Deep Learning of Potential Outcomes

**Bernard Koch¹, Tim Sainburg², Pablo Geraldo Bastias¹
Song Jiang³, Yizhou Sun³, Jacob Foster¹**

1. UCLA Dept. of Sociology, 2. UCSD Dept. of Psychology, 3. UCLA Dept. of Computer Science

Abstract

This review systematizes the emerging literature for causal inference using deep neural networks under the potential outcomes framework. It provides an intuitive introduction on how deep learning can be used to estimate/predict heterogeneous treatment effects and extend causal inference to settings where confounding is non-linear, time varying, or encoded in text, networks, and images. To maximize accessibility, we also introduce prerequisite concepts from causal inference and deep learning. The survey differs from other treatments of deep learning and causal inference in its sharp focus on observational causal estimation, its extended exposition of key algorithms, and its detailed tutorials for implementing, training, and selecting among deep estimators in Tensorflow 2 available at github.com/kochbj/Deep-Learning-for-Causal-Inference.

Keywords: deep learning, causal inference, potential outcomes, machine learning.

Contents

1	Introduction	4
1.1	Why use deep learning for causal inference?	5
2	Primer on Deep Learning	7
2.1	Artificial Neural Networks	7
2.2	Representation Learning and Multitask Learning	11
3	Causal Identification and Estimation Strategies	12
3.1	Identification of Causal Effects	12
3.2	Estimation of Causal Effects	16
3.2.1	Outcome Modeling	17
3.2.2	Non-Parametric Matching	18
3.2.3	Treatment Modeling	18
4	Four Different Approaches to Deep Causal Estimation	20
4.1	Deep Outcome Modeling	20
4.2	Balancing through Representation Learning	22
4.2.1	Extending Representation Balancing with IPMs	23
4.2.2	Extending Representation Balancing with Matching	26
4.3	Extensions with Inverse Propensity Weighting	26
4.3.0	Treatment Modeling with Dragonnet	27
4.4	Adversarial Training of Generative Models, Representations, IPW	31
4.4.1	The Origins of Adversarial Training in GANs	31
4.4.2	GANs as Generative Models of Treatment Effect Distributions	32
4.4.3	Adversarial Representation Balancing	33
5	Extending Causal Estimation to Non-tabular Data	35
5.1	Conditioning on Time-Varying Confounding	35
5.2	Relaxing Strong Ignorability: Controlling for Latent Confounders in Text, Graphs, and Images	36
5.2.1	Conditioning on Latent Confounding Encoded in Networks	36
5.2.2	Conditioning on Latent Confounding in Text Data	38
5.2.3	Causal Inference on Images	40
6	Conclusion	40

Boxes

1	Box 1: Reading Machine Learning Papers: Computational Graphs and Loss Functions	8
2	Box 2: Training and Regularizing Supervised Deep Learning Models	9
3	Box 3: Basic Introduction to Causal Inference	13
4	Box 4: Notation for Causal Inference and Estimation	15
5	Box 5: Other Flavors of TarNet	30
6	Box 6: Generative Adversarial Networks (GAN)	31
7	Box 7: Recurrent Neural Networks (RNN)	36
8	Box 8: Graph Neural Networks (GNN)	37
9	Box 9: Transformers	39

1. Introduction

In this paper, we systematize the emerging literature for estimating causal effects using deep neural networks within the potential outcomes framework. In recent years, both causal inference frameworks and deep learning have seen rapid adoption across science, industry, and medicine. Causal inference has a long tradition in the social sciences (for a basic introduction, see Box 3), but deep learning (and machine learning more generally) is conspicuously underutilized. This review aims to introduce social and data scientist readers to an exciting literature within the machine learning community exploring how deep learning might be used to estimate causal effects. Because this literature is growing rapidly, we organize proposed deep causal estimators into four basic categories that reflect the causal estimation strategies employed. We assume the reader has limited familiarity with causal inference and neural networks, so key concepts from both paradigms are introduced throughout the paper. To streamline the experience for more advanced readers, these concepts are contained to boxes.

The review is organized as follows. We first provide a brief, intuitive primer on deep learning and representation learning, and then review assumptions needed for causal identification when applying deep learning models for causal estimation under a selection on observables strategy. To motivate the typology used to organize deep learning models, we then discuss three distinct approaches to causal identification in the selection on observables setting: outcome modeling, treatment modeling through non-parametric matching, and treatment modeling using propensity scores.

In the main body of the paper, we organize deep learning algorithms into four basic categories, three of which correspond with the approaches above. First, we discuss the use of neural networks as plug-in *outcome modelers* for conditional average treatment effects, a ubiquitous technique in this literature that is generally combined with other approaches. Second, we explain how representation learning can be used to *balance* covariate distributions. Third we discuss the usage of neural networks to generate *inverse propensity score weights* (IPW). The fourth section describes *adversarial training* regimes inspired by Generative Adversarial Networks (GANs) to build generative models of counterfactual outcome distributions or improve

performance of the above-described techniques (Goodfellow et al. 2014). For each of these strategies, we take “Deep Dives” into a representative algorithm.

In the final section of the paper, we discuss models that extends estimation under selection on observables to complex settings: data with time-varying treatments or scenarios where confounders, mediators, and treatments might be latently represented in graphs, text, or images. We conclude by presenting the pros and cons of these deep causal estimators compared to established approaches in the social sciences. To aid researchers in implementing these approaches themselves, we provide extensive tutorials in Tensorflow 2 available at <https://github.com/kochbj/Deep-Learning-for-Causal-Inference>. Like the review, these tutorials assume no past experience with Tensorflow.

1.1. Why use deep learning for causal inference?

Deep learning estimators present several advantages compared to existing linear and machine learning causal estimators in scientists’ arsenal:

- **(Nearly) non-parametric modeling of relationships between covariates, treatments, and outcomes.** Using generalized linear models for causal estimation requires the analyst to make strong assumptions about the functional relationship between observed covariates (outcome predictors, confounders, mediators, colliders), treatment assignment, and outcomes. Machine learning estimators relax these parametric assumptions by exhaustively exploring non-linear interactions that may correlate covariates, treatment, and outcomes. In particular, neural networks have two distinct advantages compared to other machine learning approaches to causal inference (e.g. decision tree/forest-based approaches, LASSO regression, support vector machines). First, neural networks naturally extract relevant information from covariates through representation learning (discussed extensively below), allowing the analyst to potentially incorporate dozens or hundreds of observed variables that predict/confound treatment assignment and outcome. Second, the flexibility of neural networks allow analysts to extend nearly non-parametric estimation to scenarios where not many viable non-parametric estimation

strategies exist (e.g., observational data with time varying treatments, data with multiple treatments).

- **State of the art estimation of heterogeneous treatment effects.** A recent trend in causal inference has been to focus on how heterogeneous treatment effects vary across subpopulations. The emergence of machine learning causal estimators has been a driving force behind this trend. Because deep neural networks can theoretically approximate any continuous function, neural networks appear to substantially outperform other machine learning approaches to causal inference for the estimation of heterogeneous/conditional treatment effects with respect to bias, in both simulated and real data.
- **Moving from inference to prediction.** The deep causal estimators surveyed below are designed not just for in-sample inference, but also for out-of-sample prediction. Predictive modeling would allow social scientists to train a model on observational data where treatment of some units is observed, and estimate effects in new datasets where treatment assignment is unobserved.
- **Causal inference in quantitative data, text, images, and graphs.** Through representation learning, deep neural network models can adjust for confounding not just in quantitative data, but also extract latent confounders encoded in text, networks, and graphs. To motivate the use of these models, we discuss some example causal scenarios that can be addressed by using deep neural estimators:

Traditional Data. The companion tutorials use a naturalistic simulation based on the Infant Health and Development Program (IHDP) example from [Hill \(2011\)](#). One of the goals of the original IHDP study was to estimate the causal effect of specialized childcare interventions on cognitive outcomes for premature infants. The treatment (T) is attendance at a special child development center for premature infants. The outcome is some measure of cognitive development for infants after (Y). Measured covariates (X) such as socioeconomic status are predictive of both seeking treatment and cognitive development.

Text. As a motivating example, [Veitch et al. \(2019a\)](#) consider the effect of the

author’s reported gender (T) on the number of upvotes a Reddit post receives (Y). However gender may also “affect the text of the post, e.g., through tone, style, or topic choices, which also affects its score $[(X)]$.” Controlling for a representation of the text would allow the analyst to more accurately estimate the direct effect of gender.

Images. Todorov et al. (2005) showed that split second-judgments of a politician’s competence (T) from pictures (X) of their face is predictive of their electability (Y). When attempting to replicate this study using machine learning classifiers rather than human classifiers, Joo et al. (2015) suggest that the age of the face (Z) is a not-so-obvious confounder: while older individuals are more likely to appear competent, they are also more likely to be incumbents. Even if age is unknown, using neural networks to control for confounders implicitly encoded in the image (like age) could reduce bias.

Networks. Nagpal et al. (2020) explore the question of which types of prescription opioids (e.g., natural, semi-synthetic, synthetic) (T) are most likely to cause long term addiction (Y). Because of predisposition to different injuries, type of employment (X) could be a common cause of both treatment and outcome. Suppose job type is unobserved, but we know that patients are likely to associate with coworkers through homophily. To capture some of the effects of this latent unobserved confounder, analysts might choose to control for a representation of the patient’s position in their social network when estimating the causal effect.

While the main body of this review focuses on algorithm for causal inference/prediction from traditional quantitative data, models for dealing with non-traditional data are discussed in Section 5.

2. Primer on Deep Learning

2.1. Artificial Neural Networks

Artificial neural networks (ANN) are statistical models inspired by the human brain (Brand et al. 2020). In an ANN, each “neuron” in the network takes the weighted sum of its inputs (the outputs of other neurons) and transforms them using a twice differentiable, non-linear function (e.g. sigmoid, rectified linear unit) that outputs a value between 0 and 1 if the transformed value is above some threshold. Neurons are arrayed in layers where an input layer takes the raw data, and each neuron in subsequent layers take the weighted sum of outputs in previous layers as input. An “output” layer contains a neuron for each of the predicted outcomes with transformation functions appropriate to those outcomes. For example, a regression network that predicts one outcome will have a single output neuron without a transformation function so that it produces a real number. A regression network without any hidden layers corresponds exactly to a generalized linear model (Fig. 1A). When additional “hidden” layers are added between the input and output layers, the architecture is called a **feed-forward network** or **multi-layer perceptron** (Fig. 1B). A neural network with multiple hidden layers is called a “deep” network, hence the name “deep learning” (LeCun et al. 2015). A neural network with a single, large enough hidden layer can theoretically approximate any continuous function (Cybenko 1989).

Box 1: Reading Machine Learning Papers: Computational Graphs and Loss Functions

Within the machine learning literature, novel algorithms are often presented in terms of their computational graph and loss function. A computational graph (not to be confused with a causal graph) uses arrows to depicts the flow of data from the inputs of a neural network, through parameters, to the outputs. Layers of neurons or specialized sub-architectures are often generically abstracted as shapes. In our diagrams, we use purple to represent observables, orange for representation layers of the network, white for produced outputs, and red and blue for outcome modeling layers. Operations that are computed *after* prediction (i.e., for which an error gradient is not calculated) are shown with dashed lines (e.g., plug-in estimation of causal estimands).

Along with the architecture, the loss function of a neural network is the primary means for the analyst to dictate what types of representations a neural network learns and what types of outputs it produces. In multi-task learning settings, we denote joint loss functions for an entire network as a weighted sum of the losses for substituent tasks and modules. These specific losses are weighted by hyperparameters. For example, we might weight the joint loss for a network that predicts outcomes and propensity scores as:

$$\arg \min_{h, \pi} \mathcal{L} = \mathcal{L}_h + \lambda \mathcal{L}_\pi = \text{MSE}(Y, \hat{Y}) + \lambda \text{BCE}(T, \hat{\pi}(X, T))$$

where $\hat{\pi}(X, T)$ is the predicted propensity score, λ is a hyperparameter and MSE and BCE stand for mean squared error and binary cross entropy (i.e., log loss), common losses for regression and binary classification respectively.

Neural networks are trained to predict their outcomes by optimizing a **loss function** (also called objective or cost function). During training, error in the loss function is distributed proportionally (i.e. **backpropagated**) to weight parameters in previous layers in the network. An optimizer, such as the stochastic gradient descent algorithm or the currently popular ADAM algorithm (Kingma and Ba 2014), then moves the parameters in the opposite direction of this error gradient. Neural networks first rose to popularity in the 1980s but fell out of favor compared to other machine learning model families (e.g., support vector machines) due to their expense of training. By the late 2000s, the improvements to backpropagation, advances in graphical processing units (i.e., graphic cards), and access to larger datasets, collectively enabled a deep learning revolution where ANNs began to significantly outperform other model families. Today, deep learning is the hegemonic machine learning approach in industries and fields other than social science. For further discussion on how deep learning models are trained and regularized in a supervised machine learning framework, see Box 2.

Box 2: Training and Regularizing Supervised Deep Learning Models

Supervised Training. In supervised learning, data is split into a training set and a validation set. Model parameters are optimized on the training set before out-of-sample performance is assessed on the validation set. Performance on the validation set is typically used to choose hyperparameter settings. Deep learning models differ from other machine learning approaches in that the loss function is typically non-convex and trained models may not converge on the same optima. Thus unlike other machine learning approaches which train first on the complete training set and are then evaluated subsequently on the complete validation set, neural networks are typically trained on small batches of training data at a time. Because a batch of data is only a sample of a sample (the training dataset), the optimizer only adjusts weight parameters by a fraction of the error gradient (the **learning rate**) to avoid overfitting. When a model has cycled through a set of batches that cover the complete training set, this is called a training **epoch**. After each training epoch, the network is typically tested over a validation epoch (i.e. a complete iteration of batches for the validation set) without updating the weights.

Figure 1: **A: Generalized linear model represented as a computational graph.** Observable covariates X_1, X_2, X_3 and treatment status T depicted in purple. Each of the lines between the purple inputs and the orange box represents a parameter (i.e., a β in a generalized linear model equation). The orange box is an “output neuron” that sums its weighted inputs, performs a transformation g (the link function in GLM; in this case the identity function), and predicts the conditional outcome $\hat{Y}(T)$. Instead of theoretically interpreting these parameters from an inferential statistics perspective, machine learning approaches typically use the predicted observed and unobserved potential outcomes for plug-in estimation of causal estimands (e.g., the conditional average treatment effect $C\hat{ATE}$). After training, setting T to $1 - T$ for each observation can predict the unobserved potential outcome $\hat{Y}(1 - T)$. Because this operation occurs after prediction and does not feed a gradient back to the network to optimize the parameters, it is depicted here with a dotted line. Plug-in calculation of $C\hat{ATE}$ is similarly shown with a dotted line.

B: Feed-forward neural network. In a feed-forward neural network, additional fully connected (parameterized) layers of neurons are added between the inputs and output neuron. The size of the input covariates and hidden layers are generically abstracted as boxes. The final hidden layer before the output neuron is denoted Φ because the hidden layers collectively encode a representation function (see section 2.2). In causal inference settings, this architecture is sometimes called a S(ingle)-learner because one feed-forward network learns to predict both potential outcomes.

Regularization. Neural networks are highly susceptible to overfitting, and early stopping of training once the validation error begins to rise is a fundamental regularization technique. Other common regularization techniques include **weight decay** (i.e., ℓ^2 norm, ridge, or Tikhonov) penalties on the weight parameters, dropout of neurons during training, and batch normalization. **Dropout** is a regularization technique in deep learning where certain nodes are randomly “dropped out” from training during a given epoch [Srivastava et al. \(2014\)](#). The general idea of dropout is to force two neurons in the same layer to learn different aspects of the covariate/feature space and reduce overfitting. **Batch normalization** is another regularization technique applied to a layer of neurons ([Ioffe and Szegedy 2015](#)). By standardizing (i.e. z-scoring) the inputs to a layer on a per-batch basis and then rescaling them using trainable parameters, batch normalization smooths the optimization of the loss function.

2.2. Representation Learning and Multitask Learning

One comparative advantage of deep learning over other machine learning approaches has been the ability of ANNs to encode and automatically compress informative features from complex data into flexible, relevant “**representations**” or “embeddings” that make downstream supervised learning tasks easier ([Goodfellow et al. 2016](#); [Bengio 2013](#)). While other machine learning approaches may also encode representations, they often require extensive pre-processing to create useful features for the algorithm. Through the lens of representation learning, a geometric interpretation of the role of each layer in a supervised neural network is to transform its inputs (either raw data or output of previous layers) into a typically lower (but possibly higher) dimensional vector space. As a means to share statistical power, encoded representations can also be jointly learned for two tasks at once in **multi-task learning**.

The simplest example of a representation might be the final layer in a feed-forward network, where the early layers of the network can be understood as non-linearly encoding the inputs into an array of latent linear features for the output neuron ([Goodfellow et al. 2016](#)) (Fig. 1B). A famous example of representation learning is the use of neural networks for face detection. Examining the representations produced by each layer of these networks shows that subsequent layer seems to capture increasingly abstract features of a face (first edges, then noses and eyes, and finally whole faces) ([LeCun et al. 2015](#)). A more familiar example of representation learning to social sciences might be word vector models like Word2Vec ([Mikolov](#)

Figure 2: **Balancing through representation learning.** The promise of deep learning for causal inference is that a neural network encoding function Φ can transform the treated and control covariate distributions into a representation space such that they are indistinguishable. Used with permission from [Johansson and Shen \(2018\)](#).

[et al. 2013](#)). Word2Vec is a two-layer neural network where words that are semantically similar are closer together in the representation produced by the hidden layer of the network.

The novel contribution of deep learning to causal estimation is the proposal that a neural network can learn a function Φ that produces representations of the covariates decorrelated from the treatment. Fundamentally, the idea is that Φ can transform the treated and control covariate distributions into a representation space such that they are indistinguishable (Fig. 2). To ensure that these representations are also still predictive of the outcome (multi-task learning), multiple loss functions are generally applied simultaneously to balance these objectives. This approach is applied in a majority of the algorithms presented in the main body of this review (Section 4).

3. Causal Identification and Estimation Strategies

3.1. Identification of Causal Effects

The papers described in this review are primarily framed within the Potential Outcomes

causal framework (Neyman-Rubin causal model) (Rubin 1974; Imbens and Rubin 2015). This framework is concerned with identifying the “potential outcomes” of each unit i in the sample, had they both received treatment ($Y(1)$) and not received treatment ($Y(0)$). However, because each unit can only receive one treatment regime in reality (being treated or remaining untreated), it is not possible to observe both potential outcomes for each individual (often termed “the fundamental problem of causal inference” (Holland 1986)). While we cannot thus identify individual treatment effects $\tau_i = Y_i(1) - Y_i(0)$ for each unit, causal inference frameworks allow us to probabilistically estimate average treatment effects (ATE) and average treatment effects conditional on select covariates ($CATE$) across samples of treated and control units. Within this literature, the motivation of many papers is to present algorithms that can both infer CATEs from observational data, but also predict them for out-of-sample units where treatment status is unknown.

Box 3: Basic Introduction to Causal Inference

Correlation does not equal causation, and causal statistics is concerned with the identification of causal relationships between random variables. Many causal questions we would like to ask about social data can be framed as counterfactual questions with the general format: “What would have been the outcome Y for a unit with X characteristics, if T had happened or not happened?” Equivalently, this can be reworded to “What is the causal effect of T on Y for units with characteristics X .”

Causal inference frameworks usually take randomized control trials (RCTs, also known as A/B testing in data science and industry applications), where each unit with covariate or features X is randomly assigned to the treatment or control groups and outcome Y is subsequently measured, as the ideal approach to answering this type of question. But in many scenarios it is prohibitively expensive or unethical (e.g., randomly assigning students to attend college or not) to collect experimental data. In these cases, we can statistically adjust observational data (e.g., survey data on college attendance) to approximate the experimental ideal. The methods described in this paper are designed to answer counterfactual questions with primarily non-experimental observational data.

There are at least three different schools of causal inference that have been introduced in social statistics (Rubin 1974; Imbens and Rubin 2015), epidemiology (Robins 1986, 1987; Hernán and Robins 2020), and computer science (Goldschmidt and Pearl 1996; Pearl 2009). The goal of these causal frameworks is to describe and correct for biases in data or study design that would prevent one from making a true causal claim. If these biases are correctable and the causal effect can be uniquely expressed in terms of the distribution of observed data, then we say that the causal effect is *identifiable* (Kennedy 2016). If a causal effect is identifiable, we can use statistical tools that cor-

rect for identified biases to *estimate* the causal effect (e.g., inverse propensity score weighting, g-computation, deep learning).

The focus of the algorithms presented in this paper is on estimating causal effects while correcting for *confounding bias*. Loosely speaking, a confounding covariate/feature is one that is correlated with both the treatment and the outcome, misleadingly suggesting that the treatment has a causal effect on the outcome, or obscuring a true causal relationship between the treatment and outcome. Often times, the confounder is a cause of the treatment *and* outcome. As an example, estimating the causal effect of attending college (treatment) on adult income (outcome) requires controlling for the fact that parental income may be a common cause of both college attendance and adult income.

The ATE is defined as:

$$ATE = \mathbb{E}[Y_i(1) - Y_i(0)] = \mathbb{E}[\tau_i]$$

where $Y(1)$ and $Y(0)$ are the potential outcomes had the unit i received or not received the treatment, respectively. The CATE is defined as,

$$CATE = \mathbb{E}[Y_i(1) - Y_i(0)|X = x] = \mathbb{E}[\tau_i|X = x]$$

where X is the set of selected, observable covariates, and $x \in X$.

Within the machine learning literature on causal inference surveyed here, the primary strategy for causal identification is **selection on observables**. A challenge to identifying causal effects is the presence of confounding relationships between covariates associated with both the treatment and the outcome. The key assumption allowing the identification of causal effects in the presence of confounders is:

1. **Conditional Ignorability/Exchangability** The potential outcomes $Y(0)$, $Y(1)$ and the treatment T are conditionally independent given X ,

$$Y(0), Y(1) \perp\!\!\!\perp T | X$$

Conditional Ignorability specifies that there are no unmeasured confounders that affect both treatment and outcome outside of those in the observed covariates/features X . Additionally

X may contain predictors of the outcome, but should not contain instrumental variables or colliders within the conditioning set.

Other standard assumptions invoked to justify causal identification are:

2. **Consistency/Stable Unit Treatment Value Assumption (SUTVA)**. Consistency specifies that when a unit receives treatment, their observed outcome is exactly the corresponding potential outcome (and the same goes for the outcomes under the control condition). Moreover, the response of any unit does not vary with the treatment assignment to other units (i.e., no network or spillover effects), and the form/level of treatment is homogeneous and consistent across units (no multiple versions of the treatment). More formally,

$$T = t \rightarrow Y = Y(T)$$

3. **Overlap**. For all $x \in X$ (i.e., any observed covariate value), all treatments $t \in \{0, 1\}$ have a non-zero probability of being observed in the data, within the “strata” defined by such covariates,

$$\forall x \in X, t \in \{0, 1\} : p(T = t | X = x) > 0$$

4. An additional assumption often invoked at the interface of identification and estimation using neural networks is:

Invertability

$$\Phi^{-1}(\Phi(X)) = X$$

In words, there must exist an inverse function of the representation function Φ encoded by a neural network that can reproduce X from representation space. This is required for the Conditional Ignorability assumption to hold when using representation learning.

For reference, we describe the full notation used within the review in Box 4.

Box 4: Notation for Causal Inference and Estimation

We use uppercase to denote general quantities (e.g., random variables) and lowercase to denote specific quantities for individual units (e.g., observed variable values).

Causal identification

- Observed covariates/features: X
- Potential outcomes: $Y(0)$ and $Y(1)$
- Treatment: T
- Unobservable Individual Treatment Effect: $\tau_i = Y_i(1) - Y_i(0)$
- Average Treatment Effect: $ATE = \mathbb{E}[Y_i(1) - Y_i(0)] = \mathbb{E}[\tau_i]$
- Conditional Average Treatment Effect: $CATE(x) = \mathbb{E}[Y_i(1) - Y_i(0)|X = x] = \mathbb{E}[\tau_i|X = x]$

Deep learning estimation

- Predicted potential outcomes: $\hat{Y}(0)$ and $\hat{Y}(1)$
- Outcome modeling functions: $\hat{Y}(T) = h(X, T)$
- Propensity score function: $\pi(X, T) = P(T|X)$ (where $\pi(X, 0) = 1 - \pi(X, 1)$)
- Representation functions: $\Phi(X)$ (producing representations ϕ)
- Loss functions: $\mathcal{L}(true, predicted)$
- Loss abbreviations: MSE (mean squared error), BCE (binary cross-entropy), CCE (categorical cross-entropy)
- Loss hyperparameters: λ, α, β
- Estimated CATE*: $CATE_i = \hat{\tau}_i = \hat{Y}_i(1) - \hat{Y}_i(0) = h(X, 1) - h(X, 0)$
- Estimated ATE: $\hat{ATE} = \frac{1}{N} \sum_{i=1}^N \hat{\tau}_i$

Beyond the ATE and $CATE$ there is an additional metric commonly used in the machine learning literature, first introduced by Hill (2011) called the *Precision in Estimated Heterogeneous Effects (PEHE)*. PEHE is the average error across the predicted $CATE$ s.

- Precision in Estimated Heterogeneous Effects: $PEHE = \frac{1}{N} \sum_{i=1}^N (\tau_i - \hat{\tau}_i)^2$

Beyond being a metric for simulations with known counterfactuals, the *PEHE* has theoretical significance in the formulation of generalization bounds within this literature Shalit et al. (2017); Johansson et al. (2018, 2020); Zhang et al. (2020).

*Note that we use $\hat{\tau}$ to refer to the estimated CATE because truly individual treatment effects cannot be described only by the observed covariates X .

3.2. Estimation of Causal Effects

Once a strategy for isolating causal effects from available data has been developed (arguably the harder and more important part of causal inference), statistical methods can be used to estimate causal effects by controlling for confounding bias. There are two fundamental approaches to estimation: “treatment modeling” to control for correlations between the co-

Figure 3: **Two fundamental approaches to deconfounding.** Blunted arrows indicate blocked causal paths. Treatment modeling approaches like inverse propensity weighting, balancing, and representation learning adjust for correlations between the covariates X and the treatment T . Outcome modeling approaches like generalized linear models or machine learning regressors adjust for correlations between X and the outcome Y .

variates X and the treatment T , and “outcome modeling” to control for correlations between the treatment X and the outcome Y (Fig. 3). Below we briefly review three traditional techniques for removing confounding bias to motivate our systematization of deep learning models. First, we discuss outcome modeling through regression. Next, we consider treatment modeling through non-parametric matching. Finally, we discuss treatment modeling through inverse propensity score weighting (IPW).

Outcome Modeling: Regression

Assuming the treatment effect is constant across covariates/features or the probability of treatment is constant across all covariates/features (both improbable assumptions), the simplest consistent approach to estimating the ATE is to regress the outcome on the treatment indicator and covariates using a linear model.¹ The ATE is then the coefficient of the treatment indicator. Without loss of generality, we call outcome models of this nature, linear or non-linear, h :

¹Another outcome modeling approach that could be used to estimate the outcome, not discussed here, is g-computation (Robins 1986; Hernán and Robins 2020).

$$\hat{Y}(T) = h(X, T)$$

A slightly more sophisticated semi-parametric approach to **outcome modeling**, used widely in the application of machine learning to causal inference, is to use $h(X, T)$ to impute $\hat{Y}(1)$ and $\hat{Y}(0)$, and calculate the CATE for each unit as a plug-in estimator:

$$CATE_i = \hat{\tau}_i = Y_i(1) - Y_i(0) = h(X_i, 1) - h(X_i, 0)$$

and the ATE as:

$$ATE = \frac{1}{N} \sum_{i=1}^N \hat{\tau}_i$$

Treatment Modeling: Non-Parametric Matching

Another common approach is balancing the treated and control covariate distributions through matching. Matching requires the analyst to select a distance measure that captures the difference in observed covariate distributions between a treated and untreated unit ([Austin 2011](#)). Units with treatment status T can then be matched with one or more counterparts with treatment status $1 - T$ using a variety of algorithms [Stuart \(2010\)](#). In a one-to-one matching scenario where each treated unit has an otherwise identical untreated counterpart, the covariate distribution of treated and control units is indistinguishable.

Treatment Modeling: Inverse Propensity Score Weighting

A common treatment-modeling strategy is **inverse propensity score weighting (IPW)**. In IPW, units are weighted on their inverse propensity to receive treatment. Without loss of generality, we call the propensity function π . The propensity score is calculated as the probability of receiving treatment conditional on covariates:

$$\pi(X) = P(T|X)$$

The simplest IPW estimator of the ATE is then:

$$A\hat{T}E = \frac{1}{N} \sum_{i=1}^N \left\{ \frac{T_i Y_i}{\pi(X_i, T_i)} + \frac{(1 - T_i) Y_i}{\pi(X_i, T_i)} \right\} \quad (1)$$

Note that only one of the two terms is active for any given unit. Furthermore, this presentation looks different than how the IPW is generally presented because we use π as a function with different outputs depending on the value of T rather than a scalar (Box 4).

To de-emphasize the contribution of units with extreme weights due to sparse data, sometimes the stabilized IPW is used:

$$A\hat{T}E = \frac{1}{N} \sum_{i=1}^N \left\{ \frac{T_i \cdot Y_i \cdot P(T_i = 1)}{\pi(X_i, T_i)} + \frac{(1 - T_i) \cdot Y_i \cdot P(T_i = 0)}{\pi(X_i, T_i)} \right\} \quad (2)$$

IPW weighting is attractive because if the propensity score π is specified correctly, it is an unbiased estimator of the ATE. Moreover, the IPW is consistent if π is estimated consistently (Rosenbaum and Rubin 1983; Glynn and Quinn 2010).

Double Robustness

Because different models make different assumptions, it is not uncommon to combine outcome modeling with propensity modeling or matching estimators to create **doubly-robust** estimators. For example, one of the most widely used doubly-robust estimators is the Augmented-IPW (AIPW) estimator.

$$A\hat{T}E = \frac{1}{N} \sum_{i=1}^N \underbrace{h(X_i, T_i)}_{\text{Outcome model}} + \underbrace{\left[\frac{T_i}{\hat{\pi}(X_i, T_i)} + \frac{(1 - T_i)}{\hat{\pi}(X_i, T_i)} \right]}_{\text{Prop. score}} \cdot \underbrace{[Y_i - h(X_i, T_i)]}_{\text{Residual bias}} \quad (3)$$

The first term is the outcome model, while the third term accounts for any residual bias left over by the outcome model. The propensity score (second term) weights the relative importance of each unit's residual bias to estimation of $A\hat{T}E$. As expected, this estimator is unbiased if the IPW and regression estimators are consistently estimated. However, the model is attractive because it will be consistent if *either* the propensity score $\pi(X, T)$ is correctly

specified or the regression model h is consistently specified (Glynn and Quinn 2010). The model also provide efficiency gains with respect to the use of each model separately, and especially with respect to weighting alone. Many of the algorithms introduced below combine outcome regression with some type of treatment modeling using multi-task learning for double robustness.

4. Four Different Approaches to Deep Causal Estimation

The architectures proposed in the deep learning literature draw inspiration from existing approaches to estimation under selection on observables: outcome modeling via deep regression, balancing via representation learning, and IPW adjustment after representation learning. Nearly every algorithm discussed below contains some form of outcome modeling, and most contain some form of representation learning. In addition to these three strategies we describe an approach unique to deep learning: generative modeling and adversarial training. Generative models estimate the joint distribution of covariates, treatment, and outcome and/or modeling of the treatment effect or counterfactual distribution. This section also describes other uses of GAN-like adversarial training to enhance performance of other deep causal estimators. Throughout the review, algorithms are presented via their loss functions and architectures (see Box 1).

4.1. Deep Outcome Modeling

S-Learners and T-Learners

Because at most one potential outcome is unobserved, it is not possible to apply supervised models to directly learn treatment effects. Across econometrics, biostatistics, and machine learning, a common approach to this challenge has been to instead use machine learning to model each potential outcome separately and use plug-in estimators for treatment effects (Chernozhukov et al. 2016; Van der Laan and Rose 2011; Wager and Athey 2018) . As with linear models, a single neural model can be trained to learn both potential outcomes (“S[ingle]-learner”) (Fig. 1B), or two independent models can be trained to learn each potential outcome

Figure 4: A. **T-learner**. In a T-learner, separate feed-forward networks are used to model each outcome. We denote the function encoded by these outcome modelers h . B. **TARNet**. TARNet extends the T-learner with shared representation layers. The motivation behind TARNet (and further elaborations of this model) is that the multi-task objective of accurately predicting both the treated and control potential outcomes forces the representation layers to learn a balancing function Φ such that the $\Phi(X|T = 0)$ and $\Phi(X|T = 1)$ are overlapping distributions in representation space.

(a “T-learner”) (Johansson et al. 2020) (Fig. 4A). In both cases, the neural network estimators would be feed-forward networks tasked with minimizing the MSE in the prediction of observed outcomes. The joint loss function for a T-learner can be written as:

$$\mathcal{L}(Y, h(X, T)) = \text{MSE}(Y, h(X, 0)) + \text{MSE}(Y, h(X, 1)) \quad (4)$$

After training, inputting the same unit into both networks of a T-learner will produce predictions for both potential outcomes: $\hat{Y}(T)$ and $\hat{Y}(1 - T)$. We can plug-in these predictions

to estimate the *CATE* for each unit,

$$\hat{\tau}_i = (1 - 2T_i)(\hat{Y}_i(T_i) - \hat{Y}_i(1 - T + i))$$

and the average treatment effect as,

$$\hat{ATE} = \frac{1}{N} \sum_{i=1}^N \hat{\tau}_i$$

Nearly all of the models described below combine this plug-in outcome modeling approach with other forms of treatment adjustment.

4.2. Balancing through Representation Learning

Balancing is a treatment adjustment strategy that aims to deconfound the treatment from outcome by forcing the treated and control covariate distributions closer together (Johansson et al. 2016). *The novel contribution of deep learning to the selection on observables literature is the proposal that a neural network can transform the covariates into a representation space Φ such that the treated and control covariate distributions are indistinguishable (Fig. 2).* To encourage a neural network to learn balanced representations, the seminal paper in this literature (Shalit et al. (2017)) propose a simple two-headed neural network called Treatment Agnostic Regression Network (TARNet) that extends the outcome modeling T-learner with shared representation layers (Fig. 4B). Each head models a separate potential outcome. One head learns the function $\hat{Y}(1) = h(\Phi(X), 1)$, and the other head learns the function $\hat{Y}(0) = h(\Phi(X), 0)$. Both heads backpropagate their gradients to shared representation layers that learn $\Phi(X)$. The idea is that these representation layers must learn to balance the data because they are tasked with predicting both outcomes.

The complete objective for the network is to minimize the parameters of h and Φ for all n units in the training sample such that,

$$\arg \min_{h, \Phi} \frac{1}{N} \sum_{i=1}^N MSE(Y_i(T_i), \underbrace{h(\Phi(X_i), T_i)}_{\hat{Y}_i(T_i)}) + \lambda \underbrace{\mathcal{R}(h)}_{L_2} \quad (5)$$

where $\mathcal{R}(h)$ is a model complexity term (e.g., for L_2 regularization) and λ is a hyperparameter chosen through model selection.

Extending Representation Balancing with IPMs

Deep Dive: CFRNet (Shalit et al. (2017); Johansson et al. (2018, 2020))

Beyond receiving outcome modeling gradients for both potential outcomes, the authors have subsequently extended TARNet with additional losses that explicitly encourage balancing by minimizing a statistical distance between the two covariate distributions in representation space. These distances are called integral probability metrics (Müller 1997).² Johansson et al. (2016); Shalit et al. (2017); Johansson et al. (2018) propose two possible IPMs, the Wasserstein distance and the maximum mean discrepancy distance (MMD) for use in these architectures.

The Wasserstein or “Earth Mover’s” distance fits an interpretable “map” (i.e. a matrix) showing how to efficiently convert from one probability mass distribution to another. The Wasserstein distance is most easily understood as an optimal transport problem (i.e., a scenario where we want to transport one distribution to another at minimum cost). The nickname “Earth mover’s distance” comes from the metaphor of shoveling dirt to terraform one landscape into another. In the idealized case in which one distribution can be perfectly transformed into another, the Wasserstein map corresponds exactly to a perfect one-to-one matching on covariates strategy (Kallus 2016).

The MMD is the normed distance between the means of two distributions, after a kernel function ϕ has transformed them into a high-dimensional space called a reproducing kernel Hilbert Space (RKHS) (Gretton et al. 2012). The MMD with an L^2 norm in RKHS $\mathcal{H}$ can be specified as:

²Zhang et al. (2020) criticize the usage of IPMs because they make no restrictions on the moments of the transformed distributions. Thus while the covariate distributions may have a high percentage of overlap in representation space, this overlap may be substantially biased in unknown ways.

$$MMD(P, Q) = \|\mathbb{E}_{X \sim P} \phi(X) - \mathbb{E}_{X \sim Q} \phi(X)\|_{\mathcal{H}}^2 \quad (6)$$

The metric is built on the idea that there is no function that would have differing Expected Values for P and Q in this high-dimensional space if P and Q are the same distribution (Huszar 2015). The MMD is inexpensive to calculate using the “kernel trick” where the inner product between two points can be calculated in the RKHS without first transforming each point into the RKHS.³

When an IPM loss is applied to the representation layers in TARNet, the authors call the resulting network “CounterFactual Regression Network” (CFRNet) (Fig. 5) (Shalit et al. 2017). The loss function for this network is

$$\min_{h, \Phi, IPM} \frac{1}{N} \sum_{i=1}^N \underbrace{MSE(Y_i, h(\Phi(X_i), T_i))}_{\text{Outcome Loss}} + \lambda \underbrace{IPM(\Phi(X|T=1), \Phi(X|T=0))}_{\text{Dist. b/w T \& C covar. distributions}} + \alpha \underbrace{\mathcal{R}(h)}_{L_2} \quad (7)$$

where $R(h)$ is a model complexity term and λ and α are hyperparameters.

These two papers also make important theoretical contributions by providing bounds on the generalization error for the PEHE (Hill 2011). In Shalit et al. (2017), they show that the PEHE is bounded by the sum of the factual loss, counterfactual loss, and the variance of the conditional outcome. Adding an L_2 penalty penalizes large weights for units at the propensity extremes that might bias the results.

In Johansson et al. (2020), the authors introduce estimated IPW weights $\pi(\Phi(X), T)$ to CFRNet to provide consistency guarantees (Fig. 5B). Theoretically, they also use these weights to relax the overlap assumption as long as the weights themselves obey the positivity assumption. From a practical standpoint, adding weights that are optimized smoothly across the whole dataset each epoch reduces noise created by calculating the IPM score in small batches. Weighted CFRNet minimizes the following loss function:

³This kernel trick is also what makes support vector machines computationally tractable.

Figure 5: **A. CFRNet** architecture originally introduced in Shalit et al. (2017). CFRNet adds an additional integral probability metric (IPM) loss to TARNet to explicitly force representations of the treated and control covariates closer in representation space.

B. Weighted CFRNet adds a propensity score head to CFRNet to predict IPW-weighted outcomes. During training, the propensity score is used to reweight both the predicted outcomes $\hat{Y}(0)$ and $\hat{Y}(1)$, as well as the represented covariate distributions in calculation of the IPM loss. This allows the authors to provide consistency guarantees and relax the overlap assumption. Figures adapted from Johansson et al. (2020).

C. Dragonnet Dragonnet also adds a propensity score head to TARNet and a free “nudge” parameter ϵ . In an adaptation of Targeted Maximum Likelihood Estimation, $\hat{\pi}$ and ϵ are used to re-weight the outcomes to provide lower biased estimates of the ATE .

$$\begin{aligned}
& \arg \min_{h, \Phi, IPM, \pi, \lambda_h, \lambda_w} \frac{1}{N} \sum_{i=1}^N \underbrace{\frac{\hat{P}(T_i)}{\pi(\Phi(X_i), T_i)}}_{IPW} \cdot \underbrace{MSE(Y_i, h(\Phi(X_i), T_i))}_{\text{Outcome Loss}} + \lambda_h \underbrace{\mathcal{R}(h)}_{L_2 \text{ Outcome}} + \\
& \underbrace{\alpha \cdot IPM\left(\underbrace{\frac{\hat{P}(1)}{\pi(\Phi(X, 1))}}_{IPW} \cdot \Phi(X|T=1), \underbrace{\frac{\hat{P}(0)}{\pi(\Phi(X, 0))}}_{IPW} \cdot \Phi(X|T=0)\right)}_{\text{Distance between IPW weighted T \& C covar. distributions}} + \lambda_\pi \underbrace{\frac{\|\pi\|_2}{N}}_{L_2 \text{VAR}(\pi)} \quad (8)
\end{aligned}$$

where $R(h)$ is a model complexity term and λ_h , λ_π and α are hyperparameters. The final term is a regularization term on the variance of the weight parameters.

Extending Representation Balancing with Matching

Beyond IPMs, other approaches have directly embraced matching as a balancing strategy. Yao et al. (2018) train their TARNet on six point mini-batches of propensity score-matched units with additional reconstruction losses designed to preserve the relative distances between these points when projecting them into representation space. Schwab et al. (2018) takes an even simpler approach by feeding random batches of propensity-matched units to the TarNet outcome structure.

4.3. Extensions with Inverse Propensity Score Weighting

Rather than applying losses directly to the representation function, IPW methods estimate propensity scores from representations using the function $\pi(\Phi(X), T) = P(T|X)$. As in traditional IPW estimators, these methods exploit the sufficiency of correctly-specified propensity scores to reweight the plugged-in outcome predictions and provide unbiased estimates of the ATE (Rosenbaum and Rubin 1983). Because these models combine outcome modeling with IPW, they retain the attractive statistical properties of doubly robust estimators discussed in section 3.2.2. Atan et al. (2018) combines adversarial learning with IPW estimation, while Shi et al. (2019)’s Dragonnet model adapts semi-parametric estimation theory for batch-wise neural network training in a procedure they call “Targeted Regularization” (TarReg) (Kennedy 2016). We discuss Dragonnet and Targeted Regularization in more detail below, including a

brief introduction to semi-parametric theory and targeted maximum likelihood (TMLE) for context.

Deep Dive: Dragonnet (Shi et al. (2019))

Rather than adding an IPM loss, another trivial extension to TARNet is to add a third head to predict the propensity score. This third head could use multiple neural network layers or just a single neuron, as proposed in Dragonnet (Fig. 5) (Shi et al. 2019).

The loss function for this network looks like this:

$$\arg \min_{\Phi, \pi, h} \underbrace{MSE(Y, h(\Phi(X), T))}_{\text{Outcome Loss}} + \alpha \underbrace{\text{BCE}(T, \pi(\Phi(X), T))}_{\pi \text{ Loss}} + \lambda \underbrace{\mathcal{R}(h)}_{L_2} \quad (9)$$

with α being a hyperparameter to balance the two objectives.

Semi-Parametric Theory

The application of semi-parametric theory to causal inference is focused on estimating a target parameter of a distribution P of treatment effects $T(P) := ATE$. While we do not know the true distribution of treatment effects because we lack counterfactuals, we do know some parameters of this distribution (e.g., the treatment assignment mechanism). We can encode these constraints in the form of a likelihood that parametrically defines a set of possible approximate distributions of P from our existing data called $\mathcal{P}$. Within this set there is a sample-inferred distribution $\tilde{P} \in \mathcal{P}$, that can be used to estimate $T(P)$ using $T(\tilde{P})$.

Regardless of $\tilde{P}$ chosen, $\tilde{P} \neq P \rightarrow T(\tilde{P}) \neq T(P)$. We do not know how to pick $\tilde{P}$ with finite data to get the best estimate $T(\tilde{P})$. We can maximize a likelihood function to pick $\tilde{P}$, but there may be “nuisance” parameters in the likelihood that are not the target and we do not care about estimating accurately. Maximum likelihood optimization may provide lower-biased estimates of these nuisance terms at the cost of better estimates of $T(P)$.

To sharpen the likelihood’s focus on $T(P)$, we define a “nudge” parameter ϵ that moves $\tilde{P}$ closer to P (thus moving $T(\tilde{P})$ closer to $T(P)$). An influence curve of $T(P)$ tells us how changes in ϵ will induce changes in $T(P + \epsilon(\tilde{P} - P))$. We’ll use this influence curve to fit

ϵ to get a better approximation of $T(P)$ within the likelihood framework. In particular, there is a specific **efficient influence curve (EIC)** that provides us with the lowest variance estimates of $T(P)$. In causal estimation, solving the EIC for the ATE yields estimates that are asymptotically unbiased, efficient, and have confidence intervals with (asymptotically) correct coverage.

The EIC for the ATE is,

$$EIC_{ATE} = \frac{1}{N} \sum_{i=1}^N \underbrace{\left[\underbrace{\left(\frac{T}{\pi(x,1)} - \frac{1-T}{\pi(x,0)} \right)}_{\text{Treatment Modeling}} \times \underbrace{(Y - h(x,T))}_{\text{Residual Confounding}} \right]}_{\text{Adjustment}} + \underbrace{[h(x,1) - h(x,0)]}_{\text{Outcome Modeling}} - ATE \quad (10)$$

Minimizing EIC_{ATE} to 0,

$$ATE = \frac{1}{N} \sum_{i=1}^N \underbrace{\left[\underbrace{\left(\frac{T}{\pi(x,1)} - \frac{1-T}{\pi(x,0)} \right)}_{\text{Treatment Modeling}} \times \underbrace{(Y - h(x,T))}_{\text{Residual Confounding}} \right]}_{\text{Adjustment}} + \underbrace{[h(x,1) - h(x,0)]}_{\text{Outcome Modeling}} \quad (11)$$

The underbraces illustrate how EIC_{ATE} resembles a doubly robust estimator. When the EIC is minimized (set to 0) as in equation 11, the ATE is equal to the outcome modeling estimate plus a treatment modeling estimate proportional to the residual error.

From TMLE to Targeted Regularization

Targeted Regularization (TarReg) is closely modeled after “Targeted Maximum Likelihood Estimation” (TMLE) (Van der Laan and Rose 2011). TMLE is an iterative procedure where a nuisance parameter ϵ is used to nudge the outcome models towards sharper estimates of the ATE when minimizing the EIC as in Equation 11.

1. Fit h by predicting outcomes (e.g., using TARNet) and minimizing $MSE(Y, h(X, T))$
2. Fit π by predicting treatment (e.g., using logistic regression) and $BCE(T, \pi(X, T))$

3. Plug-in h and π functions to fit ϵ and estimate $h^*(X, T)$ where,

$$\underbrace{h^*(X, T)}_{Y^*} = \underbrace{h(X, T)}_{\hat{Y}} + \underbrace{\left(\frac{T}{\pi(X, T)} - \frac{1-T}{\pi(X, 1-T)} \right)}_{\text{Treatment Modeling Adjustment}} \times \underbrace{\epsilon}_{\text{"nudge"}}$$

by minimizing $MSE(Y, h^*(X, T))$. This is equivalent to minimizing the “Adjustment” part in equation 11.

4. Plug-in $h^*(X, T)$ to estimate $\hat{ATE}$:

$$\hat{ATE}_{TMLE} = \frac{1}{N} \sum_{i=1}^N \underbrace{h^*(X_i, 1)}_{Y_i^*(1)} - \underbrace{h^*(X_i, 0)}_{Y_i^*(0)}$$

Targeted Regularization takes TMLE and adapts it for a neural network loss function. The main difference is that steps 1 and 2 above are done concurrently by Dragonnet, and that the loss functions for the first three steps are combined into a single loss applied to the whole network at the end of each batch. It requires adding a single free parameter to the Dragonnet network for ϵ .

At a very intuitive level, Targeted Regularization is appealing because it introduces a loss function to TARNet that explicitly encourages the network to learn the mean of the treatment effect distribution, and not just the outcome distribution. The Targeted Regularization procedure proceeds as follows:

In each epoch:

1. (a) Use Dragonnet to predict $h(\Phi(X), T)$ and $\pi(\Phi(X), T)$.
- (b) Calculate the standard ML loss for the network using a hyperparameter α :

$$\arg \min_{\Phi, \pi, h} \underbrace{MSE(Y, h(\Phi(X), T))}_{\text{Outcome Loss}} + \alpha \underbrace{\text{BCE}(T, \pi(\Phi(X), T))}_{\pi \text{ Loss}} + \lambda \underbrace{\mathcal{R}(h)}_{L_2}$$

2. (a) Compute $h^*(\Phi(X), T)$ as above,

$$\underbrace{h^*(\Phi(X), T)}_{Y^*} = \underbrace{h(\Phi(X), T)}_{\hat{Y}} + \underbrace{\left(\frac{T}{\pi(\Phi(X), T)} - \frac{1-T}{\pi(\Phi(X), 1-T)} \right)}_{\text{Treatment Modeling Adjustment}} \times \underbrace{\epsilon}_{\text{"nudge"}}$$

- (b) Calculate the targeted regularization loss: $MSE(Y, h^*(\Phi(X), T))$

3. Combine and minimize the losses from 1 and 2 using a hyperparameter β ,

$$\arg \min_{\Phi, h, \epsilon} = \underbrace{MSE(Y, h(\Phi(X), T))}_{\text{Outcome Loss}} + \underbrace{\alpha \cdot BCE(T, \pi(\Phi(X, T)))}_{\pi \text{ Loss}} + \underbrace{\lambda \mathcal{R}(h)}_{L_2} + \underbrace{\beta \cdot MSE(Y, h^*(\Phi(X), T))}_{\text{Targeted Regularization Loss}}$$

Step 3 of Targeted Regularization is exactly equivalent to minimizing the EIC up to a constant β .

At the end of training, we can thus estimate the targeted regularization estimate of the ATE ATE_{TR} as in TMLE:

$$ATE_{TR} = \frac{1}{N} \sum_{i=1}^N \underbrace{h^*(\Phi(x_i), 1)}_{y_i^*(1)} - \underbrace{h^*(\Phi(x_i), 0)}_{y_i^*(0)}$$

Other approaches to estimating IPW weights using adversarial training are discussed in the next section (Ozery-flato et al. 2018; Kallus 2018). We note that a number of other losses for the basic TarNet/Dragonnet architecture have been proposed with differing theoretical motivations. In the interest of space, these approaches are discussed briefly in Box 5.

Box 5: Other Flavors of TarNet

A number of additional losses have been proposed for the representation layers in the two-headed TARNet or three-headed Weighted CFRNet/Dragonnet architectures:

- **Reconstruction Loss.** Some authors have proposed that reconstruction losses should be applied to representation layers to improve confidence in the invertability assumption (Du et al. 2019; Zhang et al. 2020). These losses simply minimize an L_2 norm between inputs and outputs to force the representation function to be able to reconstruct its inputs, along with its other tasks: $\mathcal{L}(X, X') = ||X - X'||^2$
- **Adversarial Loss.** Rather than learn to predict the propensity score Du et al. (2019), apply an adversarial gradient to force the representation layers to “unlearn” information about treatment assignment. This approach is also applied in Bica et al. (2020a).
- **Propensity Dropout.** While not a loss *per se*, Alaa et al. (2017) propose probabilistically applying dropout to neurons based on the Shannon Entropy in propensity score predictions. This penalty forces the network to attend comparatively more to data where overlap is greatest and the propensity score is not close to either the 0 or 1 extremes.

Note that because they solve the EIC estimating equation for the ATE, both TMLE and

Targeted Regularization are doubly robust estimators.

4.4. Adversarial Training of Generative Models, Representations, and IPW

The Origins of Adversarial Training in GANs

Adversarial training approaches include a wide variety of architectures where two networks or loss functions compete against each other. Adversarial approaches are inspired by Generative Adversarial Networks (GANs) (Box 6) (Goodfellow et al. 2014). In the machine learning literature on causal inference, adversarial training has been applied both to trade off outcome modeling and treatment modeling tasks during representation learning, as well as to trade off estimation and regularization of IPW weights. GANs have also been used directly as generative models for counterfactual and treatment effect distributions.

Box 6: Generative Adversarial Networks (GAN)

In GANs, two networks, a discriminator network D and a generator network G , play a zero-sum game like cops and robbers. The generator network's job is to learn a distribution from which the training data X could have credibly been generated. In each training batch, the generator produces a new outcome (originally images, but could be IPW weights, counterfactuals or treatment effects) by drawing a random noise sample from a known distribution Z (e.g. Gaussian) and transforming it into outcomes with the function $G(Z) = \hat{X}$. The discriminator's job is to learn a function $D(X) = P(X \text{ is real})$ that can distinguish whether the outcome is from the training data X , or whether it is a "fake" $\hat{X}$ created by the generator. The generator then receives a negative version of the discriminator's loss, a penalty that is proportional to how well it was able to "deceive" the discriminator. The discriminator's loss can be the log loss, Jensen-Shannon divergence (Goodfellow et al. 2014), the Wasserstein distance (Arjovsky et al. 2017; Gulrajani et al. 2017), or any number of divergences and IPMs. Formally, the generator attempts to minimize the following loss function,

$$\arg \min_G = \underbrace{\mathbb{E}_X [\mathcal{L}(D(X))]}_{\text{EV data } P(X \text{ is real})} + \underbrace{\mathbb{E}_Z [1 - \mathcal{L}(D(G(Z)))]}_{\text{EV fakes } P(\hat{X} \text{ is real})}$$

where the first quantity is the discriminator's estimated probability data from X is indeed real, and the second quantity is the discriminator's estimate that a generated quantity from the distribution Z is real.

Because the discriminator is trying to catch the generator, its objective is to maxi-

mize the same function,

$$\arg \max_D = \underbrace{\mathbb{E}_X [\mathcal{L}(D(X))]}_{\text{EV data } P(X \text{ is real})} + \underbrace{\mathbb{E}_Z [1 - \mathcal{L}(D(G(Z)))]}_{\text{EV fakes } P(\tilde{X} \text{ is real})}$$

In practice, the discriminator and the generator are trained either alternately or simultaneously, with the discriminator increasing its ability to discern between real and fake outcomes over time, and the generator increasing its ability to deceive the discriminator. The idea is that the adaptive loss function created by the discriminator can coax the generator out of local minima to generate superior outcomes. Results by these models have been impressive, and many of the fake portraits and “deepfake” videos circulating online in recent years are generated by this architectures. The advantage of GANs is that they can impressively learn very complex generative distributions with limited modeling assumptions. The disadvantage of GANs is that they are difficult and unreliable to train, often plateauing in local optima.

GANs as Generative Models of Treatment Effect Distributions (GANITE)

Deep Dive: GANITE (Yoon et al. (2018))

Although a generative model of the treatment effect distribution is generally unknown, a natural application of GANs is to try to machine learn this model from data. GANITE uses two GANs: GAN_1 to model the counterfactual distribution and GAN_2 to model the *CATE* distribution (Yoon et al. 2018) (Fig. 6). The training procedure for GAN_1 is as follows:

1. Taking X, T , and generative noise Z as input, generator G_1 generates both potential outcomes $\{\tilde{Y}(T), \tilde{Y}(1-T)\}$. A factual loss $MSE(Y(T), \tilde{Y}(T))$ is applied.
2. Create a new vector $\mathbf{C} = \{Y(T), \tilde{Y}(1-T)\}$ by combining the observed potential outcome and the counterfactual predicted by G_1 .
3. Taking X and $\mathbf{C}$ as inputs, the discriminator rates each value in $\mathbf{C}$ for the probability that it is the observed outcome using the categorical cross entropy:

$$\mathcal{L}(D) = CCE(\underbrace{\{P(\mathbf{C}_0 = Y(T)), P(\mathbf{C}_1 = Y(T))\}}_{\text{Prob first idx is real } \quad \text{Prob sec idx is real}}, \underbrace{\{\mathbf{C}_0 == Y(T), \mathbf{C}_1 == Y(T)\}}_{\text{1 if idx 0 is real } \quad \text{1 if idx 1 is real}}) \quad (12)$$

4. This loss is then fed back to G_1 such that the total loss for the generator is now

$$\arg \min_{G_1} = MSE(Y(T), \tilde{Y}(T)) - \lambda \mathcal{L}(D_1) \quad (13)$$

After generator G_1 is trained to completion, the authors use $\mathbf{C}$ as a “complete dataset” containing both a factual outcome and a counterfactual outcome to train GAN_2 , which generates

treatment effects:

1. Taking only X and generative noise Z as input, G_2 generates a new potential outcome vector $\mathbf{R} = \{\hat{Y}(T), \hat{Y}(1 - T)\}$. G_2 receives an MSE loss to minimize the difference between it's predictions and the "complete dataset" $\mathbf{C}$: $MSE(\mathbf{C}, \mathbf{R})$.
2. Discriminator D_2 takes X , $\mathbf{C}$, and $\mathbf{R}$ as inputs and estimates a probability that $\mathbf{C}$ is the "complete" dataset, and that $\mathbf{R}$ is the "complete dataset":

$$\mathcal{L}(D) = CCE(\underbrace{\{P(\mathbf{C} = \mathbf{C})\}}_{\text{Prob } \mathbf{C} \text{ is "CD"}}, \underbrace{\{P(\mathbf{R} = \mathbf{C})\}}_{\text{Prob } \mathbf{R} \text{ is "CD"}}, \underbrace{\{\mathbf{C} == \mathbf{C}\}}_{1 \text{ if idx 0 is } \mathbf{C}}, \underbrace{\{\mathbf{C}_1 == Y(T)\}}_{1 \text{ if idx 1 is } \mathbf{C}}) \quad (14)$$

3. This loss is then fed back to the generator G_2 such that the total loss for the generator is now

$$\arg \min_{G_2} = MSE(\mathbf{C}, \mathbf{R}) - \lambda \mathcal{L}(D_2) \quad (15)$$

At the end of training, G_2 should be able to predict treatment effects with only covariates X and noise Z as inputs. An evolution of GANITE, SCIGAN, extends this framework to settings with more than one treatment and continuous dosages (Bica et al. 2020b).

Adversarial Representation Balancing

The use of the IPM loss in CFRNet (Shalit et al. 2017) may also be viewed as an adversarial approach in that the representation layers are forced to maximize performance on two competing tasks: predicting outcomes and minimizing an IPM. Rather than using an IPM loss, other authors have trained propensity score estimators that send positive (rather than negative) gradients back to the representation layers (Atan et al. 2018; Du et al. 2019). This strategy explicitly decorrelates the covariate representations from the treatment (see Box 5).

Bica et al. (2020a) extend this approach to settings with treatment over time using a recurrent neural network. In their medical setting, decorrelating treatment from patient covariates and history allows them to estimate treatment effects at each individual snapshot. This algorithm is described in Section 5.

Adversarial IPW Learning

In a simple adversarial IPW model, Ozery-flato et al. (2018) estimate IPW weights for the ATE adversarially. A discriminator is presented with two sets of weights: one uniform and

Figure 6: **GANITE** has two GANs. The first generator G_1 generates counterfactuals $\tilde{Y}(T)$. The discriminator D_1 attempts to discriminate between these predictions and real data ($Y(T)$). The second generator proposes pairs of potential outcomes $\hat{Y}(0)$ and $\hat{Y}(1)$ (i.e., treatment effects), a vector we call R . Discriminator D_2 attempts to discern between R and a “complete dataset” C created by pairing each observed/factual outcome $Y(T)$ with a synthetic outcome $\tilde{Y}(1-T)$ proposed by G_1 . Although we do not show gradients in other figures, we make an exception for GANs (red line).

the other estimated, and tasked with distinguishing between the two distributions. The “generator” updates the estimated weights to minimize the ability of the discriminator to distinguish between them. The generator uses exponential gradient ascent during training to regularize the weights and minimize their variance. [Kallus \(2018\)](#) similarly proposes a GAN set-up where a discriminator network attempts to minimize a “discriminative distance”, while the “generator” of the weights is itself a deep neural network.

5. Extending Causal Estimation to Complex Data

Below we describe extensions of causal inference that are not possible with other types of machine learning. First we address estimation of treatment effects with time varying treatments and confounding. We also see great potential for neural networks in causal inference when latent confounders, treatments, and mediators are encoded in complex data (e.g., text, networks, and images) as well.

5.1. Conditioning on Time-Varying Confounding

One natural extension of deep causal estimation is measuring treatment effects at discrete time points when treatment is administered over time. While models developed by Robins et al. for estimating effects with time-varying treatments and confounding have existed for decades, they require strong parametric assumptions for estimation ([Robins 1994](#); [Robins et al. 2000, 2009](#)). Within the machine learning community, researchers have begun to extend the above-described representation learning approaches to temporal settings using recurrent neural networks (RNN) (Box 7) ([Lim et al. 2018](#); [Bica et al. 2020a](#); [Cheng et al. 2021](#)). For example, [Cheng et al. \(2021\)](#) propose to replace each outcome modeling head of CFRNet with an RNN, one modeling the treated outcomes over time and one modeling the control outcomes. [Bica et al. \(2020a\)](#) more explicitly deal with time-varying confounding by having each unit of the RNN take time-varying X , non-varying features V , T , and representations from the temporally-previous unit as inputs. Each unit predicts an outcome at the time step, but also receives an adversarial gradient to “unlearn” information about the treatment at this and previous timesteps.

Box 7: Recurrent Neural Networks (RNN)

Recurrent neural networks are a specialized architecture created for learning outcomes from sequential data (e.g. time series, biological sequences, text). In a classic RNN, each “unit” u in the network takes as input its own covariates X (or possibly a representation) and a representation produced by the previous unit, encoding cumulative information about earlier states in the sequence. These units are not just simple hidden layers: there is a set of weights within each unit for its raw inputs, the representation from the previous time step, and its outputs. Different RNN variants have different operations for integrating past representations with present inputs. Recurrent neural networks may be directed acyclic graphs or feedback on themselves. Commonly used variants include Gated Recurrent Unit networks (GRU) and Long-term Short-term memory networks (LSTM) (Cho et al. 2014; Hochreiter and Schmidhuber 1997).

5.2. Relaxing Strong Ignorability: Controlling for Latent Confounders in Text, Graphs, and Images

Modeling latent confounders explicitly from observed data is an active area of research within both the machine learning and statistics communities (Louizos et al. 2017; Mayer et al. 2020; Jesson et al. 2020; Witty et al. 2020; Wang and Blei 2019). However, the identification assumptions required by these approaches are still somewhat controversial Rissanen and Marttinen (2021); D’Amour (2019). In this section, we focus on another exciting approach: extracting implicit information about unmeasured confounders from networks, texts, and images.

Conditioning on Latent Confounding Encoded in Networks

Causal inference from networked data is difficult for two reasons. First, the causal effect of treating nodes on a network’s structure is difficult to measure because the lack of treatment contagion/homophily is a fundamental assumption in all causal frameworks (SUTVA) (Shalizi and Thomas 2011). Second, we generally do not have compelling generative models for networks that describe how features/covariates might lead to network topologies. Instead of generative network models, early literature in this area has thus largely leveraged represen-

tations of nodes that encode information about both the node’s covariates and the node’s structural position. The type of neural network that produces these representations is called a graph neural network (GNN) (Box 8).

Box 8: Graph Neural Networks (GNN)

Graph neural networks (GNNs) are the current state-of-the-art approach for creating representations for nodes in graphs. Compared to previous approaches that relied on “shallow” embeddings based only on a node’s local context (e.g., random walks to nearby nodes), GNNs are attractive because their node representations are aggregated from the structural position and covariates of all nodes n degrees away from the target node, where n is the number of graph neural network layers.

The most intuitive understanding of how graph neural networks work is as a message passing system (Gilmer et al. 2017). We use one of the first GNN papers, the Graph Convolutional Network as an example (Kipf and Welling 2016). In this interpretation, each node has a message that it passes to its neighbors through a graph convolution operation. In the first layer of a GNN this message would consist of the node’s covariates/features. In consecutive layers of the network, these messages are actually representations of the node produced by the previous layer. During graph convolution, each node multiplies incoming messages by its own set of weights and combines these weighted inputs using an aggregation function (e.g., summation). By the n -th GNN layer, these messages will contain structure and covariate information from all nodes n degrees away. For interested readers, there is also a spectral interpretation of this process. Typically GNNs are trained to produce representations of graphs by predicting the probability that two nodes are linked in the network, and then used for something else. One variant of the GNN uses an “attention” mechanism to vary the extent that nodes value messages from different neighbors (the graph attention network or GAT) (Veličković et al. 2017). This network is closely related to the transformer architecture discussed below in Box 9.

Ma and Tresp (2021) explore violations of the SUTVA assumption (i.e., treatment “spillover” or “interference”) by creating node representations from a GNN that aggregate treatment information of local neighbors and the structural position of nodes separately. These node representations are then fed to CFRNet (TARNet + an IPM loss) as a secondary input along with the central unit’s covariates X . Theoretically, they show that the error of this estimator is bounded by the maximum degree of the network.

There is also considerable interest in exploring whether the structural position of units in a network encodes information about unobserved/unmeasured confounders, in scenarios where SUTVA holds (i.e., relaxing the strong ignorability assumption). For example, even if age or

gender is unmeasured, these features might be inferable from a person’s homophilic friendship ties.

A first pass approach to this problem is proposed in [Guo et al. \(2019\)](#). They train two independent feed-forward networks for $h(\Phi(X))$ and $\pi(\Phi(X), T)$. Instead of using standard feed-forward layers as the early layers for these representation functions (as described in the main body) they prepend GNN architectures that take the graph structure and node covariates as input. They then combine these independent outcome and propensity estimates using a doubly robust estimator. A shared loss between the networks force Φ_h and Φ_π to produce similar representations of the nodes/covariates. [Chu et al. \(2021\)](#) and [Guo et al. \(2020\)](#) improve on this approach by dealing with the technical challenges of learning representations that are both faithful to the graph and can estimate causal effects.

[Veitch et al. \(2019b\)](#) take a formal semiparametric approach to identify a causal effect on graphs. They assume that a node representation need only encode necessary structural information about unobserved confounders X that allows one to make consistent estimates of the treatment and outcome, even if it does not encode all information about X completely. Under these conditions, along with an assumption that representations of the nodes are only weakly correlated (a strong assumption), they show that a variant of the AIPW estimator is asymptotically normal, unbiased, and a $\frac{1}{\sqrt{n}}$ consistent ATE estimator. The authors conduct a naturalistic simulation using 79,000 individuals in an online social network in Slovakia called Pokec.

Conditioning on Latent Confounding in Text Data

Compared to research on latent confounding encoded in networks, causal inference from text data is a much larger research area within both the computer science and social science communities. There are methods that treat text as treatments, mediators, proxies for latent confounding, or outcomes (see [Keith et al. \(2020\)](#) and [Feder et al. \(2021\)](#) for exhaustive reviews). In this article we narrowly focus on articles that use deep representations such as transformers (Box 9), but further note that methods relying on other quantitative representations (eg., topics, word counts) may address similar problems.

Box 9: Transformers

As of 2021, transformers are the hegemonic architecture used in natural language processing. After their introduction in 2017, these models improved performance on many high-profile NLP tasks across the board. Several enterprise-scale transformers have been featured in the media for their impressive performance in text generation and question answering (e.g. OpenAI’s GPT-3). Smaller models in broad use are based on the BERT architecture (Devlin et al. 2018).

The connection between GNNs, and specifically GATs, is the focus on attention mechanisms (Box 8). From this perspective, words in sentences are akin to nodes in networks, with their relative positions to each other being analogous to their structural positions in the graph. Transformers improved on previous sequential approaches to text analysis (i.e. RNNs) by having each word (or representation of a word) receive messages from not just adjacent words, but all words heterogeneously. Attention mechanisms throughout the architecture allow each layer of a transformer to attend to words or aggregated representation mechanisms heterogeneously. Architectures such as BERT stack transformer layers to create models with hundreds of millions of parameters. These models are expensive to train, both computationally and with respect to data, so they are often pretrained on large datasets and then “fine-tuned” (lightly re-trained) with smaller datasets for specific tasks.

Veitch et al. (2019a) proposes some conditions under which a causal effect would be identifiable using text representations under a semi-parametric framework. The motivation for the paper is that causal inference directly from text (or even topic models) poses a curse of dimensionality problem, and thus it would be convenient to use pre-trained representations from transformer algorithms such as BERT (Devlin et al. 2018) (see Box 9).

The main contribution of the paper is a theorem that states that if adjusting for confounding encoded in the text W is sufficient to identify a causal effect, then adjusting for confounding Z encoded in a representation $\Phi(W)$ is also sufficient for causal identification. They demonstrate this estimator on a Science of Science question testing the causal effect of equations on getting papers accepted to computer science conferences.

Pryzant et al. (2020) explores the scenario where both treatment (some linguistic property) and confounders are encoded in text. They note that an undeveloped idea in observational text identification is that the reader, not the writer, must be able to identify the treatment and produce the outcome. Since the reader’s perception of the treatment is also unobservable, we must use a proxy label (e.g. number of stars in a product review) to measure the treatment.

They show both that the ATE is identifiable in this scenario, and that any bias induced by differences between the proxy label and the actual perception of treatment may attenuate the ATE but not change its sign. In their proposed architecture, a transformer representation is fed to TARNet along with the non-text covariates X . The TARNet loss is simultaneously used to train itself and as a fine-tuning objective for the transformer.

Causal Inference on Images

(IN PROGRESS)

6. Conclusion

In this review we organize the emerging machine learning literature on deep learning for potential outcomes under four estimation strategies: outcome modeling, balancing, propensity score modeling, and adversarial training, with representative examples for each. The paper provides basic introductions to deep learning and some of the most common architectures used in deep learning today (e.g. recurrent neural networks, generative adversarial networks, graph neural networks, and transformers). Finally, we highlight promising work on controlling for time-varying confounder and confounders implicitly encoded in text, network, and image representations.

Deep causal estimators have a lot of advantages. They are low bias, excel at the estimation of heterogeneous treatment effects and provide opportunities for social science to predict causal effects in untreated populations beyond the original sample. However, there is still little consensus on important practical considerations needed to deploy these tools in the wild. First, the non-convexity of deep learning loss functions means that a model can end up in different optima even when it is trained on the same data with the same hyperparameters. It is not obvious how to perform traditional model selection (i.e., cross-validation) when one potential outcome is unobserved, but there is limited research on this topic ([Johansson et al. 2020](#); [Alaa and Van Der Schaar 2019](#)). Developing confidence intervals with either asymptotic guarantees or generalization bounds is another important direction for both rigor and ethical considerations. Recent work has proposed to use MC Dropout to identify CATE predictions

with low overlap (Jesson et al. 2020; Gal and Ghahramani 2016), but there is as-to-date sparse work on other possible approaches (e.g., bootstrapping, negative sampling, crossfitting).

To address some of these gaps between basic research and implementation, the included tutorials provide examples for social scientists to implement two consistent estimators from this literature, Dragonnet and Weighted CFRNet, as well as generic TARNet architectures (Shi et al. 2019; Johansson et al. 2020). The tutorials assume no previous experience with Tensorflow, and address practical considerations like model selection and uncertainty in training and regularizing deep causal estimators. The tutorials are available at github.com/kochbj/Deep-Learning-for-Causal-Inference.

References

- Alaa, A. and Van Der Schaar, M. (2019). Validating causal inference models via influence functions. In *International Conference on Machine Learning*, pages 191–201. PMLR.
- Alaa, A. M., Weisz, M., and Schaar, M. V. D. (2017). Deep Counterfactual Networks with Propensity-Dropout. In *International Conference on Machine Learning*.
- Arjovsky, M., Chintala, S., and Bottou, L. (2017). Wasserstein gan. *arXiv preprint arXiv:1701.07875*.
- Atan, O., Jordon, J., and Schaar, M. V. D. (2018). Deep-Treat : Learning Optimal Personalized Treatments from Observational Data Using Neural Networks. In *Association for the Advancement of Artificial Intelligence*.
- Austin, P. C. (2011). An introduction to propensity score methods for reducing the effects of confounding in observational studies. *Multivariate behavioral research*, 46(3):399–424.
- Bengio, Y. (2013). Deep learning of representations: Looking forward. In *International Conference on Statistical Language and Speech Processing*, pages 1–37. Springer.
- Bica, I., Alaa, A. M., Jordon, J., and van der Schaar, M. (2020a). Estimating counterfac-

- tual treatment outcomes over time through adversarially balanced representations. *arXiv preprint arXiv:2002.04083*.
- Bica, I., Jordon, J., and van der Schaar, M. (2020b). Estimating the effects of continuous-valued interventions using generative adversarial networks. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., and Lin, H., editors, *Advances in Neural Information Processing Systems*, volume 33, pages 16434–16445. Curran Associates, Inc.
- Brand, J. E., Koch, B., and Xu, J. (2020). Sage research methods.
- Cheng, L., Guo, R., and Liu, H. (2021). Long-term effect estimation with surrogate representation. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, WSDM '21, page 274–282, New York, NY, USA. Association for Computing Machinery.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., and Newey, W. K. (2016). Double machine learning for treatment and causal parameters.
- Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. (2014). Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.
- Chu, Z., Rathbun, S. L., and Li, S. (2021). Graph infomax adversarial learning for treatment effect estimation with networked observational data. *arXiv preprint arXiv:2106.02881*.
- Cuturi, M. (2013). Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in neural information processing systems*, pages 2292–2300.
- Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 5:455.
- Daza, D. (2019). Approximating wasserstein distances with pytorch. <https://dfdazac.github.io/sinkhorn.html>. Last accessed 2019-08-01.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

- Du, X., Duivesteijn, W., Sun, L., Nikolaev, A., and Pechizkiy, M. (2019). Adversarial Balancing-based Representation Learning for Causal Effect Inference with Observational Data.
- D’Amour, A. (2019). Comment: Reflections on the deconfounder. *Journal of the American Statistical Association*, 114(528):1597–1601.
- Feder, A., Keith, K. A., Manzoor, E., Pryzant, R., Sridhar, D., Wood-Doughty, Z., Eisenstein, J., Grimmer, J., Reichart, R., Roberts, M. E., et al. (2021). Causal inference in natural language processing: Estimation, prediction, interpretation and beyond. *arXiv preprint arXiv:2109.00725*.
- Gal, Y. and Ghahramani, Z. (2016). Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR.
- Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. (2017). Neural message passing for quantum chemistry. In *International conference on machine learning*, pages 1263–1272. PMLR.
- Glynn, A. N. and Quinn, K. M. (2010). An introduction to the augmented inverse propensity weighted estimator. *Political analysis*, 18(1):36–56.
- Goldszmidt, M. and Pearl, J. (1996). Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artificial Intelligence*, 84(1):57–112.
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press. <http://www.deeplearningbook.org>.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. (2012). A kernel two-sample test. *Journal of Machine Learning Research*, 13(Mar):723–773.

- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. (2017). Improved training of wasserstein gans.
- Guo, R., Li, J., Li, Y., Candan, K. S., Raglin, A., and Liu, H. (2020). Ignite: A mini-max game toward learning individual treatment effects from networked observational data. In Bessiere, C., editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 4534–4540. International Joint Conferences on Artificial Intelligence Organization. Special Track on AI in FinTech.
- Guo, R., Li, J., and Liu, H. (2019). Counterfactual evaluation of treatment assignment functions with networked observational data. *arXiv preprint arXiv:1912.10536*.
- Hernán, M. and Robins, J. (2020). *Causal Inference: What If. 2020*. Chapman & Hall.
- Hill, J. L. (2011). Bayesian Nonparametric Modeling for Causal Inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240.
- Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. 9(8):1735–1780.
- Holland, P. W. (1986). Statistics and causal inference. *Journal of the American statistical Association*, 81(396):945–960.
- Huszar, F. (2015). Another favourite machine learning paper: Adversarial networks vs kernel scoring rules. <https://www.inference.vc/another-favourite-machine-learning-paper-adversarial-networks-vs-kernel-scoring-rules/>. Last accessed 2019-08-01.
- Imbens, G. W. and Rubin, D. B. (2015). *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press.
- Ioffe, S. and Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. *ArXiv*, abs/1502.03167.
- Jesson, A., Mindermann, S., Shalit, U., and Gal, Y. (2020). Identifying causal effect inference failure with uncertainty-aware models. *CoRR*, abs/2007.00163.

- Johansson, F. and Shen, M. (2018). Causal inference & deep learning. MIT IAP.
- Johansson, F. D., Kallus, N., Shalit, U., and Sontag, D. (2018). Learning Weighted Representations for Generalization Across Designs.
- Johansson, F. D., Shalit, U., Kallus, N., and Sontag, D. A. (2020). Generalization bounds and representation learning for estimation of potential outcomes and causal effects. *CoRR*, abs/2001.07426.
- Johansson, F. D., Shalit, U., and Sontag, D. (2016). Learning Representations for Counterfactual Inference. In *International Conference on Machine Learning*, volume 48.
- Joo, J., Steen, F. F., and Zhu, S.-C. (2015). Automated facial trait judgment and election outcome prediction: Social dimensions of face. In *Proceedings of the IEEE international conference on computer vision*, pages 3712–3720.
- Kallus, N. (2016). Generalized optimal matching methods for causal inference. *arXiv preprint arXiv:1612.08321*.
- Kallus, N. (2018). DeepMatch : Balancing Deep Covariate Representations for Causal Inference Using Adversarial Training.
- Keith, K., Jensen, D., and O’Connor, B. (2020). Text and causal inference: A review of using text to remove confounding from causal estimates. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5332–5344, Online. Association for Computational Linguistics.
- Kennedy, E. H. (2016). Semiparametric theory and empirical processes in causal inference. In *Statistical causal inferences and their applications in public health research*, pages 141–167. Springer.
- Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.
- Kipf, T. N. and Welling, M. (2016). Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.

- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *nature*, 521(7553):436.
- Lim, B., Alaa, A. M., and van der Schaar, M. (2018). Forecasting treatment responses over time using recurrent marginal structural networks. *NeurIPS*, 18:7483–7493.
- Louizos, C., Shalit, U., Mooij, J. M., Sontag, D., Zemel, R., and Welling, M. (2017). Causal effect inference with deep latent-variable models. In *Advances in Neural Information Processing Systems*, pages 6446–6456.
- Ma, Y. and Tresp, V. (2021). Causal inference under networked interference and intervention policy enhancement. In *International Conference on Artificial Intelligence and Statistics*, pages 3700–3708. PMLR.
- Mayer, I., Josse, J., Raimundo, F., and Vert, J.-P. (2020). Missdeepcausal: Causal inference from incomplete data using deep latent variable models. *arXiv preprint arXiv:2002.10837*.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G., and Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2*, NIPS’13, page 3111–3119, Red Hook, NY, USA. Curran Associates Inc.
- Müller, A. (1997). Integral probability metrics and their generating classes of functions. *Advances in Applied Probability*, 29(2):429–443.
- Nagpal, C., Wei, D., Vinzamuri, B., Shekhar, M., Berger, S. E., Das, S., and Varshney, K. R. (2020). Interpretable subgroup discovery in treatment effect estimation with application to opioid prescribing guidelines. CHIL ’20, page 19–29, New York, NY, USA. Association for Computing Machinery.
- Ozery-flato, M., Thodoroff, P., and El-Hay, T. (2018). Adversarial balancing for causal inference.
- Pearl, J. (2009). *Causality*. Cambridge university press.
- Pryzant, R., Card, D., Jurafsky, D., Veitch, V., and Sridhar, D. (2020). Causal effects of linguistic properties. *arXiv preprint arXiv:2010.12919*.

- Rissanen, S. and Marttinen, P. (2021). A critical look at the consistency of causal estimation with deep latent variable models.
- Robins, J. (1986). A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Mathematical modelling*, 7(9-12):1393–1512.
- Robins, J. (1987). A graphical approach to the identification and estimation of causal parameters in mortality studies with sustained exposure periods. *Journal of chronic diseases*, 40:139S–161S.
- Robins, J., Hernán, M., Fitzmaurice, G., Davidian, M., Verbeke, G., and Molenberghs, G. (2009). Longitudinal data analysis. *Handbooks of Modern Statistical Methods*, pages 553–599.
- Robins, J. M. (1994). Correcting for non-compliance in randomized trials using structural nested mean models. *Communications in Statistics-Theory and methods*, 23(8):2379–2412.
- Robins, J. M., Hernan, M. A., and Brumback, B. (2000). Marginal structural models and causal inference in epidemiology.
- Rosenbaum, P. R. and Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55.
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, 66(5):688.
- Schwab, P., Linhardt, L., and Karlen, W. (2018). Perfect match : A simple method for learning representations for counterfactual inference with neural networks.
- Shalit, U., Johansson, F. D., and Sontag, D. (2017). Estimating individual treatment effect : generalization bounds and algorithms. In *International Conference on Machine Learning*.
- Shalizi, C. R. and Thomas, A. C. (2011). Homophily and contagion are generically confounded in observational social network studies. *Sociological Methods & Research*, 40(2):211–239. PMID: 22523436.

- Shi, C., Blei, D. M., and Veitch, V. (2019). Adapting Neural Networks for the Estimation of Treatment Effects.
- Srivastava, N., Hinton, G. E., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15:1929–1958.
- Stock, M. (2017). Notes on optimal transport. <https://michielsestock.github.io/OptimalTransport/>. Last accessed 2019-08-01.
- Stuart, E. A. (2010). Matching methods for causal inference: A review and a look forward. *Statistical science: a review journal of the Institute of Mathematical Statistics*, 25(1):1.
- Todorov, A., Mandisodza, A. N., Goren, A., and Hall, C. C. (2005). Inferences of competence from faces predict election outcomes. *Science*, 308(5728):1623–1626.
- Van der Laan, M. J. and Rose, S. (2011). *Targeted learning: causal inference for observational and experimental data*. Springer Science & Business Media.
- Veitch, V., Sridhar, D., and Blei, D. M. (2019a). Using text embeddings for causal inference.
- Veitch, V., Wang, Y., and Blei, D. M. (2019b). Using embeddings to correct for unobserved confounding in networks.
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., and Bengio, Y. (2017). Graph attention networks. *arXiv preprint arXiv:1710.10903*.
- Wager, S. and Athey, S. (2018). Estimation and Inference of Heterogeneous Treatment Effects using Random Forests. *Journal of the American Statistical Association*, 113(523):1228–1242.
- Wang, Y. and Blei, D. M. (2019). The blessings of multiple causes. *Journal of the American Statistical Association*, 114(528):1574–1596.
- Witty, S., Takatsu, K., Jensen, D., and Mansinghka, V. (2020). Causal inference using Gaussian processes with structured latent confounders. In III, H. D. and Singh, A., editors,

Proceedings of the 37th International Conference on Machine Learning, volume 119 of *Proceedings of Machine Learning Research*, pages 10313–10323. PMLR.

Yao, L., Li, S., Li, Y., Huai, M., Gao, J., and Zhang, A. (2018). Representation Learning for Treatment Effect Estimation from Observational Data. In *Advances in neural information processing systems*.

Yoon, J., Jordon, J., and Van Der Schaar, M. (2018). Ganite: Estimation of individualized treatment effects using generative adversarial nets. In *International Conference on Learning Representations*.

Zhang, Y., Bellot, A., and Schaar, M. (2020). Learning overlapping representations for the estimation of individualized treatment effects. In *International Conference on Artificial Intelligence and Statistics*, pages 1005–1014. PMLR.

A. Integral Probability Metrics

A.1. Wasserstein Distance

Following (Stock 2017; Daza 2019), suppose we have two discrete distributions (treated and control) with marginal densities $p(x)$ and $q(x)$ captured as vectors t and c , with dimensions n and m respectively. To compute the Wasserstein distance, we must define a "mapping matrix" P that defines the mapping of "earth" in $p(x)$ to corresponding piles in $q(x)$. Let $\mathbf{U}(t, c)$ be the set of positive, $n \times m$ mapping matrices where the sum of the rows is t and the sum of the columns is c .

$$\mathbf{U}(t, c) = \{P \in \mathbb{R}_{>0}^{n \times m} | P \cdot \mathbf{1}_m = t, P^T \cdot \mathbf{1}_n = c\} \quad (16)$$

In words, this matrix maps the probability mass from points in the support of $p(x)$ (i.e, the elements of t) to points in the support of $q(x)$ (the elements of c) (note that the mapping need not be one-to-one). We also have a "cost" matrix C that describes the cost of applying P (i.e. the cost of shoveling dirt according to the map described in P). The cost matrix can be computed using a norm ℓ (most commonly ℓ^2) between the points in t being mapped to c

in the mapping matrix P . Finally, the ℓ -norm Wasserstein distance dW_ℓ can be defined as

$$dW_\ell = \min_{P \in \mathcal{C}(t,c)} \sum_{i,j} P_{ij} C_{ij} \quad (17)$$

In other words, the Wasserstein distance is the smallest Frobenius inner product of a mapping matrix P that fits the above constraints, and its associated cost matrix C . Although this problem can be solved via linear programming, the Wasserstein distance is often implemented in a different form that works with continuous distributions and can be optimized by gradient descent (Arjovsky et al. 2017; Gulrajani et al. 2017). There is also a variant of the Wasserstein distance that imposes an entropy-based regularization on the coupling matrix to make it smoother or sparser called the Sinkhorn distance (Cuturi 2013).

B. Semi-Parametric Theory

The application of semi-parametric theory to causal inference is focused on estimating a target parameter of a distribution P of treatment effects $T(P) := ATE$. While we do not know the true distribution of treatment effects because we lack counterfactuals, we do know some parameters of this distribution (e.g., the treatment assignment mechanism). We can encode these constraints in the form of a likelihood that parametrically defines a set of possible approximate distributions of P from our existing data called $\mathcal{P}$. Within this set there is a sample-inferred distribution $\tilde{P} \in \mathcal{P}$, that can be used to estimate $T(P)$ using $T(\tilde{P})$.

Regardless of $\tilde{P}$ chosen, $\tilde{P} \neq P \Rightarrow T(\tilde{P}) \neq T(P)$. We do not know how to pick $\tilde{P}$ with finite data to get the best estimate $T(\tilde{P})$. We can maximize the likelihood function to pick $\tilde{P}$, but there are a lot of "nuisance" parameters in the likelihood that are not the target and we do not care about estimating accurately, so this will not necessarily give us the best estimate of $T(P)$.

To sharpen the likelihood's focus on $T(P)$, we define a "nudge" parameter ϵ that moves $\tilde{P}$ closer to P (thus moving $T(\tilde{P})$ closer to $T(P)$). An influence curve of $T(P)$ tells us how changes in ϵ will induce changes in $T(P + \epsilon(\tilde{P} - P))$. We'll use this influence curve to fit ϵ to get a better approximation of $T(P)$ within the likelihood framework. In particular, there is a specific **efficient influence curve (EIC)** that provides us with the lowest variance estimates

of $T(P)$.

Affiliation:

Bernard Koch

UCLA Department of Sociology

E-mail: bernardkoch@ucla.edu

URL: <https://soc.ucla.edu/grads/bernard-koch>

SocArXiv Website

<https://socopen.org/>

SocArXiv Preprints

<https://osf.io/preprints/socarxiv>

Preprint

Submitted: October 9, 2021

Accepted: October 9, 2021