

Dirac Delta Regression: Conditional Density Estimation with Clinical Trials

Eric V. Strobl

*Department of Psychiatry
Vanderbilt University Medical Center
Nashville, TN 37212, United States*

ERIC.STROBL@VUMC.ORG

Shyam Visweswaran

*Department of Biomedical Informatics
University of Pittsburgh
Pittsburgh, PA 15206, United States*

SHV3@PITT.EDU

Editor: Thuc Le, Jiuyong Li, Gregory Cooper, Sofia Triantafyllou, Elias Bareinboim, Huan Liu, Negar Kiyavash

Abstract

Personalized medicine seeks to identify the causal effect of treatment for a particular patient as opposed to a clinical population at large. Most investigators estimate such personalized treatment effects by regressing the outcome of a randomized clinical trial (RCT) on patient covariates. The realized value of the outcome may however lie far from the conditional expectation. We therefore introduce a method called *Dirac Delta Regression* (DDR) that estimates the entire conditional density from RCT data in order to visualize the probabilities across all possible outcome values. DDR transforms the outcome into a set of asymptotically Dirac delta distributions and then estimates the density using non-linear regression. The algorithm can identify significant differences in patient-specific outcomes even when no population level effect exists. Moreover, DDR outperforms state-of-the-art algorithms in conditional density estimation by a large margin even in the small sample regime. An R package is available at <https://github.com/ericstrobl/DDR>.

Keywords: Randomized Clinical Trial, Conditional Density Estimation, Causality, Small Sample Regime

1. The Problem

Randomized clinical trials (RCTs) are the gold standard for inferring causal effects of treatment in biomedicine. These trials utilize the notion of *potential outcomes*, where we consider a set of treatments \mathcal{T} and a random variable T denoting the treatments assigned from \mathcal{T} . We then postulate the existence of a random vector $\mathbf{Y} \triangleq (Y(t))_{t \in \mathcal{T}}$ representing the causal outcomes of all treatments (Rubin, 1974). The density $f(Y(t))$ intuitively summarizes the probabilities of all possible outcome values for $T = t$, so we seek to estimate $f(Y(t))$ for each treatment $t \in \mathcal{T}$. We can solve this problem by directly measuring \mathbf{Y} in an ideal world. However, we can only assign each patient to a single treatment $T = t$ and observe $Y(T = t)$ in practice. We thus can only sample from $f(Y(T)|T)$, even though we want to sample from $f(\mathbf{Y})$. RCTs circumvent this problem by randomizing T so that T and \mathbf{Y} are probabilistically independent, or *unconfounded*. The independence implies $f(Y(t)|T) = f(Y(t))$ for any $t \in \mathcal{T}$ which in

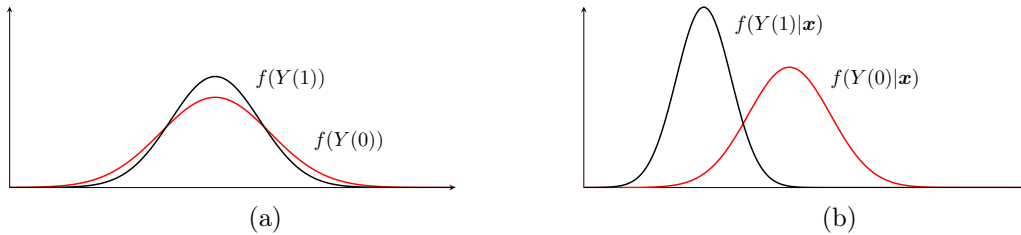


Figure 1: Treatments 0 and 1 are clinically equivalent according to the unconditional densities shown in (a). However, a patient with covariates $\mathbf{X} = \mathbf{x}$ will likely benefit more from treatment 1 as shown by the reduction in symptom severity with the conditional densities in (b).

turn implies $f(Y(T)|T) = f(Y(T))$. As a result, we can use RCT data to estimate $f(Y(t))$ for each $t \in \mathcal{T}$.

Biomedical investigators have traditionally attempted to summarize $Y(T)$ by estimating the aforementioned densities or computing simple statistics of $Y(T)$, such as the mean or variance, which average over patients. However, averages fail to capture patient heterogeneity when patients within the same clinical population respond to treatment differently. For example, all patients with breast cancer do not respond better to the medication raloxifene ($T = 1$) relative to placebo ($T = 0$), but patients who are estrogen or progesterone receptor positive often do (Martinkovich et al., 2014); the densities $f(Y(0))$ and $f(Y(1))$ therefore only differ when restricted to breast cancer patients with particular molecular receptors in this case. Biomedical investigators are thus also interested in computing statistics of $Y(T)$ given a set of patient covariates \mathbf{X} in order to elucidate patient-specific outcomes.

Most investigators in particular estimate the *conditional expectation* by regressing $Y(T)$ on \mathbf{X} (e.g., (Zhang et al., 2012b; Luedtke and van der Laan, 2016; Zhang et al., 2012a)). The conditional expectation however does not capture the uncertainty in outcome value because it assigns a single point prediction $\mathbb{E}(Y(T)|\mathbf{X} = \mathbf{x})$ to each patient. The realized outcome value for a particular patient may lie far from the conditional expectation. As a result, patients and healthcare providers instead prefer to know the conditional probabilities associated with *all* possible values of $Y(T)$ in order to make more informed clinical decisions (Akobeng, 2007; Diamond and Forrester, 1979; Bhise et al., 2018; Mah et al., 2016; Berger, 2015).

Conditional densities can fortunately summarize the desired probabilities in a single intuitive graph, especially when the potential outcomes are univariate. We therefore consider recovering $f(Y(T)|\mathbf{X})$ from RCT data. Consider for example two treatments $\mathcal{T} = \{0, 1\}$. We can estimate the unconditional densities $f(Y(0))$ and $f(Y(1))$ using RCT data as shown in Figure 1 (a). Both of the densities are very similar because they summarize the treatment outcomes across everyone recruited in the RCT, some of whom may respond well to treatment while others may not. In contrast, Figure 1 (b) displays the densities $f(Y(0)|\mathbf{x})$ and $f(Y(1)|\mathbf{x})$ estimated from the original RCT, as if we had run an RCT personalized towards a particular patient where everyone recruited into the trial had the exact same covariate values $\mathbf{X} = \mathbf{x}$. Notice that most of the patients in this trial who received $T = 1$ had a reduction in symptom severity, as measured using univariate outcomes, compared to those who received $T = 0$.

Conditional densities thus can summarize the results of a personalized RCT which may differ substantially from the results of the original RCT.

Conditional densities are unfortunately much more difficult to estimate than conditional expectations in the non-parametric setting, especially with the small sample sizes seen in RCTs (e.g., at most a few hundred samples per treatment). Existing methods struggle in this setting.

We therefore propose a new procedure called Dirac Delta Regression (DDR) that accurately estimates non-parametric conditional densities even in the small sample regime in order to intuitively visualize the probabilities associated with all possible outcome values using RCT data.

DDR transforms $Y(T)$ into a set of asymptotically Dirac delta distributions and then regresses the distributions on \mathbf{X} as described in the Section 3. The algorithm also *directly* estimates $f(Y(T)|\mathbf{X})$ using established regression procedures and tunes a minimal number of hyperparameters automatically without access to the true conditional density. We prove consistency of the procedure as well as establish the bias rate with respect to the transformation in Section 4. Experiments show that DDR (a) leads to state-of-the-art performance in practice and (b) allows biomedical investigators to estimate the densities of a personalized RCT from the original RCT data without additional data collection.

2. Related Work

We can estimate non-parametric conditional densities using several other methods besides DDR in order to visualize the probabilities associated with all values of $Y(T)$. Conditional Kernel Density Estimation (CKDE) for example estimates $f(Y(T)|\mathbf{X})$ using the ratio $f(Y(T), \mathbf{X})/f(\mathbf{X})$, where the numerator and denominator are approximated using kernel density functions (Rosenblatt, 1969; Hyndman et al., 1996; Fosgerau and Fukuda, 2010). Other algorithms first estimate the conditional cumulative distribution function (CDF) of $Y(T)$ given \mathbf{X} and then recover the associated density by smoothing the estimated CDF (Takeuchi et al., 2006; Cattaneo et al., 2018). These indirect approaches to estimating $f(Y(T)|\mathbf{X})$ however degrade accuracy in practice because they can incur error when estimating the joint density or the CDF.

Investigators therefore later proposed to directly estimate $f(Y(T)|\mathbf{X})$ instead. The earliest methods in this category, such as the Mixture Density Network (MDN), estimate $f(Y(T)|\mathbf{X})$ by performing expectation-maximization over a mixture of parametric densities (Bishop, 2006; Tresp, 2001). These methods however are very time consuming, and their accuracy depends heavily on the choice of the component densities. Sugiyama et al. (2010) and Kanamori et al. (2012) thus proposed the Least Squares Conditional Density Estimation (LSCDE) algorithm which estimates $f(Y(T)|\mathbf{X})$ in closed form using reproducing kernels. The authors nevertheless found that LSCDE performs poorly when \mathbf{X} contains multiple variables.

Researchers suggested handling the higher dimensional setting by integrating both dimension reduction and direct density estimation into a single procedure (Shiga et al., 2015; Izbicki and Lee, 2016; Tangkaratt et al., 2015). These methods, such as Series

	Small samples	Direct estimation	Non-parametric	High dimensions
CKDE			✓	
MDN		✓		
LSCDE		✓	✓	
NFE		✓	✓	
SCDE		✓	✓	✓
FC		✓	✓	✓
DDR	✓	✓	✓	✓

Table 1: Properties of prior approaches compared to DDR.

Conditional Density Estimation (SCDE), outperform LSCDE on average when sparsity or a low dimensional manifold exist, but they still struggle to accurately estimate $f(Y(T)|\mathbf{X})$ when the signal does not follow a simple structure. As a result, Izbicki and Lee (2017) introduced an algorithm called FlexCode (FC) which utilizes non-linear regression on an orthogonal series (e.g., the Fourier series). FC admits a variety of regression procedures and therefore can capitalize on the successes of regression in high dimensional estimation. Many of the proposed regression procedures however only help FC produce accurate estimates of $f(Y(T)|\mathbf{X})$ in special cases, so it is often unclear how to choose the best regressor in practice. FC also tends to over-smooth its conditional density estimates by enforcing a small number of orthogonal bases. The method therefore can have trouble handling the complexities inherent in real data like its predecessors.

The normalizing flow estimator (NFE) attempts to directly estimate $f(Y(T)|\mathbf{X})$ in the high dimensional setting using a different strategy called the normalizing flow (Trippe and Turner, 2018). A normalizing flow automatically ensures that the density integrates to one by transforming an initial density with a series of invertible transformations using a deep Bayesian neural network. NFE must tune its prior distributions, flow parameters and network architecture in addition to the network weights. The algorithm therefore tends to overfit in practice, especially in the small sample regime seen with the majority of RCTs.

In this paper, we improve upon the aforementioned works by proposing a method called DDR that (1) estimates $f(Y(T)|\mathbf{X})$ directly without approximating $f(Y(T), \mathbf{X})$ first, (2) handles the high dimensional setting using established regression procedures, and (3) balances overfitting and underfitting even with small RCTs using an automated cross-validation procedure over few hyperparameters. We summarize the differences between DDR and prior approaches in Table 1.

3. Algorithm Design

3.1 Setup

We now provide a detailed description of the DDR algorithm. We first consider independent and identically distributed (i.i.d.) data in the form of triples $\{(\mathbf{x}_i, t_i, y_i(t_i))\}_{i=1}^n$ collected from an RCT. The data corresponds to instantiations of the random variables $(\mathbf{X}, T, Y(T))$, where $Y(T)$ is assumed to be continuous, and \mathbf{X} denotes a vector of covariates (continuous, discrete or both) measured before treatment assignment. We assume that we have $f(T) > 0$. We also adopt the *strongly ignorable treatment assignment* (SITA) or *unconfoundedness* assumption

in the conditional setting which asserts that we have $T \perp\!\!\!\perp \{Y(t)\}_{t \in \mathcal{T}} | \mathbf{X}$ (Rosenbaum and Rubin, 1983; Rubin, 1990). The SITA assumption holds with an RCT because investigators must assign treatments independently of $Y(T)$ within any stratum of the covariates. The assumption also implies that we have:

$$f(Y(T)|T, \mathbf{X}) = f(Y(T)|\mathbf{X}),$$

so that the latter term can be estimated from RCT data. In this paper, we seek to estimate the density $f(Y(t)|\mathbf{X})$ for each treatment value $t \in \mathcal{T}$ using DDR.

3.2 Dirac Delta Regression

We design DDR for the setting where $Y(t)$ is continuous. However, it is informative to first consider the discrete setting. Suppose that the discrete variable $Y(t)$ takes on d distinct values in the set $\{z_j\}_{j=1}^d$. Recall that we have $\mathbb{E}[\mathbb{1}(Y(t) = z_k) | \mathbf{X}] = f_{Y(t)}(z_k | \mathbf{X})$ for each $z_k \in \{z_j\}_{j=1}^d$. We can therefore estimate the probability mass function $f(Y(t) | \mathbf{X})$ by first transforming $Y(t)$ into a set of d indicator functions. In particular, we compute $\mathbb{1}(Y(t) = z_k)$ for each $z_k \in \{z_j\}_{j=1}^d$. With n samples, we thus convert a column vector containing n samples of $Y(t)$ into an $n \times d$ matrix with the i, k^{th} entry corresponding to $\mathbb{1}(y_i(t) = z_k)$. We finally proceed with regressing the set of d indicators on \mathbf{X} to obtain an estimate of $\mathbb{E}[\mathbb{1}(Y(t) = z_k) | \mathbf{X}]$ for each $z_k \in \{z_j\}_{j=1}^d$.

The quantity $\mathbb{1}(Y(t) = z_k)$ is unfortunately almost surely equal to zero in the continuous case. The aforementioned strategy therefore fails in this setting because the $n \times d$ matrix is a matrix of zeros almost surely. Recall however that the following relation holds when $Y(t)$ is continuous (Kanwal, 2011):

$$f_{Y(t)}(z_k | \mathbf{X}) = \mathbb{E}[\delta(z_k - Y(t)) | \mathbf{X}],$$

where $\delta(z_k - Y(t))$ denotes a Dirac delta distribution, and $\{z_j\}_{j=1}^d$ is now an arbitrary set of values on \mathbb{R} (we will specify a practical choice of this set in the next subsection). We can loosely view $\delta(z_k - Y(t))$ as a function equal to infinity on a set of Lebesgue measure zero:¹

$$\delta(z_k - Y(t)) = \begin{cases} +\infty & \text{if } Y(t) = z_k, \\ 0 & \text{otherwise.} \end{cases}$$

As a result, we cannot transform the response variable into a set of Dirac delta distributions without again obtaining an $n \times d$ matrix of zeros almost surely. We can however “stretch out” the support of the Dirac delta distribution from a single point to an interval by utilizing *asymptotically* Dirac delta distributions δ_h with the stretching parameter h such that $\lim_{h \rightarrow 0^+} \delta_h(z_k - Y(t)) = \delta(z_k - Y(t))$. We in particular choose to utilize a Gaussian density $\delta_h(z_k - Y(t)) = \frac{1}{\sqrt{2\pi h}} \exp\left(-\frac{|z_k - Y(t)|^2}{2h^2}\right)$ for an unbounded outcome, or a truncated Gaussian for a bounded outcome, although many other densities are also appropriate; we may for example utilize the biweight or tricube density instead. As a general rule, densities with mean zero and lower higher order moments tend to perform better in practice. We can

1. Note that Dirac delta is more formally a generalized function or distribution, but we aim to provide the intuition here.

therefore recommend many kernel density functions $\frac{1}{h}K\left(\frac{z_k - Y(t)}{h}\right)$ used in non-parametric unconditional density estimation as well (Wand and Jones, 1994). Regardless of our choice of δ_h , we transform $Y(t)$ into a set of d asymptotically Dirac delta distributions, or kernel density functions, by computing $\delta_h(z_k - Y(t))$ for each $z_k \in \{z_j\}_{j=1}^d$. We therefore convert a column vector containing the n samples of $Y(t)$ into an $n \times d$ matrix with the i, k^{th} entry corresponding to $\delta_h(z_k - y_i(t))$.

Let $\Delta_h = \{\delta_h(z_j - Y(t))\}_{j=1}^d$ denote the set of d kernel density functions. DDR then proceeds with non-linear regression of Δ_h on \mathbf{X} using the training set in Step 1 of Algorithm 1. Thus, while DDR uses kernel density functions like CKDE, DDR directly estimates $f(Y(T)|\mathbf{X})$ instead of first approximating $f(Y(T), \mathbf{X})$ and $f(\mathbf{X})$ – both of which may be much more complicated than $f(Y(T)|\mathbf{X})$. DDR ultimately performs a separate regression for each $z_k \in \{z_j\}_{j=1}^d$, although it does not perform them independently due to the cross-validation procedure described in the next subsection. Informally, this strategy approximates the density values at each $z_k \in \{z_j\}_{j=1}^d$ on the test set just like with the discrete case because we have:

$$\lim_{h \rightarrow 0^+} \lim_{n \rightarrow \infty} \widehat{g}_h(z_k|\mathbf{X}) \xrightarrow{P} f_{Y(t)}(z_k|\mathbf{X}), \quad (1)$$

where $g_h(z_k|\mathbf{X}) = \mathbb{E}[\delta_h(z_k - Y(t))|\mathbf{X}]$ and $\widehat{g}_h(z_k|\mathbf{X}) = \widehat{\mathbb{E}}[\delta_h(z_k - Y(t))|\mathbf{X}]$ is estimated using a consistent non-linear regression method. We formalize the above intuition in Section 4, where we specifically show that the bias incurred using δ_h instead of δ is of order h^4 under an integrated squared error loss.

Algorithm 1: Dirac Delta Regression (DDR)

Input: training set $\{\mathbf{x}_i, y_i(t)\}_{i=1}^n$, test set $\{\mathbf{x}_i\}_{i=n+1}^m$

Result: conditional density estimates $\mathcal{G}_h \triangleq \{\widehat{g}_h(Y(t)|\mathbf{x}_i)\}_{i=n+1}^m$

- 1 Regress Δ_h on \mathbf{X} ; choose the optimal value of h and all other hyperparameters $\boldsymbol{\lambda}$ using cross-validation with Equation (3)
 - 2 Enforce non-negativity and normalize each element in \mathcal{G}_h using Equation (4)
 - 3 Sharpen each element in \mathcal{G}_h with Equations (5) and (4); choose the optimal value of η again with Equation (3)
-

3.3 Loss Function & Cross-Validation

We need to choose h carefully because the stretching parameter introduces bias. If we make h too large, then DDR will recover a density that is too flat. On the other hand, if we make the problem too hard with a small h , then the regression procedure will struggle to recover a density that is ultimately “too wiggly.” We therefore must pick h as well as the standard set of hyperparameters of the non-linear regression method $\boldsymbol{\lambda}$ in a principled manner.

We will choose h and $\boldsymbol{\lambda}$ using cross-validation. Recall however that we do not have access to the ground truth conditional density $f(Y(t)|\mathbf{X})$. Fortunately, we can utilize a trick by

cross-validating over the Mean Integrated Squared Error (MISE) loss function (Fryer, 1976):

$$\begin{aligned} & \int \int \left| \widehat{g}_h(z|\mathbf{x}) - f_{Y(t)}(z|\mathbf{x}) \right|^2 d\mathbb{P}_{\mathbf{X}}(\mathbf{x}) dz \\ &= \int \int \widehat{g}_h^2(z|\mathbf{x}) d\mathbb{P}_{\mathbf{X}}(\mathbf{x}) dz - 2 \int \int \widehat{g}_h(z|\mathbf{x}) f_{Y(t)}(z, \mathbf{x}) d\mathbf{x} dz + C, \end{aligned} \quad (2)$$

where C is a constant that only depends on $f(Y(t)|\mathbf{X})$. The MISE loss thus quantifies the distance of the approximated conditional density to the true conditional density similar to the mean squared error loss for regression. The empirical MISE loss takes on the following form up to a constant:

$$\frac{1}{n} \sum_{i=1}^n \int \widehat{g}_h^2(z|\mathbf{x}_i) dz - \frac{2}{n^2} \sum_{i=1}^n \widehat{g}_h(y_i(t)|\mathbf{x}_i). \quad (3)$$

Notice that computing the above quantity does not require knowledge of the ground truth $f(Y(t)|\mathbf{X})$. We therefore can use Equation (3) to tune both h and λ in Step 1 of DDR.

Observe however that we need to estimate $f_{Y(t)}(z|\mathbf{X})$ over all values of z rather than just on the grid $\{z_j\}_{j=1}^d$ to compute Equation (3). We cannot compute $f_{Y(t)}(z|\mathbf{X})$ over an infinite number of values, so we instead accomplish this task by linearly interpolating between the estimates $\{\widehat{g}_h(z_j|\mathbf{X})\}_{j=1}^d$. Recall that the area under the curve (AUC) of such an interpolated density converges to the AUC of the true density on any finite interval at rate $O(1/d^2)$ by the trapezoidal rule (Cruz-Urbe and Neugebauer, 2003). We therefore set d to a large number (e.g., 500) and the grid $\{z_j\}_{j=1}^d$ to d equispaced points between the maximum and minimum values of $\{y_i(t)\}_{i=1}^n$ to ensure that the interpolation error is negligible.

3.4 Further Refinements

DDR improves upon the estimate $\widehat{g}_h(Y(t)|\mathbf{X})$ in Step 2 by enforcing known properties of a conditional density such as non-negativity and a total AUC of 1; these properties do not automatically hold in the finite sample setting as in unconditional density estimation with kernel density functions because a non-linear regressor may predict negative values. The improved estimate of $f(Y(t)|\mathbf{X})$ thus becomes:

$$\widehat{g}_h(Y(t)|\mathbf{X}) \leftarrow \frac{\max\{0, \widehat{g}_h(Y(t)|\mathbf{X})\}}{\mathcal{T}(\widehat{g}_h(Y(t)|\mathbf{X}))}, \quad (4)$$

where the denominator denotes the total AUC under $\max\{0, \widehat{g}_h(Y(t)|\mathbf{X})\}$ as estimated using the trapezoidal rule.

Step 3 of the DDR algorithm further modifies $\widehat{g}_h(Y(t)|\mathbf{X})$ using a parameter η by setting:

$$\widehat{g}_h(Y(t)|\mathbf{X}) \leftarrow \max\{0, \widehat{g}_h(Y(t)|\mathbf{X}) - \eta\}. \quad (5)$$

We choose the optimal η value again using Equation (3). The η value thus serves to eliminate low density regions of $\widehat{g}_h(Y(t)|\mathbf{X})$. Re-normalizing $\widehat{g}_h(Y(t)|\mathbf{X})$ using Equation (4) then elongates the high density regions. As a result, Equations (5) and (4) together “sharpen” $\widehat{g}_h(Y(t)|\mathbf{X})$. In practice, sharpening $\widehat{g}_h(Y(t)|\mathbf{X})$ substantially improves performance in the

high dimensional setting because it refines the initial conditional density estimates obtained from non-linear regression.

In summary, DDR estimates $f(Y(t)|\mathbf{X})$ by performing non-linear regression on a set of kernel density functions using an established regression procedure. The algorithm then refines the density estimate using Equations (4) and (5). DDR automatically tunes the associated hyperparameters $h, \boldsymbol{\lambda}$ and η with Equation (3). All of these improvements allow DDR to perform well even with small sample sizes and multiple dimensions.

4. Theory

We now establish the consistency of DDR as well as the bias of the density estimate as a function of the parameter h . We in particular show that the MISE loss in Equation (2) converges to zero in probability with bias rate $O(h^4)$.

The MISE loss is bounded above by two terms:

$$2 \int \int \left| \widehat{g}_h(z|\mathbf{x}) - g_h(z|\mathbf{x}) \right|^2 d\mathbb{P}_{\mathbf{X}}(\mathbf{x}) dz + 2 \int \int \left| g_h(z|\mathbf{x}) - f_{Y(t)}(z|\mathbf{x}) \right|^2 d\mathbb{P}_{\mathbf{X}}(\mathbf{x}) dz$$

The first term corresponds to the *variance* while the second corresponds to the squared *bias*. We require the following assumption for the variance:

Assumption 1 *The regression procedure is consistent for any fixed $h > 0$ and any fixed z :*

$$Q_n(z) \triangleq \int \left| \widehat{g}_h(z|\mathbf{x}) - g_h(z|\mathbf{x}) \right|^2 d\mathbb{P}_{\mathbf{X}}(\mathbf{x}) = o_p(1). \quad (6)$$

Note that $Q_n(z)$ is a function of $\widehat{g}_h(z|\mathbf{x})$ (a stochastic regression estimate that is a function of the random dataset), and therefore $Q_n(z)$ is itself stochastic. Many non-linear regression methods satisfy the above assumption from standard results in the literature. Kernel ridge regression (KRR) for example satisfies it under some mild conditions (Corollary 5 in (Smale and Zhou, 2007)). K-nearest neighbor, local polynomial, Rodeo and spectral regression also meet the requirement under their respective assumptions (Izbicki and Lee, 2017). Assumption 1 therefore allows the user to incorporate different regression procedures into DDR.

We next assume that the partial derivative of Q_n with respect to (w.r.t.) z is stochastically bounded:

Assumption 2 *$Q_n(z)$ is differentiable w.r.t. z and $\sup_z \left| \frac{\partial Q_n(z)}{\partial z} \right| = O_p(1)$.*

The differentiability of $Q_n(z)$ generally holds when we choose an asymptotically Dirac delta distribution that is everywhere differentiable such as with the aforementioned Gaussian, biweight or tricube densities. The stochastic bound is equivalent to requiring a finite supremum in the deterministic setting. The assumption therefore ensures a controlled derivative across all values of z .

Now consider the squared bias term:

$$\int \int \left| g_h(z|\mathbf{x}) - f_{Y(t)}(z|\mathbf{x}) \right|^2 d\mathbb{P}_{\mathbf{X}}(\mathbf{x}) dz. \quad (7)$$

The above quantity depends on the choice of the asymptotically Dirac delta distribution as well as the choice of h . We therefore impose some additional assumptions on the distribution:

Assumption 3 *The asymptotically Dirac delta distribution is a kernel density function that admits the form $\delta_h(z - Y(t)) = \frac{1}{h}K\left(\frac{z - Y(t)}{h}\right)$ such that $\int uK(u) du = 0$ and $\int u^2K(u) du < \infty$.*

The reader may recall an identical form used in unconditional kernel density estimation, where we also require a centered and finite variance kernel density function (Vaart, 1998). We finally impose a mild smoothness assumption on the density $f(Y(t)|\mathbf{X})$ which is similarly required in the unconditional case:

Assumption 4 *The conditional density $f_{Y(t)}(z|\mathbf{X})$ is twice continuously differentiable for any z and satisfies $\int \int |f''_{Y(t)}(z|\mathbf{x})|^2 d\mathbb{P}_{\mathbf{X}}(\mathbf{x}) dz < \infty$.*

We are now ready to state the main result:

Theorem 1 *Under Assumptions 1-4, we have:*

$$\int_a^b \int \left| \widehat{g}_h(z|\mathbf{x}) - f_{Y(t)}(z|\mathbf{x}) \right|^2 d\mathbb{P}_{\mathbf{X}}(\mathbf{x}) dz \leq o_p(1) + Ch^4,$$

for any $a < b$ where C is a constant that does not depend on n or h .

The proof is located in Appendix 7.1. Notice that the above bound holds over all possible finite intervals $[a, b]$ on \mathbb{R} . The parameter h also imposes bias at rate h^4 . We thus achieve consistency as $h \rightarrow 0^+$ as expected from the intuition summarized in Equation (1).

5. Experiments

5.1 Algorithms & Hyperparameters

We next compared DDR against six other algorithms:

1. LSCDE is a least squares procedure for directly estimating the conditional density (Sugiyama et al., 2010). This method was shown to outperform CKDE and MDN, but LSCDE in general only performs well with a few variables in \mathbf{X} .
2. Series Conditional Density Estimation (SCDE) improves upon LSCDE in the high dimensional setting by integrating direct conditional density estimation with dimensionality reduction (Izbicki and Lee, 2016).
3. FC estimates conditional densities by utilizing non-linear regression on an orthogonal series (Izbicki and Lee, 2017). FC admits a variety of regression procedures, so we equipped the algorithm with popular choices including Gaussian kernel ridge regression (FC-KRR), k-nearest neighbor (FC-NN) and the Lasso (FC-Lasso).
4. Normalizing Flow Estimator (NFE) combines a deep neural network with a normalizing flow, whereby a simple initial density is transformed in a more complex one by applying a sequence of invertible transformations (Trippe and Turner, 2018).

Recall that we can instantiate DDR with a variety of non-linear regressors like FC. Choosing the right regressor is however a non-trivial task. To prevent the user from cherry picking and optimize performance, we strategically instantiate DDR with KRR equipped with the first degree INK spline kernel, a universal regressor that works very well in the sample size regime of nearly all RCTs (tens to few thousands) (Vapnik, 2013; Izmailov et al., 2013). The INK spline kernel is much less well known than the Gaussian one, but INK spline is tuning free and therefore does not introduce an extra hyperparameter that can increase the probability of overfitting. DDR runs in $O(dn^3)$ time in this case, and KRR allows us to compute leave-one-out predictions in closed form via the Sherman-Morrison formula to speed up cross-validation (Seber and Lee, 2012). We selected the h and η parameters from 20 and 50 equispaced points between 0 and 0.5, respectively. We also selected the ridge penalty from $\{1E-1, 1E-2, 1E-3, 1E-4, 1E-5, 1E-6\}$.

We tuned the hyperparameters of the other algorithms according to the original authors' source codes with one exception. LSCDE uses the Kullback-Leibler divergence to optimize hyperparameters by default, but we found that this method consistently resulted in worse performance using our evaluation criteria. We therefore ensured that all methods optimized their hyperparameters according to the MISE loss in order to ensure a fair comparison.

5.2 Accuracy

Synthetic Data. We first evaluated the seven algorithms by simulating experimental datasets with different conditions. Most RCT datasets contain up to a few hundred samples, so we generated 200 samples of $Y(t)$ from the following models: (1) homoskedastic $\mathcal{N}(\mathbf{X}_1, 0.1)$, (2) heteroskedastic $\mathcal{N}(\mathbf{X}_1, 0.1(|\mathbf{X}_2| + 0.5))$, (3) bimodal $\frac{1}{2}\mathcal{N}(\mathbf{X}_1, 0.1) + \frac{1}{2}\mathcal{N}(\mathbf{X}_2, 0.1)$, and (4) skewed $\mathbf{X}_1 + \Gamma(k = 2, \theta = 0.4)$. We instantiated each of the 4 models 100 times. We also set the cardinality of \mathbf{X} to an element of $\{2, 4, 6, 8, 10, 15, 20\}$ by sampling i.i.d. from a standard Gaussian. We therefore compared the algorithms across a total of $4 \times 100 \times 7 = 2800$ datasets.

We evaluated the algorithms with the MISE loss in Equation (2) as averaged over the datasets. We summarize the results in Figure 2. DDR achieved the lowest average MISE scores across all four model types (Figure 2 (a)) and all variable numbers (Figure 2 (b)). Moreover, DDR outperformed FC-KRR to a significant degree ($t=-9.57$, $p<2.2E-16$). Every other pairwise comparison was also significant at a Bonferonni corrected threshold of 0.05/6 according to paired t-tests. FC-Lasso performed poorly, since the linear Lasso is generally not a consistent estimator of the conditional density. NFE performed the worst because deep learning usually requires large sample sizes inaccessible to most RCTs. We conclude that DDR achieves the best performance on average across a variety of situations.

Real Data. We next ran the algorithm on 40 real observational datasets downloaded from the UCI Machine Learning Repository with sample sizes ranging from 80 to 666. We downloaded as many pre-processed UCI datasets we could find with tens to hundreds of samples in tabular format. Predicting from observational data is equivalent to having one treatment value $T = 0$ and predicting $f(Y(0)|\mathbf{X})$. Note that we can only compute the MISE loss up to a constant with the real data because we do not have access to the ground truths. We therefore computed Equation (3) with 10-fold cross-validation and then averaged over the folds for each dataset. Since we cannot compare the loss values between datasets, we

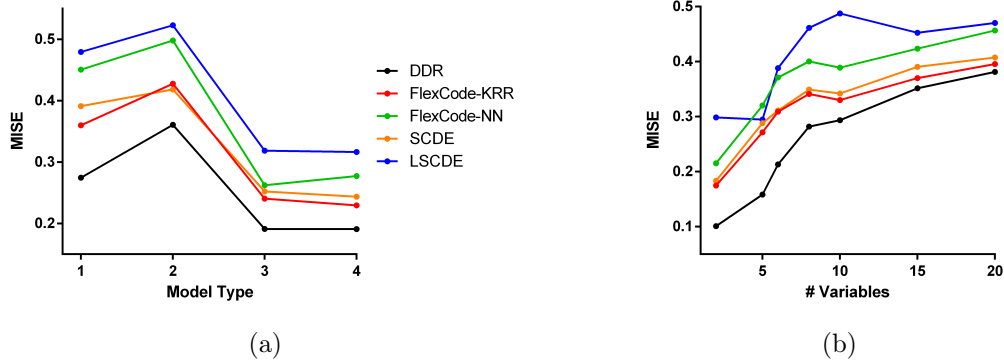


Figure 2: The average MISE loss values as a function of (a) model type and (b) number of variables. We have not included FlexCode-Lasso or NFE in the plots, since they performed substantially worse. DDR achieved the lowest average MISE in all cases.

ranked the values from 1 to 7 instead. We summarize the results in Table 2. DDR achieved the lowest average rank across all datasets as shown in the last row.

We repeated the above procedure with 7 clinical trial datasets downloaded from the National Institute of Drug and Alcohol (NIDA) Data Share. We computed the MISE rank for each value of T . DDR again achieved the lowest average rank as seen in Table 3. We conclude that the results seen with both the real observational and clinical trial datasets replicate the results seen with the synthetic data.

5.3 Detailed Clinical Trial Application

We now provide a detailed application of DDR to real RCT data in order to recover conditional densities summarizing treatment outcome. The discovery reported below is novel even in the medical literature. We investigated the effect of transdermal nicotine patches (TNPs) on long-term smoking cessation using the CTN-0009 dataset from the NIDA Data Share (Reid et al., 2008). In this trial, 166 subjects were randomized to receive either TNP or treatment-as-usual (TAU; motivational interviewing and supportive therapy) without TNP over a period of 8 weeks. Smoking increases carbon monoxide (CO) levels in the lungs, so the investigators objectively monitored smoking cessation by measuring CO levels with a breathalyzer.

TNPs only mildly increase smoking cessation after treatment ends (see Appendix 7.2). The small effect size may exist because only a minority of patients benefit from TNP. We in particular hypothesized that patients who smoke irregularly experience more intense nicotine cravings than those who smoke on a regular basis. Patients who smoke irregularly may want to stop smoking, but they struggle to remain in remission due to the cravings. As a result, these individuals will benefit more from TNP because TNP decreases the frequency and intensity of the craving episodes (Rose et al., 1985). We can evaluate this hypothesis using the RCT dataset, where we track the consistency in smoking using short-term changes in lung CO levels as shown on the x-axis in Figure 3 (a). The x-axis more specifically corresponds to the absolute value of the CO levels at baseline minus the CO levels at 2 weeks. The y-axis denotes the CO levels at 9 weeks after treatment ended. Based on the two different linear

	(rank)	DDR	FC-KRR	FC-NN	FC-Lasso	SCDE	LSCDE	NFE
Autism	5	-1.41E-01	-2.17E-01	-6.30E-01	-6.13E-01	<u>-6.38E-01</u>	-4.75E-02	+8.99E+00
AutoMPG	4	-2.47E-02	-1.85E-02	-3.18E-02	+2.10E-01	<u>-1.06E-01</u>	-2.74E-02	+2.23E+01
Breast	4	-9.05E-01	-1.19E+00	-4.03E+00	+6.37E+00	<u>-4.59E+00</u>	-9.80E-02	+1.89E+00
BuddyMove	1	<u>-2.23E-01</u>	-2.20E-01	-2.04E-01	-1.30E-01	-1.98E-01	-1.86E-01	+3.09E-01
Caesarian	2	-9.13E-03	-7.71E-03	<u>-1.45E-02</u>	+2.10E-02	+9.83E-03	-7.85E-03	+5.51E+01
CKD	4	-3.34E+00	-5.78E+00	<u>-7.26E+00</u>	-7.00E+00	-1.11E+00	-2.27E+00	-1.40E+00
Coimbra	1	<u>-7.15E-01</u>	-6.16E-01	-5.96E-01	+4.01E-01	-7.08E-01	-5.11E-01	-1.60E-01
Concrete	1	<u>-3.87E-02</u>	-3.36E-02	-2.56E-02	-1.83E-02	+6.43E-02	-3.53E-02	+1.06E+02
Credit	2	-8.27E-04	<u>-8.63E-04</u>	-7.41E-04	-2.18E-04	-4.14E-04	-6.77E-04	+8.57E+02
Cryotherapy	3	-5.42E-05	<u>-7.93E-05</u>	-6.04E-05	-8.30E-06	-3.18E-06	-3.82E-05	+3.02E+05
CSM	1	<u>-1.20E+00</u>	-1.10E+00	-1.06E+00	-7.57E-01	-1.06E+00	-1.12E+00	+8.37E-02
Dermatology	1	<u>-1.85E-02</u>	-1.49E-02	-1.20E-02	-4.52E-03	-1.33E-02	-1.38E-02	+1.11E+03
Echo	1	<u>-2.37E-02</u>	-2.35E-02	-2.20E-02	-9.02E-03	-1.19E-02	-2.32E-02	+1.72E+01
EColi	2	-4.46E+00	-4.07E+00	-3.32E+00	-1.65E+00	-3.24E+00	-2.10E+00	<u>-2.16E+01</u>
Facebook	3	-7.08E-03	-6.31E-03	-2.94E-03	-7.94E-03	<u>-1.87E-02</u>	-1.99E-03	+2.86E+03
Fertility	1	<u>-1.09E-07</u>	-8.35E-08	-7.02E-08	-1.43E-08	-1.51E-08	-6.40E-08	+1.96E+07
ForestFires	4	-1.42E-02	-2.21E-02	<u>-2.63E-02</u>	+3.18E-03	+5.13E-03	-1.69E-02	+8.57E+01
ForestTypes	1	<u>-2.40E+01</u>	-2.04E+01	-1.75E+01	-8.64E+00	-1.43E+01	-1.75E+01	-1.36E+00
Glass	5	-6.60E-02	-7.25E-01	<u>-8.92E-01</u>	-2.40E-01	-8.00E-01	-3.63E-02	+1.90E+00
GPS	2	-9.93E-03	-8.38E-03	-8.54E-03	-9.31E-03	<u>-1.00E-02</u>	-8.75E-03	+3.23E+01
Hayes	4	-3.14E-03	-3.27E-03	-3.57E-03	<u>-3.67E-03</u>	-2.88E-03	-2.53E-03	+2.43E+01
HCC	2	-5.51E-01	-5.04E-01	<u>-5.54E-01</u>	-1.60E-01	+1.67E-01	-4.07E-01	+1.65E+00
Hepatitis	2	-6.06E-02	-4.12E-02	-3.73E-02	+1.61E-03	-2.83E-02	<u>-8.52E-02</u>	+4.99E+01
Immunotherapy	2	-1.89E-01	-1.77E-01	-1.64E-01	-1.11E-01	<u>-2.37E-01</u>	-1.53E-01	+4.83E-01
Inflammation	2	-1.45E-01	-1.08E-01	-9.99E-02	-2.48E-02	<u>-1.86E-01</u>	-9.66E-02	+8.33E+00
Iris	1	<u>-2.26E+02</u>	-1.83E+02	-1.57E+02	-2.67E+01	-1.48E+02	-1.95E+02	-2.39E+00
Istanbul	1	<u>-1.70E-02</u>	-1.55E-02	-1.56E-02	-1.00E-02	-6.05E-03	-1.48E-02	+4.06E+00
Leaf	1	<u>-1.99E-02</u>	-1.90E-02	-1.24E-02	-4.70E-03	-3.53E-03	-1.13E-02	+1.36E+02
Parkinson's	1	<u>-6.60E-02</u>	-5.80E-02	-5.11E-02	+8.88E-02	+7.88E-02	-5.22E-02	+3.29E+01
Planning	1	<u>-1.20E-02</u>	-1.05E-02	-1.09E-02	-9.80E-03	-7.10E-03	-1.04E-02	+7.45E+01
Seeds	3	-6.32E-02	-7.28E-02	<u>-1.02E-01</u>	+3.38E-02	+5.63E-02	-5.30E-02	+9.99E+00
Somerville	1	<u>-1.48E-01</u>	-1.27E-01	-9.58E-02	+1.39E-02	-9.68E-02	-6.33E-02	+2.29E+02
SPECTF	4	-3.71E-01	-4.89E-01	-5.14E-01	-3.66E-02	<u>-6.98E-01</u>	-3.01E-01	+1.14E+00
Statlog	1	<u>-3.50E-02</u>	-3.11E-02	-2.41E-02	-2.54E-02	+2.05E-01	-2.19E-02	+2.25E+01
Thoracic	1	<u>-2.28E-02</u>	-2.24E-02	-2.14E-02	-2.13E-02	+4.03E-02	-1.89E-02	+5.47E+01
Traffic	2	-1.27E-01	-1.18E-01	-9.61E-02	-5.74E-02	<u>-1.38E-01</u>	-6.45E-02	+3.22E+01
User	2	-5.45E-01	-5.21E-01	-3.47E-01	-2.90E-01	-5.03E-01	<u>-9.38E-01</u>	+3.01E+00
Wine	1	<u>-4.70E-03</u>	-4.69E-03	-4.65E-03	-2.54E-03	-2.76E-04	-3.22E-03	+2.60E+02
Wisconsin	5	-2.80E+00	<u>-7.95E+00</u>	-7.03E+00	-6.54E+00	-5.93E+00	-7.63E-01	+1.68E+00
Yacht	1	<u>-5.12E+00</u>	-5.08E+00	-3.94E+00	-2.18E+00	+4.16E+00	-3.12E+00	-1.46E-02
Avg Rank		2.15	2.70	3.05	5.10	4.15	4.10	6.75

Table 2: Results from 40 real observational datasets. Lower MISE loss values (up to a constant) denote better performance. Underlined values correspond to the best in each row. The second column lists the rank of DDR relative to the other algorithms. DDR performs the best on average by achieving the lowest average rank as shown in the last row.

	T	(rank)	DDR	FC-KRR	FC-NN	FC-Lasso	SCDE	LSCDE	NFE
CSP-1019	0	2	-9.69E-02	-8.24E-02	-9.50E-02	-8.49E-02	<u>-1.00E-01</u>	-8.64E-02	+4.59E-02
	1	1	<u>-9.91E-02</u>	-9.77E-02	-9.37E-02	-4.92E-02	-6.28E-02	-9.51E-02	+1.47E-01
CSP-1021	0	1	<u>-1.09E-01</u>	-8.29E-02	-9.78E-02	+1.52E-03	+3.08E-01	-7.35E-02	+2.10E+00
	1	2	-1.06E-01	-7.76E-02	-9.40E-02	-7.68E-02	-6.78E-02	<u>-1.10E-01</u>	+1.52E+00
CTN-0009	0	2	-1.66E-02	-1.39E-02	-1.21E-02	-1.93E-04	-2.36E-03	<u>-1.83E-02</u>	+1.78E+01
	1	2	-1.82E-02	-1.79E-02	-1.80E-02	-9.56E-03	-1.26E-02	<u>-1.94E-02</u>	+7.50E+00
CTN-0010	0	1	<u>-1.01E-02</u>	-5.93E-03	-1.00E-02	+2.72E-03	+2.96E-02	-3.80E-03	+4.65E+02
	1	1	<u>-1.37E-02</u>	+3.95E-03	-1.26E-02	+2.37E-02	-2.72E-03	-1.13E-02	+4.72E+01
CTO-0012	0	4	-7.08E-02	-7.52E-02	<u>-8.89E-02</u>	-2.31E-02	3.91E-02	-7.20E-02	+2.10E+01
	1	1	<u>-9.96E-02</u>	-8.97E-02	-7.60E-02	+1.38E-03	+1.72E-01	-4.19E-02	+9.58E+00
CSP-1025	0	3	-4.36E-01	-1.39E-01	-2.74E-01	-8.64E-01	<u>-1.13E+00</u>	-1.34E-01	+5.40E+00
	1	4	-2.46E-01	-2.86E-01	-1.49E-01	<u>-7.16E-01</u>	-5.63E-01	-6.96E-02	+1.69E+00
CTN-0051	0	4	-6.49E-02	-1.36E-01	-9.57E-02	<u>-1.42E-01</u>	9.43E-03	-3.72E-02	+2.83E+01
	1	4	-9.38E-02	-2.11E-01	-2.03E-01	<u>-2.36E-01</u>	-5.85E-02	-3.97E-02	+1.00E+01
Avg Rank			2.29	3.29	3.00	4.21	4.57	3.64	7.00

Table 3: Results from 7 real clinical trials presented in a similar format as Table 2. DDR again achieves the lowest average rank.

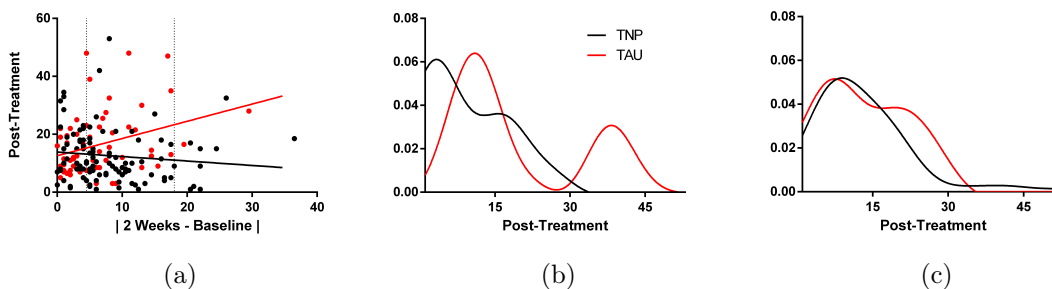


Figure 3: Detailed analysis of a clinical trial dataset for personalized outcomes of TNP. (a) We can summarize the personalized outcomes of TNP relative to TAU using two linear regression slopes. (b) The CO levels of this particular patient changed by 18 ppm, so DDR predicts that TNP is more beneficial to this patient than TAU. (c) On the other hand, CO levels only changed by 4.5 ppm in this patient, so DDR could not clearly differentiate the outcomes of the two treatments.

regression slopes, we can see that a change in CO levels while on TNP generally has no effect on post-treatment CO levels but a change in CO levels while on TAU has a detrimental effect.² We confirmed the significance of the observed trend by rejecting the null of equality in slopes between TNP and TAU ($z=-2.87$, one sided $p=0.002$). We therefore conclude that TNP reduces post-treatment CO levels among patients with large changes in short term CO levels.

The regression slopes however only provide point estimates of predicted treatment outcome. In reality, even patients with large changes may not benefit from TNP because the post-treatment CO level is stochastic. DDR allows us to visualize this uncertainty by recovering conditional densities for all possible outcome values. Consider for example patient A with a value of 18 on the x-axis of Figure 3 (a); the regression slopes suggest that this patient is expected to obtain a post-treatment CO level of around 11 if treated with TNP but a level of 23 if treated with TAU. The densities recovered by DDR in Figure 3 (b) for patient A are also significantly different (one sided $p=0.013$; see Appendix 7.3), but they imply a more modest effect because the patient also has a high probability of *not* experiencing such a large difference in post-treatment CO levels. On the other hand, patient B with a value of 4.5 on the x-axis is predicted to respond equally well to TNP and TAU in Figure 3 (a); the patient may therefore decide not to take TNP based on the regression estimates. However, DDR recovers broad densities as shown in Figure 3 (c); these densities imply that we cannot differentiate between the two treatments with the available evidence, so patient B may actually benefit more from TNP than TAU. This patient may therefore decide to try TNP after seeing the output of DDR. We conclude that the conditional densities recovered by DDR help patients make more informed treatment decisions by allowing them to visualize the probabilities associated with all possible outcome values in an intuitive fashion.

6. Conclusion

We proposed DDR for estimating conditional densities by performing non-linear regression over a set of kernel density functions. DDR differs from CKDE by directly estimating the conditional density, as opposed to estimating the joint and marginal densities first. DDR outperforms previous methods on average across a variety of synthetic and real datasets. The algorithm also generates patient-specific densities of treatment outcomes when run on RCT data; as opposed to the conditional expectation, the conditional density recovered by DDR allows patients and healthcare providers to easily compare the probabilities associated with *all* possible outcome values by effectively converting a standard RCT into a personalized one. Theoretical results further support our empirical claims by highlighting the consistency of DDR as well as quantifying the rate of bias with respect to the smoothing parameter h . We ultimately believe that this work is an important contribution to the literature because it introduces a state of the art conditional density estimation method as well as demonstrates a non-trivial application to a key problem in medicine.

2. Non-linear polynomial regression produced essentially the same conditional expectation estimates.

References

- Anthony K Akobeng. Understanding diagnostic tests 2: likelihood ratios, pre-and post-test probabilities and their use in clinical practice. *Acta Paediatrica*, 96(4):487–491, 2007.
- Zackary Berger. Navigating the unknown: shared decision-making in the face of uncertainty. *Journal of General Internal Medicine*, 30(5):675–678, 2015.
- Viraj Bhise, Ashley ND Meyer, Shailaja Menon, Geeta Singhal, Richard L Street, Traber D Giardina, and Hardeep Singh. Patient perspectives on how physicians communicate diagnostic uncertainty: an experimental vignette study. *International Journal for Quality in Health Care*, 30(1):2–8, 2018.
- Christopher M Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
- Matias D Cattaneo, Michael Jansson, and Xinwei Ma. Simple local polynomial density estimators. *arXiv preprint arXiv:1811.11512*, 2018.
- David Cruz-Uribe and CJ Neugebauer. An elementary proof of error estimates for the trapezoidal rule. *Mathematics Magazine*, 76(4):303–306, 2003.
- George A Diamond and James S Forrester. Analysis of probability as an aid in the clinical diagnosis of coronary-artery disease. *New England Journal of Medicine*, 300(24):1350–1358, 1979.
- Mogens Fosgerau and Daisuke Fukuda. Valuing travel time variability: Characteristics of the travel time distribution on an urban road. MPRA Paper 24330, University Library of Munich, Germany, 2010. URL <https://ideas.repec.org/p/pramprapa/24330.html>.
- MJ Fryer. Some errors associated with the non-parametric estimation of density functions. *IMA Journal of Applied Mathematics*, 18(3):371–380, 1976.
- Rob J. Hyndman, David M. Bashtannyk, and Gary K. Grunwald. Estimating and visualizing conditional densities. *Journal of Computational and Graphical Statistics*, 5(4):315–336, 1996. doi: 10.1080/10618600.1996.10474715. URL <https://amstat.tandfonline.com/doi/abs/10.1080/10618600.1996.10474715>.
- Rafael Izbicki and Ann Lee. Converting high-dimensional regression to high-dimensional conditional density estimation. *Electron. J. Statist.*, 11(2):2800–2831, 2017. doi: 10.1214/17-EJS1302. URL <https://doi.org/10.1214/17-EJS1302>.
- Rafael Izbicki and Ann B. Lee. Nonparametric conditional density estimation in a high-dimensional regression setting. *Journal of Computational and Graphical Statistics*, 25(4):1297–1316, 2016. doi: 10.1080/10618600.2015.1094393. URL <https://doi.org/10.1080/10618600.2015.1094393>.
- Rauf Izmailov, Vladimir Vapnik, and Akshay Vashist. Multidimensional splines with infinite number of knots as svm kernels. In *The 2013 International Joint Conference on Neural Networks (IJCNN)*, pages 1–7. IEEE, 2013.

- Takafumi Kanamori, Taiji Suzuki, and Masashi Sugiyama. Statistical analysis of kernel-based least-squares density-ratio estimation. *Machine Learning*, 86(3):335–367, 2012.
- Ram P Kanwal. *Generalized Functions: Theory and Applications*. Springer Science & Business Media, 2011.
- Alexander R. Luedtke and Mark J. van der Laan. Statistical inference for the mean outcome under a possibly non-unique optimal treatment strategy. *Ann. Statist.*, 44(2):713–742, 04 2016. doi: 10.1214/15-AOS1384. URL <https://doi.org/10.1214/15-AOS1384>.
- Hui Chin Mah, Leelavathi Muthupalaniappen, and Wei Wen Chong. Perceived involvement and preferences in shared decision-making among patients with hypertension. *Family Practice*, 33(3):296–301, 2016.
- Stephen Martinkovich, Darshan Shah, Sonia Lobo Planey, and John A Arnott. Selective estrogen receptor modulators: tissue specificity and clinical utility. *Clinical Interventions in Aging*, 9:1437, 2014.
- Whitney K Newey and Daniel McFadden. Large sample estimation and hypothesis testing. *Handbook of Econometrics*, 4:2111–2245, 1994.
- M. S. Reid, B. Fallon, S. Sonne, F. Flammino, E. V. Nunes, H. Jiang, E. Kourniotis, J. Lima, R. Brady, C. Burgess, C. Arfken, E. Pihlgren, L. Giordano, A. Starosta, J. Robinson, and J. Rotrosen. Smoking cessation treatment in community-based substance abuse rehabilitation programs. *J Subst Abuse Treat*, 35(1):68–77, Jul 2008.
- Jed E Rose, Joseph E Herskovic, Yvonne Trilling, and Murray E Jarvik. Transdermal nicotine reduces cigarette craving and nicotine preference. *Clinical Pharmacology & Therapeutics*, 38(4):450–456, 1985.
- Paul R Rosenbaum and Donald B Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.
- M Rosenblatt. Conditional probability density and regression estimators. *International Symposium on Multivariate Analysis, Multivariate analysis II Proceedings*, 1969.
- Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688, 1974.
- Donald B Rubin. [on the application of probability theory to agricultural experiments. essay on principles. section 9.] comment: Neyman (1923) and causal inference in experiments and observational studies. *Statistical Science*, 5(4):472–480, 1990.
- G.A.F. Seber and A.J. Lee. *Linear Regression Analysis*. Wiley Series in Probability and Statistics. Wiley, 2012. ISBN 9781118274422. URL <https://books.google.com/books?id=X2Y60kX18ysC>.
- Motoki Shiga, Voot Tangkaratt, and Masashi Sugiyama. Direct conditional probability density estimation with sparse feature selection. *Machine Learning*, 100(2):161–182, Sep 2015. ISSN 1573-0565. doi: 10.1007/s10994-014-5472-x. URL <https://doi.org/10.1007/s10994-014-5472-x>.

- Steve Smale and Ding-Xuan Zhou. Learning theory estimates via integral operators and their approximations. *Constructive Approximation*, 26(2):153–172, Aug 2007. ISSN 1432-0940. doi: 10.1007/s00365-006-0659-y. URL <https://doi.org/10.1007/s00365-006-0659-y>.
- Masashi Sugiyama, Ichiro Takeuchi, Taiji Suzuki, Takafumi Kanamori, Hirotaka Hachiya, and Daisuke Okanohara. Conditional density estimation via least-squares density ratio estimation. In Yee Whye Teh and Mike Titterton, editors, *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pages 781–788, Chia Laguna Resort, Sardinia, Italy, 13–15 May 2010. PMLR. URL <http://proceedings.mlr.press/v9/sugiyama10a.html>.
- Ichiro Takeuchi, Quoc V. Le, Timothy D. Sears, and Alexander J. Smola. Nonparametric quantile estimation. *J. Mach. Learn. Res.*, 7:1231–1264, December 2006. ISSN 1532-4435. URL <http://dl.acm.org/citation.cfm?id=1248547.1248592>.
- Voot Tangkaratt, Ning Xie, and Masashi Sugiyama. Conditional density estimation with dimensionality reduction via squared-loss conditional entropy minimization. *Neural Computation*, 27:228–254, 2015.
- Volker Tresp. Mixtures of gaussian processes. In *Advances in Neural Information Processing Systems*, pages 654–660, 2001.
- Brian L Trippe and Richard E Turner. Conditional density estimation with bayesian normalising flows. *arXiv preprint arXiv:1802.04908*, 2018.
- A. W. van der Vaart. *Asymptotic Statistics*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 1998. doi: 10.1017/CBO9780511802256.
- Vladimir Vapnik. *The Nature of Statistical Learning Theory*. Springer science & business media, 2013.
- Matt P Wand and M Chris Jones. *Kernel smoothing*. Chapman and Hall/CRC, 1994.
- Baqun Zhang, Anastasios A Tsiatis, Marie Davidian, Min Zhang, and Eric Laber. Estimating optimal treatment regimes from a classification perspective. *Stat*, 1(1):103–114, 2012a.
- Baqun Zhang, Anastasios A Tsiatis, Eric B Laber, and Marie Davidian. A robust method for estimating optimal treatment regimes. *Biometrics*, 68(4):1010–1018, 2012b.

7. Appendix

7.1 Proofs

Definition 1 (*Stochastic equicontinuity w.r.t. z*) For every $\varepsilon, \delta > 0$, there exists a sequence of random variables Υ_n and an integer N such that $\forall n \geq N$, we have $\mathbb{P}(|\Upsilon_n| > \varepsilon) < \delta$. Moreover, for each z , there is an open set \mathcal{N} containing z with:

$$\sup_{z' \in \mathcal{N}} |Q_n(z) - Q_n(z')| \leq \Upsilon_n, \quad n \geq N.$$

Notice that Υ_n acts like a random epsilon by bounding changes in $Q_n(z)$ w.r.t. z .

Lemma 1 (*Lipschitz continuity \implies stochastic equicontinuity; Lemma 2.9 in (Newey and McFadden, 1994)*) If $Q_n(z) = o_p(1)$ for all $z \in [a, b]$ and $B_n = O_p(1)$ such that for all $z, z' \in [a, b]$ we have $|Q_n(z) - Q_n(z')| \leq B_n|z - z'|$, then $Q_n(z)$ is stochastically equicontinuous w.r.t. z .

Lemma 2 If Assumption 2 holds, then $Q_n(z)$ is stochastically equicontinuous w.r.t. z .

Proof Because $Q_n(z)$ is differentiable on $[a, b]$ w.r.t. z , we can write:

$$|Q_n(z) - Q_n(z')| \leq \sup_{z \in [a, b]} \left| \frac{\partial Q_n(z)}{\partial z} \right| |z - z'|.$$

We have $\sup_{z \in [a, b]} \left| \frac{\partial Q_n(z)}{\partial z} \right| = O_p(1)$ by Assumption 2. Invoke Lemma 1 to conclude that Q_n is stochastically equicontinuous w.r.t. z . \blacksquare

Lemma 3 (*Stochastic equicontinuity + pointwise consistency \iff uniform consistency; Lemma 2.8 in (Newey and McFadden, 1994)*) We have $Q_n(z) = o_p(1)$ for all $z \in [a, b]$ and $Q_n(z)$ is stochastically equicontinuous w.r.t z if and only if $\sup_{z \in [a, b]} Q_n(z) = o_p(1)$.

Lemma 4 (*Uniform consistency \implies proper integral consistency*) If we have $\sup_{z \in [a, b]} Q_n(z) = o_p(1)$, then $\int_a^b Q_n(z) dz = o_p(1)$.

Proof Choose $\varepsilon > 0$. Then write:

$$\begin{aligned} & \mathbb{P}\left(\left|\int_a^b Q_n(z) dz\right| \geq \varepsilon\right) \\ & \leq \mathbb{P}\left(\int_a^b |Q_n(z)| dz \geq \varepsilon\right) \\ & \leq \mathbb{P}\left(\int_a^b \sup_{z \in [a, b]} |Q_n(z)| dz \geq \varepsilon\right) \tag{8} \\ & = \mathbb{P}\left((b-a) \sup_{z \in [a, b]} |Q_n(z)| \geq \varepsilon\right) \\ & = \mathbb{P}\left(\sup_{z \in [a, b]} |Q_n(z)| \geq \frac{\varepsilon}{b-a}\right). \end{aligned}$$

Choose $\delta > 0$. By assumption, for $\frac{\varepsilon}{b-a}$ and δ , $\exists N \in \mathbb{N}^+$ such that $\forall n \geq N$, we have:

$$\mathbb{P}\left(\sup_{z \in [a, b]} |Q_n(z)| \geq \frac{\varepsilon}{b-a}\right) \leq \delta.$$

Note that we chose ε and δ arbitrarily. The conclusion follows by the epsilon-delta definition of convergence in probability to zero. \blacksquare

Theorem 1 Under Assumptions 1-4, we have:

$$\int_a^b \int \left| \widehat{g}_h(z|\mathbf{x}) - f_{Y(t)}(z|\mathbf{x}) \right|^2 d\mathbb{P}_{\mathbf{X}}(\mathbf{x}) dz \leq o_p(1) + Ch^4,$$

for any $a < b$ where C is a constant that does not depend on n or h .

Proof We write:

$$\begin{aligned} & \int_a^b \int \left| \widehat{g}_h(z|\mathbf{x}) - f_{Y(t)}(z|\mathbf{x}) \right|^2 d\mathbb{P}_{\mathbf{X}}(\mathbf{x}) dz \\ & \leq 2 \int_a^b \int \left| \widehat{g}_h(z|\mathbf{x}) - g_h(z|\mathbf{x}) \right|^2 d\mathbb{P}_{\mathbf{X}}(\mathbf{x}) dz \\ & \quad + 2 \int_a^b \int \left| g_h(z|\mathbf{x}) - f_{Y(t)}(z|\mathbf{x}) \right|^2 d\mathbb{P}_{\mathbf{X}}(\mathbf{x}) dz \end{aligned} \quad (9)$$

The term $Q_n(z) = \int \left| \widehat{g}_h(z|\mathbf{x}) - g_h(z|\mathbf{x}) \right|^2 d\mathbb{P}_{\mathbf{X}}(\mathbf{x})$ is $o_p(1)$ for each $z \in [a, b]$ by Assumption 1. Invoke Lemmas 3, 4 and then 5 to conclude that we have:

$$\int_a^b \int \left| \widehat{g}_h(z|\mathbf{x}) - g_h(z|\mathbf{x}) \right|^2 d\mathbb{P}_{\mathbf{X}}(\mathbf{x}) dz = o_p(1).$$

We now focus on the term $\left| g_h(z|\mathbf{x}) - f_{Y(t)}(z|\mathbf{x}) \right|^2$. We write:

$$\begin{aligned} & g_h(z|\mathbf{x}) - f_{Y(t)}(z|\mathbf{x}) \\ & = \int \frac{1}{h} K\left(\frac{z-k}{h}\right) f_{Y(t)}(k|\mathbf{x}) dk - f_{Y(t)}(z|\mathbf{x}) \\ & = \int K(u) \left(f_{Y(t)}(z - hu|\mathbf{x}) - f_{Y(t)}(z|\mathbf{x}) \right) du. \end{aligned} \quad (10)$$

We then utilize a Taylorian expansion with a Laplacian representation of the remainder:

$$\begin{aligned} & f_{Y(t)}(z + h|\mathbf{x}) - f_{Y(t)}(z|\mathbf{x}) \\ & = hf'_{Y(t)}(z|\mathbf{x}) + h^2 \int_0^1 f''_{Y(t)}(z + sh|\mathbf{x})(1-s) ds. \end{aligned}$$

Substituting the above formula into Equation (10), we get:

$$\begin{aligned} & g_h(z|\mathbf{x}) - f_{Y(t)}(z|\mathbf{x}) \\ & = \int \int_0^1 K(u) [-hu f'_{Y(t)}(z|\mathbf{x}) + h^2 u^2 f''_{Y(t)}(z - shu|\mathbf{x})(1-s)] ds du \\ & = \int \int_0^1 h^2 u K(u) \left(u f''_{Y(t)}(z - shu|\mathbf{x})(1-s) \right) ds du, \end{aligned} \quad (11)$$

where the second equality follows because we assumed that K has expectation zero. We next utilize the Cauchy-Schwartz inequality $(\mathbb{E}AB)^2 \leq \mathbb{E}A^2 \mathbb{E}B^2$ with $A = U$ and $B =$

$Uf''(z - ShU|\mathbf{x})(1 - S)$; here U has density K and S is uniform on $[0, 1]$ as well as independent of U . The bottom of Equation (11) squared is therefore upper bounded by:

$$\begin{aligned} & h^4 \int K(u)u^2 du \int \int_0^1 K(u)u^2 f''_{Y(t)}(z - shu|\mathbf{x})^2 (1 - s)^2 ds du \\ &= h^4 \int K(u)u^2 du \int \int_0^1 K(u)u^2 f''_{Y(t)}(z|\mathbf{x})^2 (1 - s)^2 ds du \\ &= h^4 \left(\int K(u)u^2 du \right)^2 f''_{Y(t)}(z|\mathbf{x})^2 \frac{1}{3}. \end{aligned}$$

Integrating this with respect to $\mathbb{P}_{\mathbf{X}}(\mathbf{x})$ and z , we obtain:

$$\begin{aligned} & \int_a^b \int \left(g_h(z|\mathbf{x}) - f_{Y(t)}(z|\mathbf{x}) \right)^2 d\mathbb{P}_{\mathbf{X}}(\mathbf{x}) dz \\ & \leq h^4 \left(\int K(u)u^2 du \right)^2 \frac{1}{3} \int_a^b \int f''_{Y(t)}(z|\mathbf{x})^2 d\mathbb{P}_{\mathbf{X}}(\mathbf{x}) dz. \end{aligned}$$

We finally utilize the bound in Equation (9) to conclude that:

$$\int_a^b \int \left| g_h(z|\mathbf{x}) - f_{Y(t)}(z|\mathbf{x}) \right|^2 d\mathbb{P}_{\mathbf{X}}(\mathbf{x}) dz \leq o_p(1) + Ch^4,$$

where $C = \left(\int K(u)u^2 du \right)^2 \frac{1}{3} \int_a^b \int f''_{Y(t)}(z|\mathbf{x})^2 d\mathbb{P}_{\mathbf{X}}(\mathbf{x}) dz$. ■

7.2 Extra Clinical Trial Results

TNPs only mildly increase smoking cessation after treatment ends. We can verify this claim by plotting the estimated unconditional densities of CO levels measured 9 weeks after completion of TNP or TAU treatment (kernel density estimation with Gaussian kernel function and unbiased cross-validation; Figure 4). Notice that the TNP density is shifted to the left relative to the TAU density, indicating that patients treated with TNP eventually smoke less than those treated with TAU on average. The difference is small at 3.97 ppm but large enough to reject the null of equality in means using a t-test ($t=-2.369$, $p=0.020$).

7.3 Hypothesis Testing

The TNP density in Figure 3 (b) places higher probability at lower post-treatment CO levels than the TAU density. DDR may nevertheless recover such densities often when the TNP density does *not* place higher probability at lower post-treatment CO levels at the population level. We therefore also seek to reject the following null hypothesis:

$$H_0 : \sup_z \left[F_{Y(1)}(z|\mathbf{x}) - F_{Y(0)}(z|\mathbf{x}) \right] \leq 0,$$

where $F_{Y(t)}(z|\mathbf{X}) = \int_{-\infty}^z f_{Y(t)}(u|\mathbf{X}) du$. Notice that larger values of the above difference correspond to concentration of probability at lower values of z . We therefore reject the null

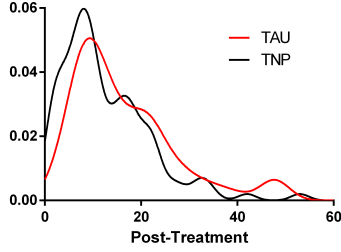


Figure 4: The unconditional densities estimated by standard kernel density estimation. The TNP density is shifted slightly to the left relative to the TAU density.

when the difference is large because lower values of z correspond to reduced post-treatment CO levels. We implement the hypothesis test by permuting the treatment labels, running DDR and then computing the following conditional statistic:

$$\mathcal{S} = \sup_z \left[\widehat{F}_{h,Y(1)}(z|\mathbf{x}) - \widehat{F}_{h,Y(0)}(z|\mathbf{x}) \right],$$

where $\widehat{F}_{h,Y(t)}(z|\mathbf{X}) = \int_{-\infty}^z \widehat{g}_{h,Y(t)}(u|\mathbf{X}) du$. We obtained a p-value of 0.0125 with 2000 permutations for patient A by setting $T = 0$ to TAU and $T = 1$ to TNP. We thus reject H_0 in this case and conclude that the TNP density places higher probability at lower post-treatment CO levels than the TAU density. Repeating the same process with patient B on the other hand led to a p-value of 0.3635.