

---

# Balance Regularized Neural Network Models for Causal Effect Estimation

---

Mehrdad Farajtabar\*

Andrew Lee

Yuanjian Feng

Vishal Gupta

Peter Dolan

Harish Chandran

Martin Szummer

DeepMind

## Abstract

Estimating individual and average treatment effects from observational data is an important problem in many domains such as healthcare and e-commerce. In this paper, we advocate balance regularization of multi-head neural network architectures. Our work is motivated by representation learning techniques to reduce differences between treated and untreated distributions that potentially arise due to confounding factors. We further regularize the model by encouraging it to predict control outcomes for individuals in the treatment group that are similar to control outcomes in the control group. We empirically study the bias-variance trade-off between different weightings of the regularizers, as well as between inductive and transductive inference.

## 1 Introduction

Causal inference practitioners are increasingly applying deep learning to estimate causal effects from observational data, capitalizing on the expressivity and representational power of neural networks. Authors in [1, 2, 3] propose neural network architectures with separate outputs ("heads") for each potential outcome [4]. However, when learning from observational data, rather than from randomized trials, we must address potential bias due to confounding factors. Fortunately, neural networks allow for flexible forms of regularization. We devise novel regularizers to tackle confounding, and also leverage existing regularizers from domain adaptation and covariate shift [1].

In this paper we extend recent deep causal inference methods by combining their approaches into one architecture, and study how different components affect causal effect estimation. We consider a multi-head model inspired by [3] that predicts the potential outcomes of data containing control and treated sets of data. Neural network models have shown improvements in estimation accuracy, but there is little understanding of how. We elucidate them by decomposing their mean squared error improvements into bias and variance. Our contributions can be summarized in three parts.

First, we evaluate the ability of the multi-head model to improve estimation accuracy by comparing it with a single-head alternative, and by analyzing the bias and variance of inductive and transductive formulations.

Second, we evaluate the ability of a maximum mean discrepancy loss (MMD) [5] on an intermediate embedding layer to regularize a shared representation of treatment and control units by penalizing covariate imbalance, inspired by [2]. We interpret this as a form of balance regularization as the prediction task of both heads influence the balancing of covariates between treated and control groups.

---

\*Correspondence to farajtabar@google.com

Third, we introduce balance regularization with a prognostic score (PRG) [6]. A prognostic score measures the similarity between the predicted untreated outcomes of treated and control groups by using standardized mean difference or a two sample Kolmogorov-Smirnov (KS) statistic, both of which can be minimized with an additional loss.

## 2 Preliminaries

The observed data  $D$  consists of  $n$  triplets  $\{(x_i, t_i, y_i)\}_{i=1\dots n}$ , where  $x_i$  is the feature or context of the  $i$ -th unit,  $t_i$  is the treatment (intervention), and  $y_i$  is the associated outcome. For simplicity we focus on binary treatment where  $t_i = \{0, 1\}$ . The potential outcome for unit  $i$  is  $Y_i^0$  (associated to  $t = 0$ ) or  $Y_i^1$  (associated to  $t = 1$ ). The observed outcome can thus be expressed as:

$$y_i = t_i Y_i^1 + (1 - t_i) Y_i^0. \quad (1)$$

For a unit  $i$  described by  $x_i$  we are interested in estimating the impact of the treatment  $im(x_i)$  (or individual treatment effect (ITE)), i.e.,

$$im(x_i) = Y_i^1(x_i) - Y_i^0(x_i). \quad (2)$$

The fundamental problem in causal inference is that only one of the two potential outcomes  $Y_i^0(x_i)$  or  $Y_i^1(x_i)$  is observed. We assume ‘no hidden confounding’ so that causal effect can be identified and confounding can be controlled for. We learn a function  $f(x, t) \rightarrow y$  that maps the context and treatment indicator to the outcome.

After learning this function, we can employ an *inductive* approach to estimate impact:

$$im(x_i) = f(x_i, 1) - f(x_i, 0). \quad (3)$$

Alternatively, for data inside our observation/training set where one factual outcome is observed, impact can be *transduced* via

$$im(x_i) = y_i - f(x_i, 1 - t_i) \quad \text{for } t_i = 1, \quad (4)$$

$$im(x_i) = f(x_i, 1 - t_i) - y_i \quad \text{for } t_i = 0. \quad (5)$$

## 3 Balance Regularized Neural Network Model

Causal inference from observational data should take confounding factors that affect both treatment and outcome into consideration. The multi-head neural network architecture (Fig. 1) adjusts for confounders and prevent biased estimates inspired by the representation learning literature [7].

If the treatment assignment were random the distribution of treated units  $p^1 \triangleq p\{x|(x, t, y), t = 1\}$  would be the same as for untreated units  $p^0 \triangleq p\{x|(x, t, y), t = 0\}$ . Therefore, one could train two separate models  $f^0(x) \triangleq f(x, 0)$  and  $f^1(x) \triangleq f(x, 1)$  to predict untreated and treated outcomes respectively. However, it is common to find that in observational studies,  $f^0$  and  $f^1$  are trained on differing  $p^0$  and  $p^1$  due to confounding and selection bias, which are the most common obstacles to discovering causal relations [8, 9]. Therefore,  $f^0$  is not accurate for samples drawn from  $p^1$ , making causal estimation,  $im(x_i) = y_i - f^0(x_i)$  on triplet  $(x_i, t_i = 1, y_i)$  biased. Similarly, for a data point in the control group  $(x_i, t_i = 0, y_i)$  the treatment model  $f^1(x_i)$  may provide a confounded estimate.

The fundamental problem of causal inference (i.e. only one of the two potential outcomes is observed) can be addressed under certain assumptions [10, 11] by matching each unit that received treatment ( $t = 1$ ) with its “nearest” units from the group that did not receive treatment ( $t = 0$ ). Propensity score matching [12] is a widely employed method for matching and grouping units based on their propensity for treatment. However, propensity score-based nearest neighbor matching has increasing bias with increasing dimension [13]. To help address this, random projections [14] and local linear embedding [15] can reduce bias by embedding the data into lower dimensional manifolds. The success of these methods inspired us using other ways of embedding and representation learning.

We learn an intermediate embedding that brings the treated and untreated distributions closer to each other before passing examples into the outcome regression part. Inspired by recent works [1, 2, 3] on deep causal inference, we employ a two-head shared-bottom model.

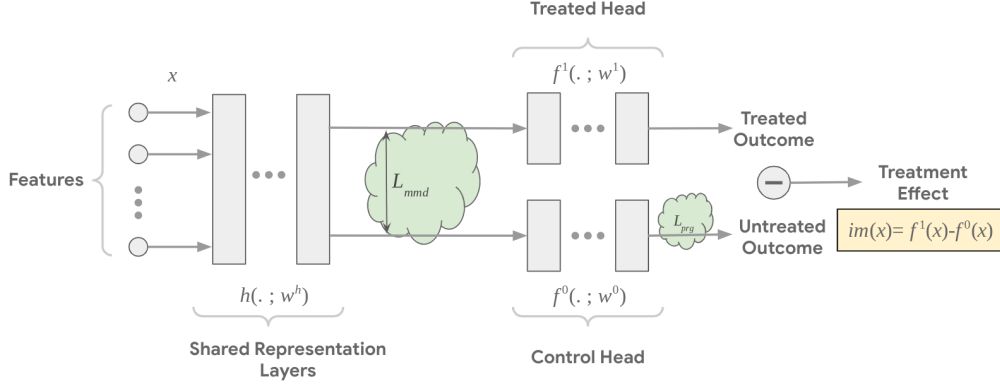


Figure 1: The architecture of our multi-head neural network model

We use the function  $h$  to map examples from  $p^0 \cup p^1$  to a  $d$ -dimensional space. Because of the layer-wise structure and implicit inductive bias that stochastic gradient descent optimizer [16, 17] imposes on the learning problem, the embedded individuals are restricted and the bottleneck can bring the two distributions together. This can be attributed to the middle layers that capture high level representations of concepts in the data [18, 19] that are shared between the two groups. The subsequent heads capture distinctions between the groups that correspond to the individual outcomes.

A reasonable estimation model should be able to fit to the factual (observed) outcomes well. In other words, similar to a standard supervised learning problem, it should minimize the following objective:

$$\mathcal{L}_{fit} = \frac{1}{2} \sum_{i:t_i=0} (f^0(h(x_i; w^h); w^0) - y_i)^2 + \rho \frac{1}{2} \sum_{i:t_i=1} (f^1(h(x_i; w^h); w^1) - y_i)^2, \quad (6)$$

where,  $h(\cdot; w^h)$  represents the shared layers with parameters  $w^h$ ,  $f^1(\cdot; w^1)$  is the treated head with parameters  $w^1$  and  $f^0(\cdot; w^0)$  is the untreated head with parameters  $w^0$ . Here  $\rho$  is a hyperparameter that weights either head in case the treated and untreated objectives have different size or importance. In our experiments we set  $\rho$  to 1. An explicit loss term is used to penalize the distance between distributions of the shared representation. We use Maximum Mean Discrepancy [5] (MMD) as the distance measure

$$\mathcal{L}_{mmd} = \sup_{\|g\|_{\mathcal{H}} \leq 1} (E_{x^0 \sim p^0}[g(h(x^0))] - E_{x^1 \sim p^1}[g(h(x^1))]), \quad (7)$$

where  $g$  is restricted to a norm-1 ball in the Hilbert space  $\mathcal{H}$  associated with a Gaussian kernel. Further theoretical study and ablation on the choice of Kernel and MMD distance is left as an interesting future direction.

For additional balance regularization, we propose a prognostic loss to balance covariates that are predictive of the outcome by assessing the similarity in the predicted untreated outcome of both treatment and control groups. Here, the motivation (i.e. the prior inductive bias of the regularizer) is different from the MMD loss. This objective helps to ensure that both groups have similar baseline outcomes without treatment. Standardized mean difference or the KS test statistic can measure differences between outcomes of groups, and differ in being parametric or nonparametric, respectively. This regularizer takes the form

$$\mathcal{L}_{prg} = \text{KS}(\{f^0(h(x_i; w^h); w^0) | i : t_i = 0\}, \{f^0(h(x_i; w^h); w^0) | i : t_i = 1\}), \quad (8)$$

where, KS is the Kolmogorov-Smirnov statistic computed on two sets of samples.

The overall objective is then

$$\min_{w^h, w^0, w^1} \mathcal{L}_{fit} + \gamma \mathcal{L}_{mmd} + \lambda \mathcal{L}_{prg}, \quad (9)$$

given the weights  $\gamma$  and  $\lambda$ , which can take any non-negative value depending on the application. In the experiments section we will see how the choice of these weights affects bias, variance, and

thus, mean squared error of the estimation. Stochastic gradient descent<sup>2</sup> is used to find minimize the objective.

The proposed model architecture improves data efficiency by shared parameters across treated and untreated units. The shared parameters store commonalities between both regression tasks when optimizing for each individual task [20].

The end-to-end nature of the training procedure offers an interesting approach to leverage increasing computation power, scalability, and the rise of large datasets to tackle causal questions, as evidenced by many recent works [21, 2, 22].

We compare inductive and transductive approaches for estimating impact. We observed that in cases when both factual or counterfactual outcomes are available, inductive inference performs better than transductive; specifically,  $f(x_i, 1) - f(x_i, 0)$  is a better estimator of ITE than  $y_i - f(x_i, 0)$  or  $f(x_i, 1) - y_i$ . Inductive inference may reduce noise, approximation error, or bias in the shared layers. Coupling inductive inference with a multitask architecture may be an effective strategy. Moreover, this procedure can be interpreted as an empirical method for removing exogenous and unobserved effects, which may be introduced in data collection or processing, in the outcome.

## 4 Experimental Results

In this section we provide experiments to substantiate the inductive inference of impact estimation. Then, we show the promise of the proposed balance regularization framework.

### 4.1 Experimental Setup

**Datasets.** We use experimental data from the Infant Health and Development Program (IHDP), a randomized experiment that targeted low-birth-weight, premature infants, and provided the treatment group with both intensive high-quality child care and home visits from a trained provider aiming at enhancing their IQ scores at the age of three [23]. The dataset comprises 747 subjects (139 treated and 608 control), with 25 features (covariates) associated with each instance. Outcomes are simulated based on the data generation process [23]. Having simulations of both factual and counterfactual observations allows us to assess the accuracy of impact estimators.

**Measures.** The main quantity that we are interested to estimate is the individual impact  $Y_i^1 - Y_i^0$  and our estimate is  $im(x_i)$ . Recall that  $Y_i^1$  and  $Y_i^0$  are the potential outcomes and  $im(x_i)$  is the estimated impact. The first measure is thus the bias of our estimate,

$$Bias^2 = \mathbb{E}[(Y_i^1 - Y_i^0) - im(x_i)]^2, \quad (10)$$

which would be more informative along with its variance,

$$Var = \mathbb{E}[(im(x_i) - \mathbb{E}[im(x_i)])^2], \quad (11)$$

to express uncertainty in the estimated quantity. The expectation is over the training data, test data, and the randomness involved in the model fitting (e.g. initialization and SGD noise). These two quantities are combined to form a Mean Squared Error (MSE) of the estimate:

$$MSE = Bias^2 + Var. \quad (12)$$

In the causal inference literature the MSE is sometimes replaced by Precision in Estimation of Heterogeneous Effect (PEHE) which is defined as

$$PEHE = \sqrt{\frac{1}{n} \sum_{i=1}^n (im(x_i) - (Y_i^1 - Y_i^0))^2}. \quad (13)$$

It is easy to verify that  $\sqrt{MSE} = PEHE$ .

**Neural network structure and training.** For all models we used Multilayer Perceptron (MLP) architecture. The two head model contains a shared bottom with two hidden layers of 20 neurons and two heads, which each have a 20-neuron hidden layer and a 1-neuron output. For the separate head

<sup>2</sup> or its many variants like ADAM, ADAGRAD, RMSPROP, etc.

model we used two MLP networks, one for fitting the treatment group and one for fitting the control group. They include three hidden layers each with 20 neurons with a one neuron output layer. The input layer has 25 neurons in both cases. All layers except the final layer use ReLU nonlinearities. None of batch norm, dropout, weight decay, data augmentation, and normalization are employed. We used ADAM [24] with a learning rate of 0.0001, a batch size of 64, and 50 epochs for training. Out of  $\approx 750$  examples, 20% are randomly held for the test set. To simulate different training sets we randomly selected 60% of data 5 times and for each we run 4 times with random initialization. Therefore, our results are averaged over 20 random runs with  $\approx 450$  and  $\approx 150$  training and test examples, respectively.

## 4.2 Results

Table 1 shows the average squared bias, average variance, and the mean and standard deviation of the MSE error over 20 iterations of a simulated set of potential outcomes over the IHDP dataset. No regularization is applied for this table. First, we note that incorporating the shared bottom in the two head model improves mean and standard deviation of the MSE. Second, the two head model improves MSE with a significant reduction in the variance and a relatively small increase in bias. Third, inductive inference consistently leads to more accurate impact estimation in both of separate and two-head models compared with transductive inference.

This lower bias of separate head model (in inductive setting) can be explained by the fact that two separate neural networks are more flexible to fit to data compared to the case where their bottom is restricted to be the same. This flexibility comes with the cost of significantly higher variance.

Table 1: Numerical comparison of the two head model with separate head model

Model	Inference Mode	$Bias^2$	$Var$	$MSE$ (average $\pm$ std)
separate head	transductive	6.15	0.72	$6.87 \pm 0.89$
	inductive	2.87	1.00	$3.87 \pm 0.85$
two head	transductive	4.55	0.23	$4.78 \pm 0.19$
	inductive	3.34	0.32	$3.66 \pm 0.23$

Inductive inference forgoes variance within the observed outcome and relies on predictions of both outcomes to estimate impact. To further study the reasons behind its superiority we have plotted the difference between real and predicted treated outcome versus the difference between real and predicted control outcome in Fig. 2. The positive correlation shows that whenever the predicted treated value ( $f^1(x)$ ) is larger than the real treated value ( $y^1$ ), then the predicted control ( $f^0(x)$ ) is larger than the observed control value ( $y^0$ ). Both prediction heads have a bias that is positively correlated and may be canceled out by the inductive difference  $f^1(x) - f^0(x)$ . Further theoretical study of this phenomenon is left for future work. Another expected, yet interesting, observation in this figure is that the predictions on the treated values underestimate the observed treated value while the predictions on the control value overestimate the observed control values. Both models attempt to statistically fit to the data by pushing the predictions towards their target values, and the optimum model converges somewhere in between.

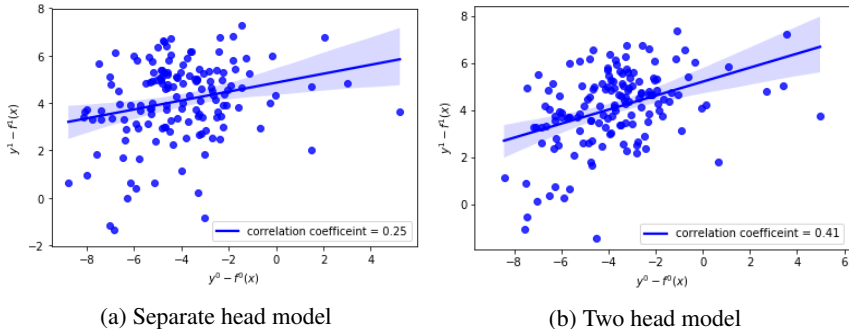


Figure 2: The correlation between real and prediction values of control and test group

Next, we show the effectiveness of balance regularization by increasing the coefficient of MMD and PRG losses from 0 to larger values in the two head model. The estimation error in terms of bias, variance and MSE are shown in Fig. 3. We can see the bias is generally decreasing from top to bottom (increasing the MMD regularization) and from left to right (increasing prognostic regularization) in Fig. 3a. Conversely, the variance of the predictions demonstrated in Fig 3b increases when enforcing more balance regularization. This creates an interesting trade-off between bias and variance that potentially leads to an optimum MSE in the middle of the coefficients grid as shown in Fig 3c. It is notable that the different ranges for MMD and PRG coefficient are due to difference in the magnitude of the associated loss terms. Fig. 3 shows that performance is more sensitive to MMD regularization than PRG.




Figure 3: The bias, variance and MSE of the two head model with respect to varying the regularization weights. The lowest MSE occurs at the suitable trade-off between the bias and variance.

## 5 Related work

Machine learning techniques have already been used in estimating causal effects from data [25, 26, 27, 28, 29]. In addition to aforementioned works [1, 3], the idea of deep representation learning has been embraced by many recent studies such as [30, 31, 32, 33, 34]. These ideas have been extended to leverage propensity score [21, 3] or to work with instrumental variables [35]. Machine learning is contributing to causal effect estimation via many recently developed deep learning techniques such as Generative Adversarial Networks (GANs) [36], Knowledge Distillation [37], Adversarial Training [38], Meta Learning [39], and Transfer Learning [40].

The problem of estimating causal effects is about learning the effect of an intervention, the treatment. Interventions modify the distribution, and here we have two, namely the treated and the untreated distributions. The machine learning field has investigated related settings. Learning with distribution shift [41] considers it generally, and categorizes the distribution shift into different types, such as a shift in the input distribution (covariate shift), or a shift in the output distribution (target shift), as well as other types. Many of these shifts can occur simultaneously in the context of causal effect estimation, and result in confounding bias.

More specific settings include those of domain invariance, domain transfer, and domain adaptation [42]. They involve learning across multiple domains, corresponding to different distributions. For example, they consider how to transfer learning from the distribution of one domain to a different one [43, 44]. Learning domain invariance is perhaps the most closely related to ours: invariant risk minimization [45] attempts to learn a representation that is invariant across domains. This has also been formulated adversarially, where representations are penalized if they differ across domains [46]. Representations can also be made similar across domains via GANs [36]. The area of causal discovery from shifted distributions [47] can also lead to interesting insights about causal effect estimation.

Our setting differs from the above in that we know that the source of the distribution shift is due to the treatment, in particular the treatment assignment, and its effect on outcome. Also, we are focused on predicting the difference in outcomes for different treatments, rather than absolute outcomes. This is what allows pinpointing the distribution shift using techniques from statistics, such as the propensity score and prognostic score [6]. Propensity score has been used as an input feature for causal inference scenarios [23], and also as an auxiliary neural network head [3]. Prognostic score has been used in epidemiology for evaluating the quality of propensity score matching [48], but has not previously been used for regularization, or with neural networks.

## 6 Conclusion and Future Work

We have proposed a deep neural network model for estimating treatment impact in observational studies. Our balance regularization is inspired by matching the treatment and control distributions and is built on the top of representation learning and enjoys the end-to-end training and sample efficiency of gradient based neural network models while being sufficiently flexible to model complex functions. We acknowledge that our empirical experiments have been conducted on a relatively small sample size. Extending this work to larger datasets and systematically studying the effects of balance regularizers combined with other regularizers (like dropout and weight decay) is left as future work. One can potentially make the prognostic loss more symmetric by enforcing it on the treatment head too. Another interesting line of work is combining these losses with treatment assignment prediction on as a third head [3].

## References

- [1] Fredrik Johansson, Uri Shalit, and David Sontag. Learning representations for counterfactual inference. In *International conference on machine learning*, pages 3020–3029, 2016.
- [2] 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.
- [3] Claudia Shi, David Blei, and Victor Veitch. Adapting neural networks for the estimation of treatment effects. In *Advances in Neural Information Processing Systems*, pages 2503–2513, 2019.
- [4] Donald B Rubin. Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, 100(469):322–331, 2005.
- [5] Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(Mar):723–773, 2012.
- [6] Ben B. Hansen. The prognostic analogue of the propensity score. *Biometrika*, 95(2):481–488, June 2008.
- [7] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.
- [8] Elias Bareinboim, Jin Tian, and Judea Pearl. Recovering from selection bias in causal and statistical inference. In *AAAI*, pages 2410–2416. Citeseer, 2014.
- [9] Juan Correa, Jin Tian, and Elias Bareinboim. Generalized adjustment under confounding and selection biases. In *Proceedings of the... AAAI Conference on Artificial Intelligence*, 2018.
- [10] Guido W Imbens and Donald B Rubin. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press, 2015.
- [11] Judea Pearl. *Causality*. Cambridge university press, 2009.
- [12] 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.
- [13] Alberto Abadie and Guido W Imbens. Large sample properties of matching estimators for average treatment effects. *econometrica*, 74(1):235–267, 2006.
- [14] Sheng Li, Nikos Vlassis, Jaya Kawale, and Yun Fu. Matching via dimensionality reduction for estimation of treatment effects in digital marketing campaigns. In *IJCAI*, pages 3768–3774, 2016.
- [15] Chen Wang, Suzhen Wang, Fuyan Shi, and Zaixiang Wang. Robust propensity score computation method based on machine learning with label-corrupted data. *arXiv preprint arXiv:1801.03132*, 2018.

- [16] Suriya Gunasekar, Jason D Lee, Daniel Soudry, and Nati Srebro. Implicit bias of gradient descent on linear convolutional networks. In *Advances in Neural Information Processing Systems*, pages 9461–9471, 2018.
- [17] Lenaic Chizat and Francis Bach. Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss. *arXiv preprint arXiv:2002.04486*, 2020.
- [18] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [19] Quoc V Le. Building high-level features using large scale unsupervised learning. In *2013 IEEE international conference on acoustics, speech and signal processing*, pages 8595–8598. IEEE, 2013.
- [20] Sebastian Ruder. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*, 2017.
- [21] Ahmed M Alaa, Michael Weisz, and Mihaela Van Der Schaar. Deep counterfactual networks with propensity-dropout. *arXiv preprint arXiv:1706.05966*, 2017.
- [22] Christos Louizos, Uri Shalit, Joris M Mooij, David Sontag, Richard Zemel, and Max Welling. Causal effect inference with deep latent-variable models. In *Advances in Neural Information Processing Systems*, pages 6446–6456, 2017.
- [23] Jennifer L Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011.
- [24] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [25] Susan Athey and Guido W Imbens. Machine learning methods for estimating heterogeneous causal effects. *stat*, 1050(5):1–26, 2015.
- [26] Yonghan Jung, Jin Tian, and Elias Bareinboim. Learning causal effects via weighted empirical risk minimization. 2020.
- [27] Daniel Malinsky and Peter Spirtes. Estimating bounds on causal effects in high-dimensional and possibly confounded systems. *International Journal of Approximate Reasoning*, 88:371–384, 2017.
- [28] Ariel Linden and Paul R Yarnold. Combining machine learning and propensity score weighting to estimate causal effects in multivalued treatments. *Journal of Evaluation in Clinical Practice*, 22(6):875–885, 2016.
- [29] Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference*. The MIT Press, 2017.
- [30] Sheng Li and Yun Fu. Matching on balanced nonlinear representations for treatment effects estimation. In *Advances in Neural Information Processing Systems*, pages 929–939, 2017.
- [31] 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.
- [32] Patrick Schwab, Lorenz Linhardt, and Walter Karlen. Perfect match: A simple method for learning representations for counterfactual inference with neural networks. *arXiv preprint arXiv:1810.00656*, 2018.
- [33] Peipei Chen, Wei Dong, Xudong Lu, Uzay Kaymak, Kunlun He, and Zhengxing Huang. Deep representation learning for individualized treatment effect estimation using electronic health records. *Journal of biomedical informatics*, 100:103303, 2019.



- [34] Patrick Schwab, Lorenz Linhardt, Stefan Bauer, Joachim M Buhmann, and Walter Karlen. Learning counterfactual representations for estimating individual dose-response curves. In *AAAI*, pages 5612–5619, 2020.
- [35] Jason Hartford, Greg Lewis, Kevin Leyton-Brown, and Matt Taddy. Deep iv: A flexible approach for counterfactual prediction. In *International Conference on Machine Learning*, pages 1414–1423, 2017.
- [36] Jinsung Yoon, James Jordon, and Mihaela van der Schaar. GANITE: Estimation of Individualized Treatment Effects using Generative Adversarial Nets. In *International Conference on Learning Representations*, February 2018.
- [37] Maggie Makar, Adith Swaminathan, and Emre Kıcıman. A distillation approach to data efficient individual treatment effect estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 4544–4551, 2019.
- [38] Xin Du, Lei Sun, Wouter Duivesteijn, Alexander Nikolaev, and Mykola Pechenizkiy. Adversarial balancing-based representation learning for causal effect inference with observational data. *arXiv preprint arXiv:1904.13335*, 2019.
- [39] Ankit Sharma, Garima Gupta, Ranjitha Prasad, Arnab Chatterjee, Lovekesh Vig, and Gautam Shroff. Metaci: Meta-learning for causal inference in a heterogeneous population. *arXiv preprint arXiv:1912.03960*, 2019.
- [40] Sören R Künzel, Bradley C Stadie, Nikita Vemuri, Varsha Ramakrishnan, Jasjeet S Sekhon, and Pieter Abbeel. Transfer learning for estimating causal effects using neural networks. *arXiv preprint arXiv:1808.07804*, 2018.
- [41] Joaquin Quinero-Candela, Masashi Sugiyama, Anton Schwaighofer, and Neil D. Lawrence. *Dataset Shift in Machine Learning*. The MIT Press, 2009.
- [42] Kun Zhang, Mingming Gong, and Bernhard Schölkopf. Multi-source domain adaptation: A causal view. In *AAAI*, volume 1, pages 3150–3157, 2015.
- [43] Sara Magliacane, Thijs van Ommen, Tom Claassen, Stephan Bongers, Philip Versteeg, and Joris M Mooij. Domain adaptation by using causal inference to predict invariant conditional distributions. In *Advances in Neural Information Processing Systems*, pages 10846–10856, 2018.
- [44] Kun Zhang, Mingming Gong, Petar Stojanov, Biwei Huang, and Clark Glymour. Domain adaptation as a problem of inference on graphical models. *arXiv preprint arXiv:2002.03278*, 2020.
- [45] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant Risk Minimization. *arXiv:1907.02893 [cs, stat]*, July 2019.
- [46] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario March, and Victor Lempitsky. Domain-Adversarial Training of Neural Networks. *Journal of Machine Learning Research*, 17(59):1–35, 2016.
- [47] Biwei Huang, Kun Zhang, Jiji Zhang, Joseph Ramsey, Ruben Sanchez-Romero, Clark Glymour, and Bernhard Schölkopf. Causal discovery from heterogeneous/nonstationary data. *Journal of Machine Learning Research*, 21(89):1–53, 2020.
- [48] Elizabeth A Stuart, Brian K Lee, and Finbarr P Leacy. Prognostic score-based balance measures can be a useful diagnostic for propensity score methods in comparative effectiveness research. *Journal of clinical epidemiology*, 66(8):S84–S90, 2013.