

Neural Networks with Causal Graph Constraints: A New Approach for Treatment Effects Estimation

Roger Pros^{1,2}, Jordi Vitrià¹

¹Departament de Matemàtica i Informàtica, Universitat de Barcelona

²Zenital

{roger.pros, jordi.vitria}@ub.edu,

Abstract

In recent years, there has been a growing interest in using machine learning techniques for the estimation of treatment effects. Most of the best-performing methods rely on representation learning strategies that encourage shared behavior among potential outcomes to increase the precision of treatment effect estimates. In this paper we discuss and classify these models in terms of their algorithmic inductive biases and present a new model, NN-CGC, that considers additional information from the causal graph. NN-CGC tackles bias resulting from spurious variable interactions by implementing novel constraints on models, and it can be integrated with other representation learning methods. We test the effectiveness of our method using three different base models on common benchmarks. Our results indicate that our model constraints lead to significant improvements, achieving new state-of-the-art results in treatment effects estimation. We also show that our method is robust to imperfect causal graphs and that using partial causal information is preferable to ignoring it.

1 Introduction

Causal inference is a method used to understand whether and how a specific action (intervention or treatment) causes a change in a particular result or outcome. This process is crucial in many fields where understanding the cause-effect relationship is important; such as in medical trials, economic policy analysis, educational program evaluations, social science research, and artificial intelligence [Schölkopf *et al.*, 2021].

A key challenge in causal inference is that when we observe an outcome, we usually do not know what would have happened with an alternative intervention. This issue is known as the “Fundamental Problem of Causal Inference” [Rubin, 1974]. Due to this, direct application of data-based function approximation methods, such as machine learning or statistical analysis, is not feasible. As a result, various indirect methods have been devised to address this challenge.

In the case of observational data, that is, information collected without any deliberate intervention or controlled ex-

periment, causal inference faces unique challenges and approaches.

The first major challenge in conducting causal inference from observational data is the identifiability problem [Pearl, 2009]: establishing that the causal effect can be uniquely inferred from the data at hand. This requires ensuring that the data, model, and assumptions used are appropriate and robust enough to distinguish between correlation and causation.

Once identifiability is established, the next step is to articulate the target causal effect using estimable quantities. Machine learning techniques have proven highly effective in calculating the types of expressions found in these estimable quantities, playing a significant role in advancing various causal inference methodologies [Kaddour *et al.*, 2022].

During recent years, there has been an increasing emphasis on developing techniques that incrementally adapt machine learning models to the task of causal effects estimation. This area, where representation learning and neural networks play a vital role [Shalit *et al.*, 2017; Shi *et al.*, 2019; Tesei *et al.*, 2023], is the focus of our paper. In this work, we review and categorize these models based on their algorithmic inductive bias – the underlying assumptions a model relies on for making predictions or generalizations from its training data. We explore their relationship with various types of biases that need to be addressed for accurate treatment effect estimation.

Furthermore, we pinpoint and address a bias that is often overlooked in the literature on non-linear models: the occurrence of spurious interactions between variables within models. These interactions are not part of the underlying causal model and can create correlational shortcuts that alter the impact of causal relationships estimated by models, particularly in situations where data is limited.

The rest of this paper is organized as follows. In Section 2, we define the problem and review related work on causal effects estimation. In Section 3, we describe our proposed method and explain how it can be implemented in neural network models. In Section 4, we present the experiments we conducted to evaluate our method and compare it with other state-of-the-art estimators. In Section 5, we discuss the results and analysis of our experiments and suggest some possible directions for future research.

1.1 Preliminary concepts and setup

We want to estimate the causal effect of a treatment variable T on an outcome variable Y , controlling for a set of covariates X . We assume that the causal relationships between these variables can be represented by a **causal graph** [Koller and Friedman, 2009]. A causal graph is a directed acyclic graph (DAG) where each node is a variable and each edge is a direct causal effect. We assume that the causal graph satisfies the causal edges assumptions, which means that all parents are causes of their children. We denote the causal graph by G , the set of parents of node x_i by $Pa(x_i)$, the children by $Ch(x_i)$ and the ancestors by $An(x_i)$.

Beyond the causal graph, the underlying causal model can be modeled through a set of structural equations. A **Structural Causal Model** (SCM) [Pearl, 2009] consists of a set of endogenous variables U , a set of exogenous variables V and a set of functions $\{f_i\}$ to generate each endogenous variable x_i as a function of other variables. Every SCM implies an associated causal graph.

An **interaction** occurs when the effect of one variable on the outcome depends on the value of another variable. However, models for treatment effect estimation can consider interactions that may be spurious, lacking a real presence in any causal mechanism within the data generation process. Hence, we define a **spurious interaction** as one that does not manifest in any of the equations f_i within the SCM. We refer to non-spurious interactions as **causally valid interactions**.

In many cases, the complete SCM may not be fully known, but there might be some expert knowledge about the nature of the interactions between variables. In such situations, the method proposed in this paper can still be utilized by incorporating this partial information.

In this work, we will focus on causal effects that can be identified using the **backdoor criterion** [Pearl, 1993]. A set of variables X satisfies the backdoor criterion relative to T and Y if: 1) X blocks all backdoor paths from T to Y , and 2) X does not contain any descendants of T .

Adjustment sets and variable interactions

Even though the backdoor criterion provides a complete identification of all **valid adjustment sets** (sets of variables can be used to identify the causal effects, possibly including variables that are not strictly necessary), it does not tell us which is the **optimal adjustment set** for practical estimation in finite-sample scenarios.

Those variables in a valid adjustment set that are not strictly required for identification are called **neutral controls** [Cinelli *et al.*, 2022]. These variables neither increase nor decrease the asymptotic bias of the estimation of treatment effects, but they can still have a huge impact in a finite-sample scenario and can improve the performance of estimators. The optimal adjustment set usually contains neutral controls.

Given a valid adjustment set, its associated spurious interactions are the potential interactions between variables included in the set that do not appear in any of the equations of the SCM associated with the data.

In this work, we assume all our datasets are composed of valid adjustment sets, i.e. the backdoor criterion can be satisfied. There are advanced methods that, given an SCM, allow

to define **efficient adjustment sets** (adjustment sets that yield the least squares estimator with the smallest asymptotic variance) [Rotnitzky and Smucler, 2020], [Henckel *et al.*, 2022]. Although this scenario is not addressed in our work, it's important to note that these sets may also be prone to spurious interaction bias. Consequently, the method outlined in this paper is also applicable to them.

2 Problem definition and related work

We are interested in estimating the $f_Y(T, X)$ function that explains the causal effect of a given treatment T on a target variable Y considering a valid adjustment set X . In this paper, we name this the *causal effects function*.

In the binary case, let us define $Y(1)$ as the outcome in the presence of treatment and $Y(0)$ as the outcome in its absence. For a given individual i , the causal effects function, or Individual Treatment Effect function (ITE), is expressed as $ITE := f_Y(T, X) = f_Y(X) = Y_i(1) - Y_i(0)$. We are also interested in the Average Treatment Effect (ATE) and the Conditional Average Treatment Effect (CATE): $ATE := \mathbb{E}[Y_i(1) - Y_i(0)]$, $CATE := \mathbb{E}[Y_i(1) - Y_i(0)|X = x]$.

In order to be able to estimate these quantities, we need to obtain a statistical estimand that depends only on observed data using the process of identification. The identified estimand for the binary case using the backdoor criterion is:

$$\mathbb{E}_X[\mathbb{E}[Y|T = 1, X] - \mathbb{E}[Y|T = 0, X]] \quad (1)$$

This estimand is apparently similar to the common estimand used in prediction tasks $\mathbb{E}[Y|T, X]$, but model selection processes and inductive biases that work for one do not necessarily apply to the other [Doutreligne and Varoquaux, 2023]. Consequently, having an identified expression might not be enough to calculate accurate causal effects. The estimand that has been identified might be prone to issues that could make it difficult to estimate and hence render any non-asymptotic estimation impossible [Maclaren and Nicholson, 2019].

The step after causal identification is causal estimation: in this step, all available data is used to compute an association that reflects the causal effect if all assumptions are met.

2.1 Causal Estimation and Inductive biases for the causal inference task

The use of algorithmic inductive biases that encode prior information about the data in causal estimation methods has been one of the most successful strategies to build better causal inference models [Künzel *et al.*, 2019]. To understand these inductive biases, we need to identify the sources of error when computing causal effects. These sources of error can be categorized into two main categories: identification bias and estimation bias [Chatton and Rohrer, 2023].

In practical situations, it is not always possible to test the assumptions that establish a statistical estimand as a causal estimand. Certain assumptions, such as the absence of confounders, cannot be tested. Below, we present several biases that arise from the potential violation of the assumptions, known as **identification bias**:

- Exchangeability bias: Bias derived from the existence of hidden confounders that impact the identification process.

METHOD	IDENTIFICATION BIAS				ESTIMATION BIAS			
	Exchangeability	Positivity	No interference	Consistency	Data efficiency	Treatment relevance	Propensity	Spurious interactions
Linear Models [Thoresen, 2019]	×	×	×	×	✓	×	×	✓
S-learner [Künzel <i>et al.</i> , 2019]	×	×	×	×	✓	×	×	×
T-learner [Künzel <i>et al.</i> , 2019]	×	×	×	×	×	✓	×	×
X-learner [Künzel <i>et al.</i> , 2019]	×	×	×	×	✓	✓	×	×
BLearner [Oprescu <i>et al.</i> , 2023]	✓	×	×	×	✓	✓	✓	×
TARNet [Shalit <i>et al.</i> , 2017]	×	×	×	×	✓	✓	×	+ CGC \rightarrow ✓
CFRNet [Shalit <i>et al.</i> , 2017]	×	×	×	×	✓	✓	✓	+ CGC \dashrightarrow ✓?
Dragonnet [Shi <i>et al.</i> , 2019]	×	×	×	×	✓	✓	✓	+ CGC \rightarrow ✓
BCAUSS [Tesei <i>et al.</i> , 2023]	×	✓	×	×	✓	✓	✓	+ CGC \rightarrow ✓
HINITE [Lin <i>et al.</i> , 2023]	×	×	✓	×	✓	✓	✓	+ CGC \dashrightarrow ✓?

Table 1: Comparison of some relevant causal inference methods in terms of the biases they address. Checkmarks indicate that the method has a mechanism to address the bias indicated in the column. Methods included in the table have mechanisms that are easy to associate with a source of error. In this paper, we present CGC, a method that can be applied on top of other representation-based learners to address the spurious interactions bias. With a solid arrow, we represent the methods evaluated in this paper. With a dashed arrow, other representation-based methods where CGC could be applied.

- **Positivity bias:** This assumption can be violated if we do not collect any data for a subpopulation.
- **Consistency bias:** Bias derived from the violation of the consistency assumption, which states that if the treatment is T , then the observed outcome Y is the potential outcome under treatment T [Neal, 2020].
- **Non-interference bias:** Bias derived from the effect of one individual treatment on another individual. The existence of this bias is particularly common in network data [Lin *et al.*, 2023].

We can also consider additional inductive biases related to the finite nature of the data. These are called **estimation bias**:

- **Propensity bias:** A causal model must generalize from behavior under one set of conditions to behavior under another set [Pearl, 2009]. This bias appears when the performance of the model differs between distinct treatments as a consequence of a difference in the distribution of the data.
- **Relevance of treatment:** Bias that appears especially in high-dimensional settings. Models tend to neglect the role of the treatment and its effect is expressed through correlated features.
- **Functional form:** Any task performance can be affected by the functional form in which it is modeled. Linear tasks are better modeled by linear models, and monotonic tasks by monotonic models.
- **Variable interactions:** Spurious variable interactions do not explain any causal relationship between the variables, but a model might use them as a correlational shortcut, modifying the effect of causal relationships.
- **Data efficiency:** A model might not make an efficient use of the information present in the data.
- **Missing data:** Bias derived from the lack of availability of some data.
- **Measurement error on some variables.**

In Table 1 some of the most relevant causal inference methods are considered along with the biases they address. It shows which methods can correct or mitigate each type of bias. Using data efficiently and ensuring treatment relevance are among

the most critical sources of error and have been addressed by most models. In the table, we can also see that one of the most interesting properties of non-linear models is that they can stack multiple bias correction strategies to address several sources of bias at the same time.

2.2 State of the Art causal inference methods

Two of the most straightforward approaches for estimating the estimand represented by Eq.(1) are the S-Learner and the T-Learner [Künzel *et al.*, 2019]. S-learner is a method that trains a single machine learning model with the treatment variable as one of the features, along with other covariates. This model is then used to predict outcomes under both scenarios - with treatment and without treatment. The main disadvantage of S-learner is that it can bias the estimate of the treatment effect towards zero, especially when the feature space is large. Unlike the S-learner, the T-learner splits the data by treatment and fits a model for each group. This addresses the issue of failing to recognize a weak treatment variable, yet it may still be prone to data efficiency bias.

S-learner and T-learner are among the simplest causal effects estimators; however, apart from these approaches, more sophisticated methods have also been developed.

Most of the best performing methods in Table 1 rely on the concept of **representation learning** [Bengio *et al.*, 2013]. They use neural networks with an architecture that can be understood in two parts: pre-representation and post-representation. The first creates a representation from the input data and the second takes over from the representation to the desired output. One reason why explicitly dealing with representations is interesting is because they can be convenient to express many general priors about the world around us. These priors can be used to incorporate inductive biases to improve the causal inference task. Recent improvements in causal effect prediction methods focus primarily on modifications of the post-representation part: either by modifying the architecture of the output layers of the network or by using specific loss functions that carry inductive biases.

For example, TARNet [Shalit *et al.*, 2017] branches off into two separate heads for each treatment group to prevent the model from disregarding T . On top of that, Dragonnet [Shi *et al.*, 2019] adds a third head to the TARNet architecture to predict propensity $g(X) = P(T = 1|X)$. This approach

reduces the bias caused by having different modeling performances for each value of the treatment, although it decreases the predictive accuracy of the model. BCAUSS [Tesei *et al.*, 2023] uses the same architecture as the Dragonnet, but modifies the loss function so that it is more robust to positivity violations. In this approach, multiple bias correction methods are stacked to obtain state-of-the-art results. Other representation learning approaches encode other inductive biases to deal with different scenarios. CFRNet [Shalit *et al.*, 2017] and BNN [Johansson *et al.*, 2016] create a balanced representation of the data for both treated and untreated observations. Similarly, importance sampling weights can be added to the CFRNet to alleviate the problem of selection bias [Hassanpour and Greiner, 2019]. HINITE [Lin *et al.*, 2023] addresses no interference assumption violations, and in [Curth and van der Schaar, 2021] the authors create a representation that encodes prior information about the structure of causal effects.

Other works have suggested inductive biases that are challenging to link with a particular source of error. For example, CEVAE [Louizos *et al.*, 2017] uses variational autoencoders (VAE) to estimate a latent-variable model that simultaneously discovers hidden confounders and infers how they affect treatment and outcome. Another example is GANITE [Yoon *et al.*, 2018], which uses generative adversarial networks (GAN) to estimate treatment effects. This approach attempts to capture the uncertainty in the counterfactual distributions and is defined for any number of treatments.

2.3 Motivation

The source of bias derived from variable interactions has been studied in the context of linear models [Thoresen, 2019], [Harring *et al.*, 2015], where the inclusion of each interaction has to be evaluated. However, the discussion has not extended to non-linear models like neural networks, despite these models frequently rely on interactions between all variables.

In this paper we propose Neural Networks with Causal Graph Constraints (NN-CGC), a method to address the variable interactions bias using causal information like the one that might be present in the causal graph and apply it to neural network models.

3 A new Inductive Bias for Estimating Causal Effects

We propose a comprehensive methodology for incorporating an inductive bias that mitigates spurious interactions in neural network models. This approach can utilize information such as expert insights about the connections between specific variables, partial or complete knowledge of the causal graph, or a blend of both.

3.1 Constraining the learned distribution

Our main assumption is that the learned distribution has to be as close as possible to the distribution defined by the underlying causal model. To do so, we remove some of the cases that do not satisfy the conditional independence structure encoded by the causal graph. Specifically, we model the distribution of Y as a function of functions of variables that are causally related to each other.

For any node $x_i \in Pa(Y)$, we define the following set of variables: $G_{x_i} := (\{x_i\} \cup An(x_i)) \setminus \{T\}$.

Then, we propose to restrict the learned model $f_Y(X)$ to this family of functions:

$$f_Y(X) \sim f_Y\left(f(Pa(Y) \setminus \{T\}), f(G_{x_1}), \dots, f(G_{x_n})\right).$$

The variable T is excluded from all the groups because the learners implemented in this paper do not use T as an input but rather as part of the structure.

Using this method, all valid interactions are included, while all discarded interactions are of spurious nature. Note that we do not use causal paths that go from root nodes to the target as groups, as these would prevent some valid interactions. See an example in Figure 1.

In the case of not having access to the causal graph, we can define G_{x_i} by using expert knowledge about existing spurious interactions or by considering causal discovery methods. This approach can help limit our model’s reliance on at least some of the spurious interactions.

Groups of variables and Independent Causal Mechanisms

The causal generative process of a system variables is composed of autonomous modules that do not inform or influence each other, called Independent Causal Mechanisms (ICM) [Schölkopf *et al.*, 2021]. The groups defined in Section 3.1 contain all the information for modeling the ICMs for the target variable Y and for each of its parents $Pa(Y)$ given the adjustment set X . Several studies, such as those by [Xia *et al.*, 2021] and [Parafita and Vitrià, 2022], have investigated the process of modeling all Independent Causal Mechanisms (ICMs) to derive a Structural Causal Model (SCM). This comprehensive method carries its own challenges and advantages, making it a suitable option for certain scenarios. The methodology presented in this paper finds a balance between fully connected modeling and SCM-based modeling.

Following this objective, in the next section, we detail how this approach can be implemented using neural networks and how it can be combined with other neural network-based learners.

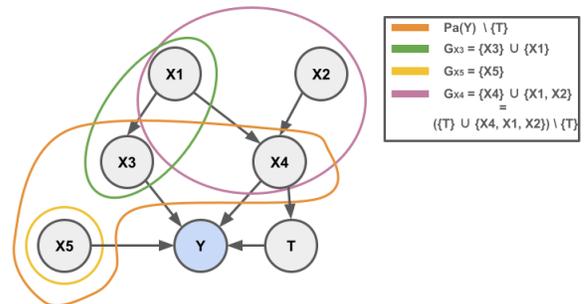


Figure 1: Illustrative example of groups of variables allowed to interact between each other. In this setting, the allowed interactions are: $(X1, X2, X4)$, $(X3, X1)$, $(X5, X3, X4)$. In this specific case, G_{x4} is equal to G_T .

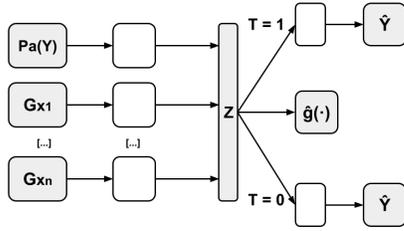


Figure 2: Model architecture when applying CGC to the Dragonnet. The post-representation part remains identical but the pre-representation layers are divided according to the groups of variables.

3.2 Implementation and Neural Network Architecture

Similarly to other methods related to representation learning ideas, our technique can be broken down into two parts. Initially, we define the inputs and the layers that lead to the representation. Subsequently, we analyze the representation layer and the output layers of the network. A pseudocode for the implementation can be found in Algorithm 1.

Pre-representation: Variable groups

The input of the network is divided into different groups of variables, as shown in Figure 2. We create a set of layers for each variable group G_{x_i} and $Pa(Y)$, that do not interact with each other. Each set of layers receives only the variables in the group as input. Note that a single variable can be part of more than one group. Each set of layers must be able to model $P(Y|G_{x_i})$ or $P(Y|Pa(Y))$. The output of all these layers is concatenated and used as input of the representation layer Z .

Post-representation: Linear layer and output heads

The representation layer Z is connected to the head of the network through a linear activation function. This ensures that the representation encoded in Z , which is used as input for the rest of the network, is free from spurious interactions.

Most state of the art methods modify the output of the network by changing the architecture after the representation layer, i.e. considering distinct heads or loss functions. In contrast, our methodology takes place in the architecture before the representation layer. This allows for a combination of both architectures, using our methodology combined with any of the output heads. In this paper we implemented the TARNet [Shalit *et al.*, 2017], Dragonnet [Shi *et al.*, 2019] and BCAUSS [Tesei *et al.*, 2023] heads, but other representation-based learners can be connected to the network.

4 Empirical testing

We are interested in determining whether models that are constrained using the proposed technique perform better than their unconstrained counterparts. Since the method can be applied to most representation-based learners, we test it using some of the best performing learners: TARNet, Dragonnet, and BCAUSS. Counting constrained and unconstrained versions, this gives us a total of six models to experiment with.

4.1 Experimental setup

Evaluating causal inference methods is hard because we usually have no access to counterfactual data. To overcome

Algorithm 1 Adding CGC procedure

Input: Train and test sets $(D_{tr}, D_{te}) \sim D$ with a treatment variable T , a target Y and a valid adjustment set X .

Parameter: A representation-based $f_Y(T, X)$ estimator NN .

- 1: Causal Knowledge Specification: Build or discover the associated causal graph using the data in D_{tr} . If this is not feasible, list all forbidden variable interactions by using expert knowledge.
 - 2: Groups: Create the groups of variables $G_{x_i} := (\{x_i\} \cup An(x_i)) \setminus \{T\} \forall x_i \in Pa(Y)$ and $Pa(Y) \setminus \{T\}$. If the causal graph is not available, form all compatible variable subsets by excluding forbidden interactions.
 - 3: Architecture pre-representation: Create an independent set of layers for each group, all connected to the representation layer Z .
 - 4: Architecture post-representation: Use the head associated to NN from the representation layer onward.
 - 5: Model: Run the causal inference workflow associated to NN using D_{tr} and D_{te} .
-

this problem, one of the common approaches is to use semi-synthetic data. We test the models on two of the most common benchmarks in causal inference: IHDP and JOBS. These benchmarks have already been used by some of the best performing methods and thus are specially relevant for the comparison of constrained and unconstrained models.

Since these benchmarks might present some problems [Curth *et al.*, 2021], we also test the models on a range of synthetic benchmarks that represent distinct situations, details in the synthetic experiments section. All the code for reproducing the experiments can be found here - <https://anonymous.4open.science/t/NN-CGC-FF9E/>.

Causal Graphs

All causal graphs in these experiments are discovered using the ICALINGAM algorithm [Shimizu *et al.*, 2006] implemented in [Zhang *et al.*, 2021]. The output of this algorithm usually involves only a subset of the variables. Since we are interested in the effect of using correct information rather than the quality of the discovery, we group the remaining variables into a single group of variables. Errors in the discovery might still translate to loss of performance of the constrained models (which can be seen in the higher-variance scenarios of the synthetic benchmarks). The choice of the discovery algorithm is out of the scope of this work, and ICALINGAM has been selected for its ability to converge under distinct scenarios. Given that the adjustment set is considered valid, it is assumed that any nodes associated with the target are in the parents of Y , as this would otherwise go against the backdoor criterion.

Hyperparameters

We use the same training hyperparameters defined in [Shi *et al.*, 2019] and [Tesei *et al.*, 2023], as they show to achieve state of the art results. These include a hidden layer size of 200 for the shared representation layers and 100 for the conditional outcome layers, train using stochastic gradient descent with

SCENARIO	NOISE	TARNET	DRAGONNET	BCAUSS
A	1/2	0.76	0.76	0.93
	1	0.91	0.92	1.02
	2	1.07	1.04	1.06
	4	1.08	1.11	1.02
B	1/2	0.73	0.72	0.63
	1	0.92	0.94	0.78
	2	1.02	1.06	0.94
	4	1.00	0.99	0.96
C	1/2	0.51	0.51	0.36
	1	0.71	0.76	0.7
	2	1.00	1.	1.09
	4	1.12	1.13	1.18
D	1/2	0.5	0.51	0.38
	1	0.76	0.76	0.72
	2	0.92	0.92	0.93
	4	1.00	1.00	1.01

Table 2: Table for PEHE Comparison. The reported value is the ratio between the constrained and the unconstrained models. Values < 1 mean the constrained version has a lower error than the unconstrained counterpart.

learning rate 10^{-5} and momentum 0.9, and batch size equal to 64 for the TARNet and Dragonnet and equal to the train-set length for BCAUSS. Details can be found in the code implementation and in appendix B.

In the constrained architectures, the hyperparameters for each set of layers are kept the same as those in the unconstrained architectures. However, the activation function of the representation layer is modified to a linear function.

Evaluation metrics

When counterfactual data is accessible, we use standard metrics such as the error in Average Treatment Effects -ATE- and the Precision in Estimation of Heterogeneous Effects -PEHE.

For the Jobs dataset, no counterfactuals are available, so we use the experimental nature of the data to compute the error of the average treatment effects on the treated ATT.

4.2 Experiments

Synthetic experiments

We use a set of synthetic experiments defined in [Nie and Wager, 2021] and implemented in [Chen *et al.*, 2020]. These experiments are composed of four scenarios: scenario A is a complex outcome regression model with an easy treatment effect, scenario B is a randomized controlled trial, scenario C has an easy propensity score with a difficult control outcome, and scenario D has an unrelated treatment arm and control arm. We run each scenario using the settings in the original paper: number of samples $n \in \{500, 1000\}$, number of variables $d \in \{6, 12\}$, and noise level $\sigma \in \{0.5, 1, 2, 4\}$. We generate an equal number of test samples as the number of train samples. Each configuration is repeated 100 times, making a total of 6400 runs. We average for the number of samples and number of variables for controlling the reporting size and show the results at the scenario-noise level because it contains the most information. Detailed results can be found in appendix A.

Synthetic experiments results. The results of the synthetic experiments are shown in the metric values Table 4 and in the

SCENARIO	NOISE	TARNET	DRAGONNET	BCAUSS
A	1/2	0.73	0.74	0.93
	1	0.83	0.81	1.02
	2	0.96	0.92	1.07
	4	1.01	1.04	1.02
B	1/2	0.87	0.88	0.69
	1	0.96	0.98	0.87
	2	0.95	0.99	0.98
	4	0.91	0.93	0.95
C	1/2	.58	0.58	0.17
	1	0.54	0.64	0.31
	2	0.70	0.72	0.54
	4	0.87	0.91	0.79
D	1/2	0.85	0.92	0.7
	1	0.98	1.01	0.88
	2	1.01	0.99	0.94
	4	1.02	1.03	1.07

Table 3: Table for ATE Comparison. The reported value is the ratio between the constrained and the unconstrained models. Values < 1 mean the constrained version has a lower error than the unconstrained counterpart.

MODEL	$\sqrt{\epsilon_{PEHE}}$				ϵ_{ATE}			
	A	B	C	D	A	B	C	D
TARNet	.363	.350	.456	.564	.285	.197	.316	.158
Dragonnet	.359	.349	.446	.568	.283	.196	.306	.160
BCAUSS	.413	.432	.424	.681	.346	.212	.405	.163
TARNet + CGC	.350	.334	.411	.475	.252	.182	.232	.158
Dragonnet + CGC	.347	.335	.408	.478	.250	.185	.238	.161
BCAUSS + CGC	.417	.366	.378	.523	.348	.195	.207	.157

Table 4: Test metrics on the synthetic benchmark averaged to the scenario level. Lower is better.

comparison Table 3 and Table 2, where each value is the ratio between the PEHE and ATE, respectively, of the constrained and unconstrained models. A value less than 1 indicates that the constrained model has a lower metric value than the unconstrained model and vice versa. We observe that constrained models generally perform better than their unconstrained counterparts. Among the models with CGC, Dragonnet + CGC achieves the lowest PEHE values in three out of four scenarios, while BCAUSS + CGC achieves the lowest ATE values in two out of four scenarios. TARNet + CGC performs slightly worse than the other two models with CGC but still better than the models without CGC. In settings with higher variance, with a noise value 4 for the ATEs and values 2 and 4 for the PEHEs, the unconstrained models performed similar or better than the constrained models. This suggests that in high-noise environments the discerning between spurious and valid interactions is more difficult, and constraining the model can be discarding causally valid information.

Semisynthetic experiments: IHDP

The Infant Health and Development Program (IHDP) is a study that used randomization to examine the effects of home visits by specialist doctors on the cognitive test scores of premature infants. The dataset¹ was first used by [Hill, 2011] to evaluate algorithms for estimating treatment effects. In order to create an observational dataset, non-random subsets of treated

¹ The dataset can be accessed at www.fredjo.com.

individuals were removed, resulting in selection bias. The outcomes in the dataset were generated using the original covariates and treatments. The dataset consists of 747 subjects and 25 variables. Following the recent literature, we used the simulated outcome implemented in the NPCI package [Dorie, 2016], which is composed of 1000 repetitions of the experiment. We averaged our results over 1000 train/validation/test splits with ratios 70/20/10.

MODEL	$\sqrt{\epsilon_{PEHE}}$ Value	Comp	ϵ_{ATE} Value	Comp
BNN [Johansson <i>et al.</i> , 2016]	2.2 ± .1	-	.37 ± .03	-
CFRW [Shalit <i>et al.</i> , 2017]	.71 ± .0	-	.25 ± .01	-
CEVAEs [Louizos <i>et al.</i> , 2017]	2.7 ± .1	-	.34 ± .01	-
GANITE [Yoon <i>et al.</i> , 2018]	1.9 ± .4	-	.43 ± .05	-
TARNet + CGC	.960 ± .032	0.78	.148 ± .004	1.04
TARNet	1.23 ± .046		.142 ± .004	
Dragonnet + CGC	.974 ± .034	0.797	.138 ± .003	.968
Dragonnet	1.22 ± .047		.143 ± .004	
BCAUSS + CGC	.663 ± .022	.805	.100 ± .003	.929
BCAUSS	.824 ± .023		.107 ± .003	

Table 5: Train metrics on the IHDP benchmark. The “Comp” value is the ratio between the constrained and the unconstrained model.

MODEL	$\sqrt{\epsilon_{PEHE}}$ Value	Comp	ϵ_{ATE} Value	Comp
BNN [Johansson <i>et al.</i> , 2016]	2.1 ± .1	-	.42 ± .03	-
CFRW [Shalit <i>et al.</i> , 2017]	.76 ± .0	-	.27 ± .01	-
CEVAEs [Louizos <i>et al.</i> , 2017]	2.6 ± .1	-	.46 ± .02	-
GANITE [Yoon <i>et al.</i> , 2018]	2.4 ± .4	-	.49 ± .05	-
TARNet + CGC	1.004 ± .039	0.783	.180 ± .006	0.884
TARNet	1.283 ± .052		.204 ± .008	
Dragonnet + CGC	1.027 ± .045	0.796	.190 ± .009	0.911
Dragonnet	1.290 ± .057		.209 ± .010	
BCAUSS + CGC	.741 ± .030	0.771	.132 ± .005	0.862
BCAUSS	.962 ± .035		.153 ± .006	

Table 6: Test metrics on the IHDP benchmark. The “Comp” value is the ratio between the constrained and the unconstrained model.

IHDP results. The constrained version of BCAUSS yields the best results in both ATE and PEHE metrics, with the unconstrained version of BCAUSS coming in second. In all cases, constrained models outperform unconstrained ones, obtaining lower errors and a similar or smaller variability in the results.

The PEHE and ATE metric values of the unconstrained models (TARNet, Dragonnet, and BCAUSS) were higher than those of the constrained models on both the train and test sets, indicating that the constraints reduced the estimation error and increased the generalization ability of the models.

Real data experiments: Jobs

The LaLonde Jobs dataset¹ [LaLonde, 1986] assesses the impact of job training as a treatment on income and employment status after training. We use the feature set defined in [Dehejia and Wahba, 2002]. Following [Shalit *et al.*, 2017], we combined the LaLonde experimental sample (297 treated, 425 control) with the PSID comparison group (2490 control). We averaged over 10 train/validation/test splits with a ratio of 62/18/20. Due to the small size of the benchmark we repeat the process 10 times and report the average of the results. For most of the runs, the discovery algorithm was unable to find a causal graph, and a default was used. The default causal graph

was obtained as the most common causal graph discovered across 100 runs.

MODEL	ϵ_{ATT}^{tr}		ϵ_{ATT}^{te}	
	Value	Comp	Value	Comp
BNN [Johansson <i>et al.</i> , 2016]	.04 ± .01	-	.09 ± .04	-
CFRW [Shalit <i>et al.</i> , 2017]	.04 ± .01	-	.09 ± .03	-
CEVAEs [Louizos <i>et al.</i> , 2017]	.02 ± .01	-	.03 ± .01	-
GANITE [Yoon <i>et al.</i> , 2018]	.01 ± .01	-	.06 ± .03	-
TARNet + CGC	.037 ± .003	0.945	.087 ± .008	1.022
TARNet	.039 ± .004		.085 ± .008	
Dragonnet + CGC	.054 ± .010	1.121	.097 ± .012	1.147
Dragonnet	.048 ± .004		.085 ± .008	
BCAUSS + CGC	.017 ± .006	0.288	.075 ± .007	0.796
BCAUSS	.058 ± .001		.094 ± .009	

Table 7: Train and test metrics on the Jobs benchmark. Lower is better. The “Comp” value is the ratio between the constrained and the unconstrained model.

Jobs results. Jobs is a small benchmark and the performance of the models varies greatly between runs. This characteristic is evidenced in the standard deviation values, which are in the same order as the mean in most cases. Because of that, we think the use of the benchmark rests upon its ability to test that models perform well on real data. In this sense, all models perform comparably well. The lowest ATT error for train is obtained by GANITE followed by the constrained version of BCAUSS while for test, the lowest error is obtained by CEVAE followed by GANITE. In this dataset, the discovery of the causal graph is especially unreliable, and this is reflected in the comparison between constrained and unconstrained models, which obtain similar results.

5 Discussion and Conclusion

We have presented NN-CGC, a novel method for incorporating causal information into the estimation of heterogeneous treatment effects. Our method leverages an inductive bias that reduces the error caused by spurious variable interactions and can be applied on top of other representation-based models. We have tested the effectiveness of our method using three different base models. The experiments indicate that constraining models using the described method leads to significant improvements, achieving new state-of-the-art results.

NN-CGC is flexible and performed well for all three base models. Following the idea of stacking mechanism to address sources of error, NN-CGC can be combined with other techniques to introduce additional inductive biases. We have also shown that our method is robust to imperfect causal graphs, and that using partial causal information is preferable to ignoring it. Having no restrictions on interactions is the same as assuming a causal graph that is fully connected, and in most cases there are at least some known edges that are not causal.

As a future work, we plan to enhance our method by using masking and a graphical conditioner, a neural network that prevents spurious interactions by design, to share some weights between the groups. This would reduce redundancy in the architecture and make a more efficient use of the data.

References

- [Bengio *et al.*, 2013] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.
- [Chatton and Rohrer, 2023] Arthur Chatton and Julia M Rohrer. The causal cookbook: Recipes for propensity scores, g-computation, and doubly robust standardization. 2023.
- [Chen *et al.*, 2020] Huigang Chen, Totte Harinen, Jeong-Yoon Lee, Mike Yung, and Zhenyu Zhao. Causallml: Python package for causal machine learning. *arXiv preprint arXiv:2002.11631*, 2020.
- [Cinelli *et al.*, 2022] Carlos Cinelli, Andrew Forney, and Judea Pearl. A crash course in good and bad controls. *Sociological Methods & Research*, page 00491241221099552, 2022.
- [Curth and van der Schaar, 2021] Alicia Curth and Mihaela van der Schaar. Nonparametric estimation of heterogeneous treatment effects: From theory to learning algorithms. In *International Conference on Artificial Intelligence and Statistics*, pages 1810–1818. PMLR, 2021.
- [Curth *et al.*, 2021] Alicia Curth, David Svensson, Jim Weatherall, and Mihaela van der Schaar. Really doing great at estimating cate? a critical look at ml benchmarking practices in treatment effect estimation. In *Thirty-fifth conference on neural information processing systems datasets and benchmarks track (round 2)*, 2021.
- [Dehejia and Wahba, 2002] Rajeev H Dehejia and Sadek Wahba. Propensity score-matching methods for nonexperimental causal studies. *Review of Economics and statistics*, 84(1):151–161, 2002.
- [Dorie, 2016] Vincent Dorie. Npci: Non-parametrics for causal inference. URL: <https://github.com/vdorie/npci>, 11:23, 2016.
- [Doutreligne and Varoquaux, 2023] Matthieu Doutreligne and Gaël Varoquaux. How to select predictive models for decision making or causal inference. Available at SSRN 4467871, 2023.
- [Harring *et al.*, 2015] Jeffrey R Harring, Brandi A Weiss, and Ming Li. Assessing spurious interaction effects in structural equation modeling: A cautionary note. *Educational and psychological measurement*, 75(5):721–738, 2015.
- [Hassanpour and Greiner, 2019] Negar Hassanpour and Russell Greiner. Counterfactual regression with importance sampling weights. In *IJCAI*, pages 5880–5887, 2019.
- [Henckel *et al.*, 2022] Leonard Henckel, Emilija Perković, and Marloes H Maathuis. Graphical criteria for efficient total effect estimation via adjustment in causal linear models. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(2):579–599, 2022.
- [Hill, 2011] Jennifer L Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011.
- [Johansson *et al.*, 2016] Fredrik Johansson, Uri Shalit, and David Sontag. Learning representations for counterfactual inference. In *International conference on machine learning*, pages 3020–3029. PMLR, 2016.
- [Kaddour *et al.*, 2022] Jean Kaddour, Aengus Lynch, Qi Liu, Matt J Kusner, and Ricardo Silva. Causal machine learning: A survey and open problems. *arXiv preprint arXiv:2206.15475*, 2022.
- [Koller and Friedman, 2009] Daphne Koller and Nir Friedman. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
- [Künzel *et al.*, 2019] Sören R Künzel, Jasjeet S Sekhon, Peter J Bickel, and Bin Yu. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the national academy of sciences*, 116(10):4156–4165, 2019.
- [LaLonde, 1986] Robert J LaLonde. Evaluating the econometric evaluations of training programs with experimental data. *The American economic review*, pages 604–620, 1986.
- [Lin *et al.*, 2023] Xiaofeng Lin, Guoxi Zhang, Xiaotian Lu, Han Bao, Koh Takeuchi, and Hisashi Kashima. Estimating treatment effects under heterogeneous interference. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 576–592. Springer, 2023.
- [Louizos *et al.*, 2017] Christos Louizos, Uri Shalit, Joris M Mooij, David Sontag, Richard Zemel, and Max Welling. Causal effect inference with deep latent-variable models. *Advances in neural information processing systems*, 30, 2017.
- [Maclaren and Nicholson, 2019] Oliver J Maclaren and Ru-anui Nicholson. What can be estimated? identifiability, estimability, causal inference and ill-posed inverse problems. *arXiv preprint arXiv:1904.02826*, 2019.
- [Neal, 2020] Brady Neal. Introduction to causal inference. *Course Lecture Notes (draft)*, 2020.
- [Nie and Wager, 2021] Xinkun Nie and Stefan Wager. Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika*, 108(2):299–319, 2021.
- [Oprescu *et al.*, 2023] Miruna Oprescu, Jacob Dorn, Marah Ghoummaid, Andrew Jesson, Nathan Kallus, and Uri Shalit. B-learner: Quasi-oracle bounds on heterogeneous causal effects under hidden confounding. *arXiv preprint arXiv:2304.10577*, 2023.
- [Parafita and Vitrà, 2022] Álvaro Parafita and Jordi Vitrà. Estimand-agnostic causal query estimation with deep causal graphs. *IEEE Access*, 10:71370–71386, 2022.
- [Pearl, 1993] Judea Pearl. Bayesian analysis in expert systems: comment: graphical models, causality and intervention. *Statistical Science*, 8(3):266–269, 1993.
- [Pearl, 2009] Judea Pearl. *Causality*. Cambridge university press, 2009.
- [Rotnitzky and Smucler, 2020] Andrea Rotnitzky and Ezequiel Smucler. Efficient adjustment sets for population average causal treatment effect estimation in graphical models.

The Journal of Machine Learning Research, 21(1):7642–7727, 2020.

- [Rubin, 1974] Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, 66(5):688, 1974.
- [Schölkopf *et al.*, 2021] Bernhard Schölkopf, Francesco Locatello, Stefan Bauer, Nan Rosemary Ke, Nal Kalchbrenner, Anirudh Goyal, and Yoshua Bengio. Toward causal representation learning. *Proceedings of the IEEE*, 109(5):612–634, 2021.
- [Shalit *et al.*, 2017] Uri Shalit, Fredrik D. Johansson, and David Sontag. Estimating individual treatment effect: generalization bounds and algorithms. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 3076–3085. PMLR, 06–11 Aug 2017.
- [Shi *et al.*, 2019] Claudia Shi, David Blei, and Victor Veitch. Adapting neural networks for the estimation of treatment effects. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [Shimizu *et al.*, 2006] Shohei Shimizu, Patrik O Hoyer, Aapo Hyvärinen, Antti Kerminen, and Michael Jordan. A linear non-gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7(10), 2006.
- [Tesei *et al.*, 2023] Gino Tesei, Stefanos Giampanis, Jingpu Shi, and Beau Norgeot. Learning end-to-end patient representations through self-supervised covariate balancing for causal treatment effect estimation. *Journal of Biomedical Informatics*, 140:104339, 2023.
- [Thoresen, 2019] Magne Thoresen. Spurious interaction as a result of categorization. *BMC medical research methodology*, 19(1):1–8, 2019.
- [Xia *et al.*, 2021] Kevin Xia, Kai-Zhan Lee, Yoshua Bengio, and Elias Bareinboim. The causal-neural connection: Expressiveness, learnability, and inference. *Advances in Neural Information Processing Systems*, 34:10823–10836, 2021.
- [Yoon *et al.*, 2018] Jinsung Yoon, James Jordon, and Mihaela Van Der Schaar. Ganite: Estimation of individualized treatment effects using generative adversarial nets. In *International conference on learning representations*, 2018.
- [Zhang *et al.*, 2021] Keli Zhang, Shengyu Zhu, Marcus Kallander, Ignavier Ng, Junjian Ye, Zhitang Chen, and Lujia Pan. gcastle: A python toolbox for causal discovery. *arXiv preprint arXiv:2111.15155*, 2021.