
Learning Overlapping Representations for the Estimation of Individualized Treatment Effects

Yao Zhang
University of Cambridge

Alexis Bellot
University of Cambridge
The Alan Turing Institute

Mihaela van der Schaar
University of Cambridge, UCLA
The Alan Turing Institute

Abstract

The choice of making an intervention depends on its potential benefit or harm in comparison to alternatives. Estimating the likely outcome of alternatives from observational data is a challenging problem as *all* outcomes are never observed, and selection bias precludes the direct comparison of differently intervened groups. Despite their empirical success, we show that algorithms that learn domain-invariant representations of inputs (on which to make predictions) are often inappropriate, and develop generalization bounds that demonstrate the dependence on domain overlap and highlight the need for invertible latent maps. Based on these results, we develop a deep kernel regression algorithm and posterior regularization framework that substantially outperforms the state-of-the-art on a variety of benchmarks data sets.

1 Introduction

Counterfactual estimation poses the question of what would have been the outcome, if a different intervention had been applied. In order to make decisions in complex domains, making predictions on the causal effects of different actions and how these may vary across individuals is critical. In this paper we focus on the problem of making these predictions based on observational data, which is increasingly available in many domains such as medicine, public policy and advertising. In this setting, past actions, outcomes and context are available, but not knowledge of the treatment assignment policy - we do not know why a given individual was intervened

on or not. The treatment assignment mechanism will often be causally affected by context variables that also causally influence the outcome. As an example, motivated unemployed individuals are more likely to both take advantage of government training programs *and* find a new job soon.

Learning from observational data requires adjusting for the covariate shift that exists between groups of individuals that are observed to have received different interventions. The challenge is how to untangle confounding factors and make valid counterfactual predictions - the response if a different treatment had been applied. Recent methods have predominantly focused on learning representations regularized to balance these confounding factors by enforcing domain invariance with distributional distances [Johansson et al. \(2016, 2018\)](#); [Yao et al. \(2018\)](#). In this paper we argue that domain invariance is often too strict a requirement; overlapping support is sufficient for identifiability of the causal effect and equality in densities is not necessary. We interpret the loss in predictive power of domain-invariant representations as loss of information in the input variables that causally influence the treatment assignment, which is also often highly predictive of the treatment effect. Consider the example above for illustration: it is precisely because motivation is predictive of job outcomes that it confounds the treatment effect.

We introduce an optimization framework based on regularizing posterior distributions of the treatment effect that includes existing representation learning algorithms for different choices of regularization terms. We take advantage of this framework to introduce a novel type of regularization criteria for the problem of treatment effect estimation: the posterior counterfactual variance for enforcing domain overlap, and invertible representations to preserve the information content of the underlying context. Such an objective enjoys much better generalization in small sample regimes, smoother representation surfaces with respect to the outcomes, as can be seen in [Figure 1](#), and a Bayesian treatment of parameters which allows consistent un-




Figure 1: T-SNE visualizations of the learned embeddings for the control potential outcomes $Y(0)$ of the IHDP dataset. Each panel shows representations regularized by different criteria and the colored heatmap represents different outcome magnitudes with different colors. The **left panel** shows representations regularized by the Wasserstein distributional distance and results in poor discrimination. The **middle panel** shows representations optimized only for the factual data with the Gaussian likelihood. The **right panel** shows representations regularized by the counterfactual variance, our proposed criteria. Much better separation in outcomes is obtained by regularizing for the predictive variance, in contrast to using integral probability metrics such as the Wasserstein distance. Similar plots for a comparison between the control and treated group, and the treated potential outcome $Y(1)$ can be found in Appendix 5.3.

certainty estimation in predictions. We summarize our contributions as follows. We develop a theory that justifies regularizing for the posterior variance to improve generalization error and establish the limitations of distributional distances. We propose to use deep kernels and posterior regularization as a general framework to learn individualized treatment effects for arbitrary regularization terms. We provide an instantiation of this algorithm informed by novel generalization bounds that substantially improves upon the performance of state-of-the-art prediction algorithms.

1.1 Related work

Due to the ability of deep neural nets to learn rich representations, recent advances in predicting individualized treatment effects have focused on learning representations invariant to the treatment assignment policy that achieve a small error on the factual data. The hope is that the learnt representation and prediction function can generalize to prediction of counterfactual outcomes. Several methods follow this approach. [Johansson et al. \(2016\)](#) proposed learning a representation of the data that makes the treated and control distributions more similar, fitting a linear ridge-regression model on top of it. [Shalit et al. \(2017\)](#) built on their approach to derive a more flexible family of algorithms including non-linear hypotheses. However, both algorithms insist on quantifying divergence between treated and control groups with integral probability metrics.

In this work, we share the need for good representations

but argue for enforcing support overlap rather than equality in densities. [Yao et al. \(2018\)](#), inspired by nearest-neighbour methods, learn representation that preserve local similarity information in feature space and were able to show a decrease in the generalization error in counterfactual estimation.

Adapting Bayesian algorithms for the problem of individualized treatment effects has attracted a lot of interest, in particular in the field of medicine where quantifying uncertainty is important. [Alaa and van der Schaar \(2017\)](#) regularize counterfactual predictions through their posterior variance and similarly stress the importance to provide confidence in their estimates using credible intervals, but did not investigate the generalization properties of their algorithm and only allowed for limited expressiveness in their algorithm. Similarly, [Jean et al. \(2018\)](#) used posterior variance regularization to learn from unlabeled data in situations where labeled data is scarce for improved performance.

Our work has also strong connections with work on domain adaptation. In particular, estimating ITE requires predictions of outcomes over a different distribution from the observed one. Our ITE error upper bound has similarities with generalization bounds in domain adaptation given by [Johansson et al. \(2016, 2018\)](#). [Johansson et al. \(2018\)](#) and [Zhao et al. \(2019\)](#) have similarly argued against enforcing domain-invariance, and related the loss of predictive power of those representations to the loss of information due to the non-invertibility of learned representations.

2 ITE Problem formulation

To estimate what would happen under different interventions, we first frame the causal problem as a statistical one. We build upon the potential outcomes framework Rubin (2005), this gives us a formal mathematical representation of the potential outcome using the notation $Y(a)$, which reads as observing Y if we do a .

Set up. Consider a population of individuals with each individual i described by a context $X_i \in \mathcal{X}$ (typically \mathbb{R}^d). An intervention T is applied to subjects in the population. An important special case is when interventions are binary, $T \in \{0, 1\}$: individual i 's response to the intervention is a random variable denoted by $Y_i(1)$, whereas i 's natural response when no intervention is applied is denoted by $Y_i(0)$. The two random variables, $Y_i(1), Y_i(0) \in \mathbb{R}$, are known as the potential outcomes. As discussed, in observational data intervention assignments generally depend on the subjects' features, i.e. $T_i \not\perp\!\!\!\perp X_i$. This dependence is quantified via the conditional distribution $p(T_i|X_i)$, also known as the propensity score of subject i .

Assumption 1 (Consistency, ignorability and overlap). *For any individual i , assigned to intervention t_i , we observe $Y_i = Y(t_i)$. Further, $\{Y(t)\}_{t \in T}$ and the data-generating process $p(X, T, Y)$ satisfy strong ignorability: $Y(0), Y(1) \perp\!\!\!\perp T|X$ and overlap: $\forall x, 0 < p(T|X) < 1$.*

Assumption 1 is a sufficient condition for causal identifiability Rubin (2005). Ignorability is also known as the no hidden confounders assumption, indicating that all variables that cause both T and Y are assumed to be measured. Under ignorability therefore, any domain shift in $p(X)$ cannot be due to variables that causally influence T and Y , other than through X . Under Assumption 1, potential outcomes equal conditional expectations: $\mathbb{E}[Y(t)|X = x] = \mathbb{E}[Y|X = x, T = t]$, and we may predict $Y(t)$ by regression.

Objective. We attempt to learn predictors $f_t : \mathcal{X} \rightarrow \mathcal{Y}$ such that $f_t(x)$ approximates $\mathbb{E}[Y|X = x, T = t]$. Predictors $f_t = w_t \circ \phi$ are formed by composing a function $w_t : \mathcal{Z} \rightarrow \mathcal{Y}$ that operates in a feature space \mathcal{Z} defined by a representation $\phi : \mathcal{X} \rightarrow \mathcal{Z}$.

The individual effect of an intervention $T = 1$ in context X is measured by the conditional average treatment effect (CATE), $\tau(x) = \mathbb{E}[Y(1) - Y(0)|X = x]$, and the error in estimation is given by the empirical precision in estimating heterogeneous effects ϵ_{PEHE} , defined as the mean squared error in the estimation of the treatment effect $\tau(x)$,

$$\epsilon_{\text{PEHE}} = \int (\hat{\tau}(x) - \tau(x))^2 p(x) dx \quad (1)$$

Predicting τ for unobserved units involves prediction of both potential outcomes, but since we only observe the "factual" outcome for a specific treatment assignment, and never observe the corresponding "counterfactual" outcome, we never observe any samples of the true treatment effect in an observational dataset. This makes the problem of causal inference fundamentally different from standard supervised learning.

Notation. We analyse the generalization properties of predictors f_t and underlying representations ϕ with respect to ϵ_{PEHE} in the next section. To do so we will make use of the following notation. Let $\mathcal{P}_{x,t}$ denote the input data distribution $p(x, t)$, $\mathcal{P}_t = p(y_t|x)p(x, t)$ the joint factual distribution of x and y_t , $\mathcal{P}_{1-t} = p(y_t|x)p(x, 1-t)$ the joint counterfactual distribution of x and y_t . Each instance in the observed (factual) dataset $\mathcal{D}_t = \{(x_{i,t}, y_{i,t})\}_{i=1}^{N_t}$, is assumed to be sampled *i.i.d* from \mathcal{P}_t . \mathcal{D}_{t-1} will be used to denote the unobserved (counterfactual) data set that results from flipping the treatment assignment for each instance. While we assume ϕ to be deterministic, in the following section we let w_t be a vector of weights with prior distribution $\pi_t = \mathcal{N}(\mathbf{0}, \lambda_t^{-1}\mathbf{I})$, $\lambda_t > 0$. In this sense, π_t defines the hypothesis space \mathcal{F} of f_t and we write $\hat{\rho}_t$ for the posterior distribution of f_t , itself a random variable. We write $\mu(x_i|\mathcal{D}_t, \Theta_t)$ and $\sigma^2(x_i|\mathcal{D}_t, \Theta_t)$ for its posterior mean and variance given context x_i . Θ_t includes all hyperparameters (both shared parameters (in ϕ) and specific parameters to each treatment group).

3 Intuition and theoretical results

Inherent to the approach of learning representations for counterfactual inference is that the representation must trade-off between containing predictive information about factual outcomes while mitigating the information content that drives the treatment selection policy to ensure good generalization on counterfactuals. In this section, we make several observations about the deficiencies of enforcing domain invariance for this purpose and propose alternatives based on the posterior counterfactual variance. We start with a simple example that illustrates the inadequacy of distributional distances and the benefits of counterfactual variance minimization.

Toy example - Counterfactual variance vs. distributional distances. In the middle panels of Figure 2 we show two simulated datasets. The left-hand/red dataset arises from truncated normal distributions with large overlap in the tails; the right-hand/green dataset arises from ordinary normal distributions with small overlap in the tails. In both cases, we show treated and control outcomes in different shades. In both cases, the

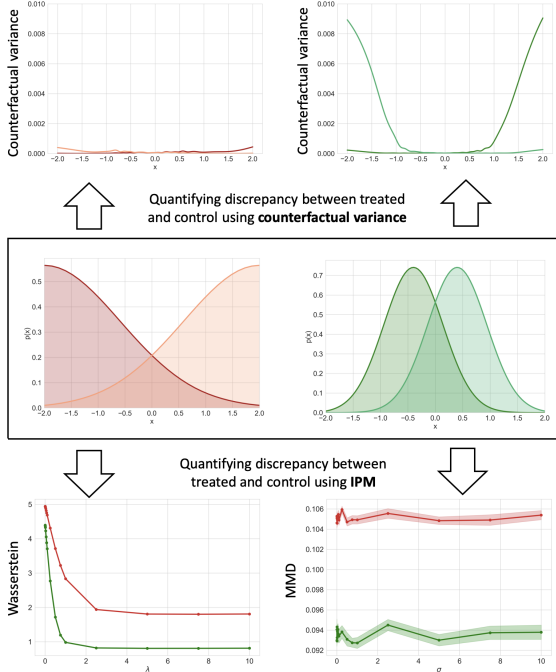


Figure 2: Toy example illustrating the shortcomings of distributional distances, like IPMs, for regularizing representations in causal inference. Despite the fact that sufficient support is satisfied in the red population and not in the green population, IPMs (bottom) give the opposite result, with larger discrepancy between groups in the red population than in the green population. The counterfactual variance (top) accurately describes the lack of support of the green population.

outcome is $y = \text{sinc}(4x)$. The red population satisfies sufficient assumptions for identifiability of causal effects; the green population does not. However, as shown in the bottom panels, both the MMD and Wasserstein distances are *smaller* in the green population than in the red population. In contrast, the top panel shows that the predictive variance of the counterfactual outcomes much better describes this lack of overlap. Counterfactual variance is adaptive to the prediction problem of interest (given that it is obtained from a fitted model), providing a data-dependent measure to quantify distances in the underlying function class, perhaps more precise when the underlying function to be estimated are unknown. IPMs are defined as worst-case distances dependent on a function class to be specified a priori. We make the observation also that IPMs need to be approximated in practice which may be inaccurate for high-dimensional representations and small training data samples [Boissard and Le Gouic \(2014\)](#).

3.1 Generalization bounds

In this subsection, we develop a PAC-Bayes generalization bound [McAllester \(1999\)](#) for the empirical

precision in estimating heterogeneous effects ϵ_{PEHE} , that shows specifically why minimizing counterfactual variance can improve generalization performance.

Theorem 1. *With the assumption that the squared loss function $L : \mathcal{F} \times \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ is sub-gaussian under π_t and \mathcal{P}_t , and using the notation introduced above, the following holds. With probability at least $1 - \delta$ and for any posterior distribution $\hat{\rho}_t$ on \mathcal{F} , ϵ_{PEHE} is upper-bounded by*

$$\begin{aligned} & \sum_{t=0}^1 \left(2\tilde{D}_\infty L_{\hat{\rho}_t}(\mathbf{X}_t, \mathbf{Y}_t) + (\tilde{D}_\infty + 1)\text{Var}_{\hat{\rho}_t}(\mathbf{X}_t) \right. \\ & \quad + \text{Var}_{\hat{\rho}_t}(\mathbf{X}_{1-t}) + \frac{C_{t,1}(N_t, N_{1-t})}{N_t N_{1-t}} \text{KL}(\hat{\rho}_t \| \pi_t) \\ & \quad \left. + \frac{C_{t,2}(N_t, N_{1-t})}{N_t N_{1-t}} \frac{1}{\delta} + C_{t,3} \right) \end{aligned} \quad (2)$$

$C_{t,1}(N_t, N_{1-t})$, $C_{t,2}(N_t, N_{1-t})$ are linear function in their arguments, $C_{t,3}$ is constant, $D_\infty(\mathcal{P}_{x,1-t} \| \mathcal{P}_{x,t}) = \sup_x \frac{p(x,1-t)}{p(x,t)}$, $L_{\hat{\rho}_t}(\mathbf{X}, \mathbf{Y}) = \frac{1}{N} \sum_{i=1}^N (\mu(x_i | \mathcal{D}_t, \Theta_t) - y_i)^2$ is the posterior prediction loss, $\text{Var}_{\hat{\rho}_t}(\mathbf{X})$ is the posterior variance on \mathbf{X} , $\text{Var}_{\hat{\rho}_t}(\mathbf{X}) = \frac{1}{N} \sum_{i=1}^N \sigma^2(x_i | \mathcal{D}_t, \Theta_t)$, $\tilde{D}_\infty = D_\infty(\mathcal{P}_{x,1-t} \| \mathcal{P}_{x,t}) + 1$, and finally $\text{KL}(\cdot \| \cdot)$ is the Kullback-Leibler divergence.

Observe the dependence on the variance with respect to the distribution of the observed data of the counterfactual outcomes $\text{Var}_{\hat{\rho}_t}(\mathbf{X}_{1-t})$. By minimizing the counterfactual variance we can expect the representations ϕ of counterfactual data to encode a relatively smoother prediction function $w \circ \phi(x)$, as can be seen in Figure 1. The estimation of treatment effects is inherently a label-scarce problem as counterfactual data is not observed, representations resulting in smooth prediction curves are especially important to generalize beyond the factual data. This view of the problem of estimating treatment effects emphasizes the need of regularization for good generalization. We note that the sub-gaussian assumption on the loss function may not hold in some cases e.g. large noise. In this case, we may relax this assumption to sub-exponential with all results and proofs otherwise unchanged [Germain et al. \(2016\)](#).

3.1.1 Why distributional distances may be inadequate?

While the toy example clearly illustrates the inability of distributional distances to capture domain overlap, we argue that by enforcing equality in full marginals, optimizing for distributional distances may also overly penalize the model’s ability to predict factual observations when sufficient data is available, that is, when domain overlap is satisfied.

In the following Theorem, we provide a bound on the generalization error on counterfactual data that illustrates the interplay between distribution mismatch and prediction loss on factual data. We show that distribution mismatch becomes decreasingly relevant with increasing data set size.

Theorem 2. *Assume the notation introduced above, for any posterior distribution $\hat{\rho}_t$ on \mathcal{F} , the expected counterfactual Gibbs risk¹ $R_{\mathcal{P}_{1-t}}(G_{\hat{\rho}_t})$ is bounded above by,*

$$\frac{1}{2}D_{\mathcal{P}_{x,1-t}}(G_{\hat{\rho}_t}) + D_{\infty}(\mathcal{P}_{x,1-t} \parallel \mathcal{P}_{x,t})L_{\mathcal{P}_t}(G_{\hat{\rho}_t}). \quad (3)$$

where $D_{\mathcal{P}_{x,1-t}}(\hat{\rho}_t) = \mathbb{E}_{x \sim \mathcal{P}_{x,1-t}}[\sigma^2(x|\mathcal{D}_t, \Theta_t)]$ is the expected counterfactual variance, $L_{\mathcal{P}_t}(\hat{\rho}_t) = \mathbb{E}_{(x,y) \sim \mathcal{P}_t}[(\mu(x|\mathcal{D}_t, \Theta_t) - y)^2 + \frac{1}{2}\sigma^2(x|\mathcal{D}_t, \Theta_t)]$ is the expected factual loss, and finally $D_{\infty}(\mathcal{P}_{x,1-t} \parallel \mathcal{P}_{x,t}) = \sup_x \frac{p(x,1-t)}{p(x,t)}$.

The bound in Equation (3) describes the interaction between the distribution mismatch and the prediction error on factual data $L_{\mathcal{P}_t}(G_{\hat{\rho}_t})$. $D_{\infty}(\mathcal{P}_{x,1-t} \parallel \mathcal{P}_{x,t})$ is large if, for some x , $\mathcal{P}_{x,t}$ has small density while $\mathcal{P}_{x,1-t}$ has large density (that is when there is poor overlap between them), which understandably, makes minimizing the counterfactual risk $R_{\mathcal{P}_{1-t}}(G_{\hat{\rho}_t})$ harder because few examples from the other population are observed for a given context. However, note that $D_{\infty}(\mathcal{P}_{x,1-t} \parallel \mathcal{P}_{x,t})$ is multiplied by the expected factual loss $L_{\mathcal{P}_t}(G_{\hat{\rho}_t})$ (which decreases as the number of factual samples increases, see Equation (10) in Theorem 7 of Appendix 1.3). The distribution mismatch $D_{\infty}(\mathcal{P}_{x,1-t} \parallel \mathcal{P}_{x,t})$ thus becomes less important for generalization, if we can minimize the expected factual loss $L_{\mathcal{P}_t}(G_{\hat{\rho}_t})$ arbitrarily well. This suggests that optimizing for distributional distances between treated and control data *at the expense* of prediction error on the factual data may be counterproductive. Representations regularized with distributional distances may thus shrink the hypothesis space and converge on solutions that, although balanced across treatment groups, lose their predictive power.

3.2 Why encourage preserving information content?

Assumption 1 gives sufficient conditions for unbiased estimation of the treatment effect using observational data, but identifiability need not hold with respect to the feature representation $Z = \phi(X)$, even if it does

¹The expected counterfactual Gibbs risk is an upper bound of the expected Bayesian counterfactual risk, a detailed discussion of Gibbs and Bayesian risk is given in Appendix 1.2.

with respect to X . For instance consider $\phi^{-1}(z) := \{x : \phi(x) = z\}$, then,

$$\begin{aligned} p(Y|\phi(x), t) &= \frac{\int_{x \in \phi^{-1}(z)} p(Y|x, t)p(x|t)}{\int_{x \in \phi^{-1}(z)} p(x|t)} \\ &= \frac{\int_{x \in \phi^{-1}(z)} p(Y|x)p(x|t)}{\int_{x \in \phi^{-1}(z)} p(x|t)} \\ &\neq p(Y|\phi(x)) \end{aligned} \quad (4)$$

in general, with equality only if ϕ is invertible. The conditional independence in the strong ignorability assumption, $Y(0), Y(1) \perp\!\!\!\perp T | \phi(X)$, required for estimating the treatment effect need not hold for non-invertible transformations. In this sense, we may be introducing unobserved confounders in representation space we hypothesize by the information lost in the map ϕ . Observe also that our objective, the conditional average treatment effect $\tau(x) = \mathbb{E}[Y(1) - Y(0) | X = x]$, is expressed in terms of expectations. Similarly, it holds that in feature space, $\mathbb{E}[Y(1) - Y(0) | \phi(x)] = \int_{x \in \phi^{-1}(z)} \mathbb{E}[Y(1) - Y(0) | x] dx$ will in general *not* be equal to our quantity of interest $\tau(x) = \mathbb{E}[Y(1) - Y(0) | x]$ unless ϕ is invertible.

4 Algorithms

In this section, we describe a model for counterfactual estimation, called DKLITE² (Deep Kernel Learning for Individualized Treatment Effects), motivated by our analysis. Our proposed prediction algorithm works in a feature space induced by ϕ indirectly via a kernel function $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, that intuitively models the correlation between inputs $x, x' \in \mathcal{X}$ in a possibly high-dimensional feature space. We extend the expressiveness of K by transforming the inputs through a non-linear mapping $\phi : \mathcal{X} \rightarrow \mathbb{R}^{d_\phi}$ to form $k(\phi(x), \phi(x'))$, a *deep kernel* Wilson et al. (2016). The dimension d_ϕ of $\phi(x)$ can be chosen arbitrarily, and may therefore differ from the dimension of x . We let ϕ_θ be parameterized by a neural network with parameters θ to encode the information content of input variables.

Prediction model. The potential outcomes given individuals covariates are assumed to take the form,

$$y_{i,t} = f_t(x_i) + \epsilon_{i,t}, \quad t \in \{0, 1\} \quad (5)$$

where $f_t(x_i) = w_t^\top \phi(x_i)$ is given as a linear combination of feature representations $\phi(\cdot)$. $w_t \in \mathbb{R}^d$ is a weight vector, $\phi(\cdot)$ the representation layer and $\epsilon_{i,t} \sim \mathcal{N}(0, \beta_t^{-1})$ is a noise variable, $t = 0, 1$. We specify our uncertainty in parameter values through the specification of prior distributions. Combined with the likelihood of observed data, these determine the posterior distribution

²An implementation of DKLITE is available at <https://bitbucket.org/mvdschaar/mlforhealthlabub>.

of parameters and predictive distributions of treatment effects. For each component of the weight vector, independently, we assume $w_t \sim \mathcal{N}(0, \lambda_t^{-1})$. With knowledge of ϕ , the posterior of w_t is given by,

$$p(w_t | \mathcal{D}_t, \Theta_t) = \mathcal{N}(\mathbf{m}_{w_t}, \mathbf{K}_{w_t}^{-1}) \quad (6)$$

where $\Theta_t = \{\phi, \beta_t, \lambda_t\}$, the posterior mean $\mathbf{m}_{w_t} = \beta_t \mathbf{K}_{w_t}^{-1} \Phi_t^\top \mathbf{Y}_t$, the posterior covariance $\mathbf{K}_{w_t}^{-1} = \beta_t \Phi_t^\top \Phi_t + \lambda_t \mathbf{I}_{d_\phi \times d_\phi}$ and Φ_t the representation of \mathbf{X}_t . However, it can be difficult to encode our knowledge about good representation spaces for counterfactual estimation in a Bayesian prior. Posterior regularization offers a more direct and flexible mechanism for guiding and controlling the posterior distribution.

Regularized Bayes. The regularized posterior is the solution of the following optimization problem [Zhu et al. \(2014\)](#); [Jean et al. \(2018\)](#):

$$\inf_{q(f|\mathcal{D})} \mathcal{L}(q(f|\mathcal{D})) + \Omega(q(f|\mathcal{D})) \quad (7)$$

where $\mathcal{L}(q(f|\mathcal{D}))$ is the (KL-divergence between the approximate posterior $q(f|\mathcal{D})$ and the true posterior $p(f|\mathcal{D})$, and $\Omega(q(f|\mathcal{D}))$ is a regularizer of the approximate posterior $q(f|\mathcal{D})$. We attempt therefore to learn a posterior distribution as close as possible to the true posterior while also fulfilling the requirements imposed by the regularization term.

Inference. The predictive distribution for a testing point x is obtained by marginalising over the posterior w_t . It is given by

$$f_t(x) \sim \mathcal{N}(\mu(x|\mathcal{D}_t, \Theta_t), \sigma^2(x|\mathcal{D}_t, \Theta_t)) \quad (8)$$

where $\mu(x|\mathcal{D}_t, \Theta_t) = \mathbf{m}_{w_t}^\top \phi(x)$ and $\sigma^2(x|\mathcal{D}_t, \Theta_t) = \phi(x)^\top \mathbf{K}_{w_t}^{-1} \phi(x)$. Point estimates are given for example by the posterior mean or median, optimal for minimizing squared loss or absolute loss, respectively. Note also that knowledge of the full posterior allows us to quantify our uncertainty around point estimates through credible intervals, especially useful in medicine and public policy, for example.

4.1 Learning ϕ

The framework of regularized Bayes described in (7) provides us with a flexible optimization objective, many previous algorithms for treatment effects can be reduced to specific instantiations of this problem. In this section we describe the likelihood and regularization terms that we use to encode the intuition and insights derived from section 3. In Appendix 2, we show that the final loss function we derive can be reformulated a specific instance of problem (7).

Likelihood of the factual data. Assuming the discriminative model introduced in 5, we encourage good

prediction of the factual data by optimizing for the negative log-likelihood of the factual data,

$$\mathcal{L}_{lik} = - \sum_{t=0}^1 \log(p(\mathbf{Y}_t | \mathbf{X}_t, \Theta_t)) \quad (9)$$

Observe that we may re-write the negative log-likelihood, \mathcal{L}_{lik} , as

$$\begin{aligned} \mathcal{L}_{lik} = & \sum_{t=0}^1 \left(\frac{N_t}{2} \ln(2\pi\beta_t^{-1}) + \text{KL}(\hat{\rho}_t \| \pi_t) \right. \\ & \left. + \frac{N_t\beta_t}{2} [\text{Var}_{\hat{\rho}_t}(\mathbf{X}_t) + L_{\hat{\rho}_t}(\mathbf{X}_t, \mathbf{Y}_t)] \right) \end{aligned} \quad (10)$$

From Theorem 1, we see that the empirically estimable quantities to optimize for in the upperbound (2), are $\text{Var}_{\hat{\rho}_t}(\mathbf{X}_{1-t})$, $\text{Var}_{\hat{\rho}_t}(\mathbf{X}_t)$, $L_{\hat{\rho}_t}(\mathbf{X}_t, \mathbf{Y}_t)$ and $\text{KL}(\hat{\rho}_t \| \pi_t)$. The latter three already exist in our objective \mathcal{L}_{lik} , as shown in Equation (10).

Posterior variance as regularization. Therefore, by including an empirical estimate of the counterfactual variance $\text{Var}_{\hat{\rho}_t}(\mathbf{X}_{1-t})$ as a regularizer in our objective we are effectively optimizing for all terms given by the PAC-Bayes upperbound in equation (2). We write this regularization term as,

$$\mathcal{L}_{var} = \sum_{t=0}^1 \frac{1}{N_{1-t}} \sum_{i=1}^{N_{1-t}} \sigma^2(x_{i,1-t} | \mathcal{D}_t, \Theta_t) \quad (11)$$

Since the deep kernel parameters are jointly learned, the neural net ϕ is encouraged to learn a feature representation in which the counterfactual examples are close to the factual examples, thereby reducing the variance on our predictions. The implication is that we are optimizing for representations where counterfactual data tend to cluster around the representations of factual data. This is a way to see how intuitively that we are encouraging overlap in support in representation space *without* enforcing equality in densities (i.e. the size of the factual and counterfactual "clusters" need not coincide).

Invertibility as regularization. While the loss due to non-invertible representations is not directly observable, we may associate it with the information content of x lost in $\phi(x)$ and we found it to be an important source of gain in performance, empirically. We can encourage information content preservation with an additional decoder $\psi : \mathbb{R}^{d_\phi} \rightarrow \mathcal{X}$ - a neural network with the reversed structure of network ϕ (the encoder in this sense) - trained to reconstruct the input x from $\phi(x)$. The reconstruction loss is given as,

$$\mathcal{L}_{rec} = \sum_{t=0}^1 \frac{1}{N_t} \sum_{i=1}^{N_t} \|x_{i,t} - \psi(\phi(x_{i,t}))\|_2^2 \quad (12)$$

	IHDP ($\sqrt{\hat{\epsilon}_{PEHE}}$)		Twins ($\sqrt{\hat{\epsilon}_{PEHE}}$)		Jobs ($\hat{\mathcal{R}}_{pol}(\pi_f)$)	
	In-sample	Out-sample	In-sample	Out-sample	In-sample	Out-sample
OLS/LR ₁	5.8 ± .3	5.8 ± .3	.319 ± .001	.318 ± .007	.22 ± .00	.23 ± .02
OLS/LR ₂	2.4 ± .1	2.5 ± .1	.320 ± .002	.320 ± .003	.21 ± .00	.24 ± .01
BLR	5.8 ± .3	5.8 ± .3	.312 ± .003	.323 ± .018	.22 ± .01	.26 ± .02
k-NN	2.1 ± .1	4.1 ± .2	.333 ± .001	.345 ± .007	.22 ± .00	.26 ± .02
BART	2.1 ± .1	2.3 ± .1	.347 ± .009	.338 ± .016	.23 ± .00	.25 ± .02
R-Forest	4.2 ± .2	6.6 ± .3	.366 ± .002	.321 ± .005	.23 ± .01	.28 ± .02
C-Forest	3.8 ± .2	3.8 ± .2	.366 ± .003	.316 ± .011	.19 ± .00	.20 ± .02
BNN	2.2 ± .1	2.1 ± .1	.325 ± .003	.321 ± .018	.20 ± .01	.24 ± .02
TARNET	.88 ± .02	.95 ± .02	.317 ± .002	.315 ± .003	.17 ± .01	.21 ± .01
CAR _{WASS}	.72 ± .02	.76 ± .02	.315 ± .007	.313 ± .008	.17 ± .01	.21 ± .01
CMGP	.63 ± .08	.74 ± .11	.320 ± .002	.319 ± .008	.22 ± .03	.24 ± .05
DKLITE	.52 ± .02	.65 ± .03	.288 ± .001	.293 ± .003	.13 ± .01	.14 ± .01

Table 1: Mean performance (lower better) of individualized treatment effect estimation and standard deviation.

We remark that some recent density estimation approaches, e.g. normalization flows [Rezende and Mohamed \(2015\)](#) and non-volume preserving transformation [Dinh et al. \(2016\)](#), can be used as an alternative method to achieve invertibility.

Final loss function. Based on the objectives described above, our final loss trades-off between maximizing the likelihood of the observed (factual) data under our model, minimizing the predictive variance of the counterfactual outcomes and minimizing the reconstruction loss of the representations. The loss is given as follows,

$$\mathcal{L}_{fin} = \mathcal{L}_{lik} + \alpha_1 \mathcal{L}_{var} + \alpha_2 \mathcal{L}_{rec} \quad (13)$$

$\alpha_1 > 0$ and $\alpha_2 > 0$ are hyperparameter. Standard methods for hyperparameter selection are not generally applicable for choosing hyperparameters because counterfactuals are never observed. As an approximation scheme, we replace the missing counterfactuals with their nearest factual neighbour to compute the treatment effects in cross-validation [Shalit et al. \(2017\)](#).

5 Experiments

The primary focus of our experiments will be on the comparison of prediction performance on benchmark tasks, the use of the posterior variance for decision-making and, a deeper analysis of learned representations and source of performance gain. In the following, we start by introducing competing prediction algorithms before giving a brief description of the data and appropriate performance metrics.

Competing algorithms. We compare DKLITE with a total of 11 algorithms. First we evaluate least squares

regression using treatment as an additional input feature OLS/LR₁), we consider separate least squares regressions for each treatment (OLS/LR₂), we evaluate balancing linear regression (BLR) [Johansson et al. \(2016\)](#), k-nearest neighbor (k-NN) [Crump et al. \(2008\)](#), Bayesian additive regression trees (BART) [Chipman et al. \(2010\)](#), random forests (R-Forest) [Breiman \(2001\)](#), causal forests (C-Forest) [Wager and Athey \(2018\)](#), balancing neural networks (BNN) [Johansson et al. \(2016\)](#), treatment-agnostic representation network (TARNET), counterfactual regression with Wasserstein distance (CAR_{WASS}) [Shalit et al. \(2017\)](#), and multi-task gaussian process (CMGP) [Alaa and van der Schaar \(2017\)](#).

Data. Causal inference models are often impossible to reliably validate using real-world data due to the absence of counterfactual outcomes. Various established approaches for evaluating causal models have been proposed, which we use for our analysis. We describe these briefly below and refer the reader to the accompanying references and section 5 of Appendix for further details. Additional results (e.g. average treatment effect estimation) are included in the same Appendix section. We consider **IHDP** (*747 instances described by 25 covariates*) [Hill \(2011\)](#); [Alaa and van der Schaar \(2017\)](#); [Shalit et al. \(2017\)](#); [Yao et al. \(2018\)](#) in which counterfactual outcomes are randomly generated via a predefined probabilistic model; **Twins** (*11300 instances described by 30 covariates*), in which outcomes are observed but the treatment assignment is simulated; and finally, **Jobs** (*3212 instances described by 7 covariates*) [Shalit et al. \(2017\)](#); [LaLonde \(1986\)](#); [Dehejia and Wahba \(2002\)](#) which is constructed from a mixture of experimental and observational data with the treatment.

Metrics. The metrics used to evaluate each data set

differ slightly depending on the available outcome (real or simulated) and the available treatment assignment mechanisms (known or unknown). For the **IHDP**, we use the empirical precision in estimating treatment effects $\hat{\epsilon}_{\text{PEHE}} = \frac{1}{N} \sum_{i=1}^N (\tau(x_i) - \hat{\tau}(x_i))^2$. For **Twins**, we use the observed precision in estimating heterogeneous effects, $\tilde{\epsilon}_{\text{PEHE}} = \frac{1}{N} \sum_{i=1}^N (y_{i,1} - y_{i,0} - \hat{\tau}(x_i))^2$. For **Jobs**, we use the policy risk that measures the expected loss if the treatment is taken according to the ITE policy prescribed by the algorithm, $\mathcal{R}_{\text{pol}} = 1 - \mathbb{E}[Y(1)|\pi(x) = 1]P(\pi(x) = 1) + \mathbb{E}[Y(0)|\pi(x) = 0]P(\pi(x) = 0)$ where $\pi(x) = 1$ if $\hat{y}(1) - \hat{y}(0) > 0$ and $\pi(x) = 0$, otherwise.

5.1 Prediction performance results

We report in-sample and out-of-sample performance in Table 1. DKLITE targets aspects of the treatment effect estimation problem that have not been considered before. Learning with such an objective significantly outperforms all competing algorithms and does so on all benchmark data sets. The most relevant comparison is perhaps with BNN Johansson et al. (2016) and CAR_{WASS} Shalit et al. (2017), neural network models that enforces domain invariance through distributional distances. The performance gain highlights the predictive power of our representations.

5.2 Source of gain

	\mathcal{L}_{lik}	$\mathcal{L}_{\text{lik}} + \mathcal{L}_{\text{var}}$	$\mathcal{L}_{\text{lik}} + \mathcal{L}_{\text{rec}}$
IHDP			
In-sample	1.46 ± .01	.98 ± .04	.96 ± .04
Out-sample	1.95 ± .14	1.24 ± .09	1.28 ± .13
Twins			
In-sample	.308 ± .004	.292 ± .002	.291 ± .001
Out-sample	.323 ± .008	.294 ± .007	.293 ± .003
Jobs			
In-sample	.18 ± .01	.14 ± .01	.15 ± .01
Out-sample	.28 ± .01	.23 ± .01	.24 ± .01

Table 2: Source of performance gain.

In this section, we analyze more deeply the contribution to the performance gain of each component of our loss. We evaluate DKLITE optimized for different components of $\mathcal{L}_{\text{fin}} = \mathcal{L}_{\text{lik}} + \alpha_1 \mathcal{L}_{\text{var}} + \alpha_2 \mathcal{L}_{\text{rec}}$. As can be seen in Table 2, including regularization based on the counterfactual variance ($\mathcal{L}_{\text{lik}} + \alpha_1 \mathcal{L}_{\text{var}}$) and reconstruction loss ($\mathcal{L}_{\text{lik}} + \alpha_2 \mathcal{L}_{\text{rec}}$), each evaluated separately, already provides a significant gain in performance with respect to optimization on the factual data only (\mathcal{L}_{lik}). Importantly though, combining them (\mathcal{L}_{fin}) improves performance further by an order of magnitude (see DKLITE in Table 1), which suggests that \mathcal{L}_{var} and \mathcal{L}_{rec} capture to some extent "orthogonal" sources of gain. The gain is especially important on relatively smaller data sets, such as IHDP with 747 individuals, and to a lesser extent on bigger data sets. These results illustrate the behaviour suggested by Equation 3

in Theorem 2: distribution mismatch between groups becomes decreasingly relevant with increasing data set size. In this setting, the error on the factual data (\mathcal{L}_{lik}) drives generalization performance.

5.3 Leveraging the predicted uncertainty

Data-driven solutions for decision support have most often been proposed without methods to quantify and control their uncertainty in a decision. In contrast, in medicine for example, a physician knows whether she is uncertain about a case and will consult more experienced colleagues if needed. We use this idea to show that uncertainty informed treatment effect estimation can improve performance. An instantiation of this approach, termed DKLITE-U, is given by referring the 10 % most uncertain predictions for further scrutiny. Performance in comparison to DKLITE is given in Table 3. Note that especially on small data sets, such as IHDP, predictions can be significantly improved.

	DKLITE	DKLITE-U
IHDP		
In-sample	.52 ± .02	.46 ± .02
Out-sample	.60 ± .03	.53 ± .02
Twins		
In-sample	.288 ± .001	.287 ± .001
Out-sample	.293 ± .003	.292 ± .002
Jobs		
In-sample	.13 ± .01	.13 ± .01
Out-sample	.14 ± .01	.14 ± .01

Table 3: Performance of DKLITE-U on all data sets.

6 Discussion

In many domains understanding the effect of interventions at an individual level is crucial, but prediction of those potential outcomes is challenging. Despite their empirical success, we find that algorithms enforcing representations to satisfy domain invariance is often too strong a requirement for causal predictions. This stems from the fact that overlapping support is sufficient for identifiability of the causal effect and equality in densities is not necessary.

We have proposed a bound on generalization performance that shows the dependence on domain overlap through the counterfactual variance which we interpret as a proxy for domain overlap, and highlighted the need for invertible latent maps. These results motivated novel regularization criteria that we incorporated in a deep kernel learning framework through posterior regularization. We hypothesize that many existing models, both frequentist (through representer theorems if optimized with empirical risk minimization) and Bayesian, can be reduced to specific instances of our framework with different regularization terms. We leave a formal derivation of this unifying theory for future work.

Acknowledgements

We thank the anonymous reviewers for valuable feedback. This work was supported by GlaxoSmithKline (GSK), the Alan Turing Institute under the EPSRC grant EP/N510129/1, the US Office of Naval Research (ONR), and the National Science Foundation (NSF): grant numbers ECCS1462245, ECCS1533983, and ECCS1407712.

References

- Fredrik Johansson, Uri Shalit, and David Sontag. Learning representations for counterfactual inference. In *International conference on machine learning*, pages 3020–3029, 2016.
- Fredrik D Johansson, Nathan Kallus, Uri Shalit, and David Sontag. Learning weighted representations for generalization across designs. *arXiv preprint arXiv:1802.08598*, 2018.
- Liuyi Yao, Sheng Li, Yaliang Li, Mengdi Huai, Jing Gao, and Aidong Zhang. Representation learning for treatment effect estimation from observational data. In *Advances in Neural Information Processing Systems*, pages 2633–2643, 2018.
- Uri Shalit, Fredrik D Johansson, and David Sontag. Estimating individual treatment effect: generalization bounds and algorithms. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3076–3085. JMLR. org, 2017.
- Ahmed M Alaa and Mihaela van der Schaar. Bayesian inference of individualized treatment effects using multi-task gaussian processes. In *Advances in Neural Information Processing Systems*, pages 3424–3432, 2017.
- Neal Jean, Sang Michael Xie, and Stefano Ermon. Semi-supervised deep kernel learning: Regression with unlabeled data by minimizing predictive variance. In *Advances in Neural Information Processing Systems*, pages 5322–5333, 2018.
- Han Zhao, Remi Tachet des Combes, Kun Zhang, and Geoffrey J Gordon. On learning invariant representation for domain adaptation. *arXiv preprint arXiv:1901.09453*, 2019.
- Donald B Rubin. Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, 100(469):322–331, 2005.
- Emmanuel Boissard and Thibaut Le Gouic. On the mean speed of convergence of empirical and occupation measures in wasserstein distance. In *Annales de l’IHP Probabilités et statistiques*, volume 50, pages 539–563, 2014.
- David A McAllester. Some pac-bayesian theorems. *Machine Learning*, 37(3):355–363, 1999.
- Pascal Germain, Francis Bach, Alexandre Lacoste, and Simon Lacoste-Julien. Pac-bayesian theory meets bayesian inference. In *Advances in Neural Information Processing Systems*, pages 1884–1892, 2016.
- Andrew Gordon Wilson, Zhiting Hu, Ruslan Salakhutdinov, and Eric P Xing. Deep kernel learning. In *Artificial Intelligence and Statistics*, pages 370–378, 2016.
- Jun Zhu, Ning Chen, and Eric P Xing. Bayesian inference with posterior regularization and applications to infinite latent svms. *The Journal of Machine Learning Research*, 15(1):1799–1847, 2014.
- Danilo Jimenez Rezende and Shakir Mohamed. Variational inference with normalizing flows. *arXiv preprint arXiv:1505.05770*, 2015.
- Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.
- Richard K Crump, V Joseph Hotz, Guido W Imbens, and Oscar A Mitnik. Nonparametric tests for treatment effect heterogeneity. *The Review of Economics and Statistics*, 90(3):389–405, 2008.
- Hugh A Chipman, Edward I George, Robert E McCulloch, et al. Bart: Bayesian additive regression trees. *The Annals of Applied Statistics*, 4(1):266–298, 2010.
- Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018.
- Jennifer L Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011.
- Robert J LaLonde. Evaluating the econometric evaluations of training programs with experimental data. *The American economic review*, pages 604–620, 1986.
- Rajeev H Dehejia and Sadek Wahba. Propensity score-matching methods for nonexperimental causal studies. *Review of Economics and statistics*, 84(1):151–161, 2002.