I}(|C_1| \geq 0.5) + C_4^2 + 0.5 \log(|C_7| + 1)$$

Scenario 5, observational trial:

$$R \sim \mathcal{N}(Q_5(\mathbf{C}, A), 1), \quad A \sim \text{Trunc}\mathcal{N}(f_5^{opt}(\mathbf{C}), 0, 2, 0.5), \quad \mathbf{c}_i \sim \mathcal{U}(-1, 1)^{10}$$

$$Q_5(\mathbf{C}, A) \equiv Q_4(\mathbf{C}, A)$$

$$f_5^{opt}(\mathbf{C}) \equiv f_4^{opt}(\mathbf{C})$$

Results of randomized study simulation (tables 3, 4) allow to conclude that our method outperforms every other competitor (often by a large margin) in a randomized trial setting. Such superiority originates from the randomization of treatment assignment, which increases the accuracy of global response surface model and, as a result, improves quality of treatment recommendations.

Table 3: Scenario 3: Mean and (std) of $\hat{\mathcal{V}}(f)$. Results with * are from Chen et al. (2016)

Sample size	LCSL.95	K-OWL*	L-OWL*	SVR*	Lasso*
50	7.65 (0.92)	4.78 (0.48)	4.83 (1.40)	-12.21 (7.53)	-15.62 (6.61)
100	7.87 (0.11)	5.69 (0.40)	5.39 (0.93)	-2.57 (6.34)	-9.76 (5.21)
200	7.93 (0.11)	6.68 (0.26)	6.85 (0.34)	3.46 (1.97)	-1.91 (4.65)
400	7.97 (0.09)	7.28 (0.15)	7.41 (0.14)	6.13 (0.47)	7.95 (0.01)
800	7.99 (0.09)	7.54 (0.08)	7.67 (0.08)	7.36 (0.12)	7.97 (0.01)

However, for simulation of observational study (table 5) we observe that our method yields superior results only when given enough data. Small samples do not allow to properly model response surface because all of the points are already concentrated in near optimal regions, giving insufficient information about treatment dependence on covariates.

Table 4: Scenario 4: Mean and (std) of $\hat{\mathcal{V}}(f)$. Results with * are from Chen et al. (2016)

Sample size	LCSL.95	K-OWL*	L-OWL*	SVR*	Lasso*
50	2.19 (1.97)	2.00 (0.29)	1.16 (0.71)	-1.96 (1.70)	-5.58 (2.27)
100	4.46 (0.76)	2.19 (0.43)	1.57 (0.52)	0.24 (1.42)	-4.12 (2.17)
200	5.30 (0.24)	2.84 (0.37)	2.02 (0.30)	2.01 (0.84)	-3.37 (1.38)
400	5.58 (0.20)	3.69 (0.27)	2.30 (0.18)	3.47 (0.37)	-0.92 (3.12)
800	5.75 (0.21)	4.41 (0.19)	2.49 (0.10)	4.35 (0.19)	1.92 (0.69)

Table 5: Scenario 5: Mean and (std) of $\hat{\mathcal{V}}(f)$. Results with * are from Chen et al. (2016)

Sample size	LCSL.95	K-OWL	K-OWL*	SVR*
50	-1.15 (2.55)	1.88 (0.60)	NA	NA
100	-0.45 (3.91)	2.43 (0.42)	NA	NA
200	2.23 (3.94)	3.27 (0.32)	3.29 (0.28)	-0.95 (1.57)
400	4.26 (3.17)	4.22 (0.28)	NA	NA
800	5.56 (0.65)	4.95 (0.21)	4.91 (0.14)	3.04 (0.52)

This effect could be mitigated to some extent by large variance penalty (figure 2, last column).

LCSL performance dependence on the variance penalty coefficient, denoted by s in objective (3), needs additional elaboration. We have varied s on standard normal quantile grid (from 0.5 to 0.99 quantile with a 0.01 step) and observed a sustainable relationship between penalty value and method performance (figure 2).

The first interesting effect, equally present in all scenarios, is that strong variance penalty always improves upon mean prediction given very small train data (positive trend on every plot in the first row of figure 2).

The second observation is that increasing train sample size in randomized studies reduces the optimal variance penalty (first two columns in figure 2). This observation could be explained by increase of model accuracy and, as a consequence, diminished need in variance penalty. Note that for more complex setup (scenario 2) optimal penalty shrinks slower than for a simpler one (scenario 1). However, we also report a consistent small improvement in average $\hat{\mathcal{V}}(f)$ irrespective of train data size for small values of penalty (figure 2). This effect can be viewed as a regularization, preventing the model of too assertive extrapolation.

The last noteworthy remark concerns observational study (last column of figure 2). In such setting we observe that large variance penalty increases performance irrespective of train sample size. That is because data generation policy ($A \sim \mathcal{N}(f_4^{opt}(\mathbf{C}), 0.5)$) is by design close to optimal. Thus, in observational studies the value of variance penalty should be treated as a proxy of prior knowledge about the degree of data generation policy optimality.

We conclude that different setups need to treat variance penalty differently. Randomized studies benefit most when treating variance penalty as a regularization. In contrast, observational studies should treat variance penalty as a means for explicit reasoning about degree of data generation policy optimality.

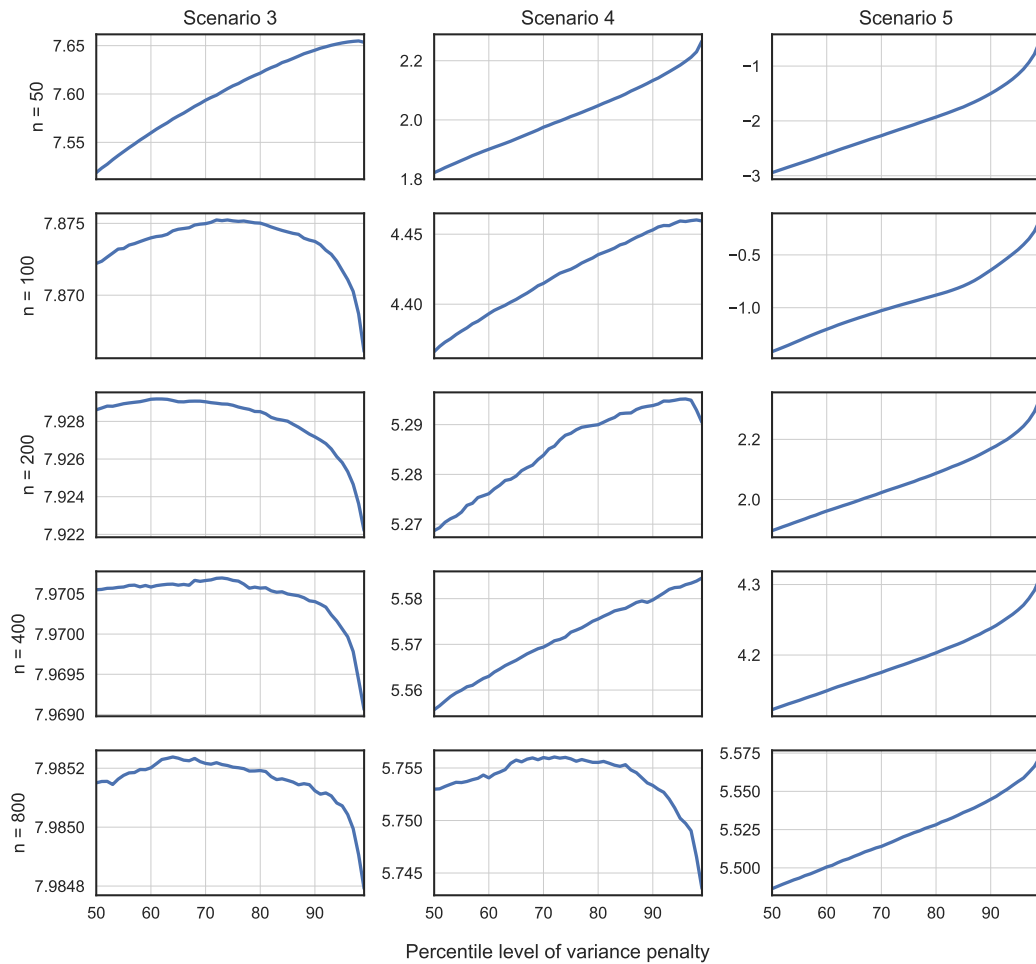


Figure 2: Dependence of average quality $\hat{V}(f)$ on s -factor values (see (3)), which are represented as normal percentile levels, i.e. 95 percentile level means $s = \Phi^{-1}(0.95)$, where Φ is normal CDF.

4. Conclusions

We have presented an *indirect* approach to IDR estimation which is based on Gaussian process regression.

In spite of obvious difficulties of *indirect* IDR estimation, we have shown that it is possible to reduce such an intractable problem to a univariate global optimization setting which is ubiquitous in Bayesian optimization and active learning.

We have discussed the advantages and disadvantages of the proposed method including explicit probability reasoning about potential outcome and computational challenges arising as a consequence of using flexible nonparametric model.

We have also reported various simulation results which not only show that our method is highly competitive with other algorithms of ITR estimation, but uncover unexpected variance penalty interpretations.

5. Future work

As was previously noted, ITR estimation shares many similarities with policy search problem arising in reinforcement learning (RL). This suggests researches dealing with ITR and DTR (dynamic treatment regime) estimation could get inspiration from already existing methods in reinforcement learning world.

For example, it would be very interesting to see an adaptation of reward weighted regression (Peters and Schaal, 2007) to *direct* DTR estimation. Also, an extension of the proposed LCSL method to multi-stage treatment assignments would be of a great interest. Such extension could probably inherit some of the useful ideas from data-efficient RL algorithm called PILCO (Deisenroth and Rasmussen, 2011).

Diffusion of statistical science and reinforcement learning may bring benefits to all parties and thus is highly desirable.

REFERENCES

- Attema, A. E., Brouwer, W. B., and Haridon, O. (2013). Prospect theory in the health domain: a quantitative assessment. *Journal of health economics*, 32(6):1057–1065.
- Bayesian Optimization (since 2014). Bayesian Optimization: Bayesian global optimization with gaussian processes in python. <https://github.com/fmfn/BayesianOptimization>.
- Bishop, C. M. (2006). Pattern recognition. *Machine Learning*, 128:1–58.
- Bolondi, L., Burroughs, A., Dufour, J.-F., Galle, P. R., Mazzaferro, V., Piscaglia, F., Raoul, J. L., and Sangro, B. (2012). Heterogeneity of patients with intermediate (bcl_c b) hepatocellular carcinoma: proposal for a subclassification to facilitate treatment decisions. In *Seminars in liver disease*, volume 32, pages 348–359. Thieme Medical Publishers.
- Brochu, E., Cora, V. M., and De Freitas, N. (2010). A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*.
- 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.
- Deisenroth, M. and Rasmussen, C. E. (2011). Pilco: A model-based and data-efficient approach to policy search. In *Proceedings of the 28th International Conference on machine learning (ICML-11)*, pages 465–472.
- Duvenaud, D. K., Nickisch, H., and Rasmussen, C. E. (2011). Additive gaussian processes. In *Advances in neural information processing systems*, pages 226–234.
- GPy (since 2012). GPy: A gaussian process framework in python. <http://github.com/SheffieldML/GPy>.
- Gramacy, R. B. (2007). An R Package for Bayesian Nonstationary, Semiparametric Nonlinear Regression and Design by Treed Gaussian Process Models. *Journal Of Statistical Software*, 19(9):1–46.
- Havilr, D. V., Kendall, M. A., Ive, P., Kumwenda, J., Swindells, S., Qasba, S. S., Luetkemeyer, A. F., Hogg,

- E., Rooney, J. F., Wu, X., et al. (2011). Timing of antiretroviral therapy for hiv-1 infection and tuberculosis. *New England Journal of Medicine*, 365(16):1482–1491.
- Hernán, M. and Robins, J. (2017). Causal inference.
- International Warfarin Pharmacogenetics Consortium et al. (2009). Estimation of the warfarin dose with clinical and pharmacogenetic data. *New England Journal of Medicine*, 2009(360):753–764.
- Ishigooka, J., Murasaki, M., Miura, S., Group, T. O. L.-P. I., et al. (2000). Olanzapine optimal dose: Results of an open-label multicenter study in schizophrenic patients. *Psychiatry and clinical neurosciences*, 54(4):467–478.
- Jones, D. R., Perttunen, C. D., and Stuckman, B. E. (1993). Lipschitzian optimization without the lipschitz constant. *Journal of optimization Theory and Applications*, 79(1):157–181.
- Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the econometric society*, pages 263–291.
- Kravitz, R. L., Duan, N., and Braslow, J. (2004). Evidence-based medicine, heterogeneity of treatment effects, and the trouble with averages. *Milbank Quarterly*, 82(4):661–687.
- Laber, E. B. and Qian, M. (2017). Evaluating personalized treatment regimes. In *Methods in Comparative Effectiveness Research*, pages 483–498. CRC Press.
- Laber, E. B. and Zhao, Y. Q. (2015). Tree-based methods for individualized treatment regimes. *Biometrika*, 102(3):501–514.
- Lipkovich, I., Dmitrienko, A., et al. (2017). Tutorial in biostatistics: data-driven subgroup identification and analysis in clinical trials. *Statistics in Medicine*, 36(1):136–196.
- Moodie, E. E., Dean, N., and Sun, Y. R. (2014). Q-learning: Flexible learning about useful utilities. *Statistics in Biosciences*, 6(2):223–243.
- Neal, R. M. (1997). Monte carlo implementation of gaussian process models for bayesian regression and classification. *arXiv preprint physics/9701026*.
- O’Donoghue, B., Munos, R., Kavukcuoglu, K., and Mnih, V. (2016). Pgg: Combining policy gradient and q-learning. *arXiv preprint arXiv:1611.01626*.
- Peters, J. and Schaal, S. (2007). Reinforcement learning by reward-weighted regression for operational space control. In *Proceedings of the 24th international conference on Machine learning*, pages 745–750. ACM.
- Plate, T. A. (1999). Accuracy versus interpretability in flexible modeling: Implementing a tradeoff using gaussian process models. *Behaviormetrika*, 26(1):29–50.
- Qian, M. and Murphy, S. A. (2011). Performance guarantees for individualized treatment rules. *Annals of statistics*, 39(2):1180.
- Rasmussen, C. E. and Williams, C. K. I. (2005). *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press.
- Schulman, J., Abbeel, P., and Chen, X. (2017). Equivalence between policy gradients and soft q-learning. *arXiv preprint arXiv:1704.06440*.
- Sobol’, I. M. (1967). On the distribution of points in a cube and the approximate evaluation of integrals. *Zhurnal Vychislitel’noi Matematiki i Matematicheskoi Fiziki*, 7(4):784–802.
- Spearmint (since 2014). Spearmint: Software package to perform bayesian optimization. <https://github.com/HIPS/Spearmint>.
- Sutton, R. S. and Barto, A. G. (1998). *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge.
- Tao, Y. and Wang, L. (2017). Adaptive contrast weighted learning for multi-stage multi-treatment decision-making. *Biometrics*, 73(1):145–155.
- Taylor, J. M., Cheng, W., and Foster, J. C. (2015). Reader reaction to a robust method for estimating optimal treatment regimes by zhang et al.(2012). *Biometrics*, 71(1):267–273.
- Teml, A., Schaeffeler, E., Herrlinger, K. R., Klotz, U., and Schwab, M. (2007). Thiopurine treatment in inflammatory bowel disease. *Clinical pharmacokinetics*, 46(3):187–208.
- Zhang, B. and Zhang, M. (2015). C-learning: a new classification framework to estimate optimal dynamic treatment regimes.
- Zhao, Y., Zeng, D., Rush, A. J., and Kosorok, M. R. (2012). Estimating individualized treatment rules using outcome weighted learning. *Journal of the American Statistical Association*, 107(499):1106–1118.
- Zhou, X., Mayer-Hamblett, N., Khan, U., and Kosorok, M. R. (2015). Residual weighted learning for estimating individualized treatment rules. *Journal of the American Statistical Association*, (just-accepted):00–00.
- Zhu, C., Byrd, R. H., Lu, P., and Nocedal, J. (1997). Algorithm 778: L-bfgs-b: Fortran subroutines for large-scale bound-constrained optimization. *ACM Trans. Math. Softw.*, 23(4):550–560.