
Finding Valid Adjustments under Non-ignorability with Minimal DAG Knowledge

Abhin Shah
MIT
abhin@mit.edu

Karthikeyan Shanmugam
IBM Research
karthikeyan.shanmugam2@ibm.com

Kartik Ahuja
Mila
kartik.ahuja@mila.quebec

Abstract

Treatment effect estimation from observational data is a fundamental problem in causal inference. There are two very different schools of thought that have tackled this problem. On the one hand, the Pearlian framework commonly assumes structural knowledge (provided by an expert) in the form of directed acyclic graphs and provides graphical criteria such as the back-door criterion to identify the valid adjustment sets. On the other hand, the potential outcomes (PO) framework commonly assumes that all the observed features satisfy ignorability (i.e., no hidden confounding), which in general is untestable. In prior works that attempted to bridge these frameworks, there is an observational criteria to identify an *anchor variable* and if a subset of covariates (not involving the anchor variable) passes a suitable conditional independence criteria, then that subset is a valid back-door. Our main result strengthens these prior results by showing that under a different expert-driven structural knowledge — that one variable is a direct causal parent of the treatment variable — remarkably, testing for subsets (not involving the known parent variable) that are valid back-doors is *equivalent* to an invariance test. Importantly, we also cover the non-trivial case where the entire set of observed features is not ignorable (generalizing the PO framework) without requiring the knowledge of all the parents of the treatment variable. Our key technical idea involves generation of a synthetic

sub-sampling (or environment) variable that is a function of the known parent variable. In addition to designing an invariance test, this sub-sampling variable allows us to leverage Invariant Risk Minimization, and thus, connects finding valid adjustments (in non-ignorable observational settings) to representation learning. We demonstrate the effectiveness and tradeoffs of these approaches on a variety of synthetic datasets as well as real causal effect estimation benchmarks.

1 INTRODUCTION

Estimating the impact of a treatment (or an action) is fundamental to many scientific disciplines (e.g., economics (Imbens and Rubin, 2015), medicine (Shalit et al., 2017; Alaa and van der Schaar, 2017), policy making (LaLonde, 1986; Smith and Todd, 2005)). In most of these fields, randomized clinical trials (RCT) is a common practice for estimating treatment effects. However, conducting a RCT could be unethical or costly, and we may only have access to observational data. Estimating treatment effects with only observational data is a challenging task and is of central interest to causal inference researchers.

A fundamental question in treatment effect estimation is: *Which subset of observed features should be adjusted for while estimating treatment effect from observational data?* Simpson’s paradox (Pearl, 2014), which is a phenomenon that is observed in many real-life studies on treatment effect estimation, underscores the value of selecting appropriate features for treatment effect estimation. Over the years, two schools of thoughts have formed on how to tackle treatment effect estimation. The Pearlian framework (Pearl, 2009) commonly assumes that an expert provides us with the causal generative model in the form of a directed acyclic graph (DAG) that relates unobserved exogenous variables to observed features, treatment variable, and outcome variables. With the knowledge

of the DAG available, the framework provides different graphical criteria (e.g., back-door criterion (Pearl, 1993), front-door criterion (Pearl, 1995)) that answers whether a subset is valid for adjustment. The DAG framework allows for the existence of confounders – unobserved variables that affect multiple observed variables. The potential outcomes (PO) framework (Rubin, 1974) makes an untestable assumption called *ignorability* – the assumption (in a rough sense) requires potential outcomes under different treatments be independent of the treatment conditioned on all (or a known subset of) observed features. In other words, ignorability implies that a subset of observed features is a valid adjustment and is known. The PO framework provides various techniques (e.g., inverse propensity weighing (Swaminathan et al., 2016), doubly robust estimation (Funk et al., 2011)) for treatment effect estimation under ignorability. One can view the Pearlian DAG framework as providing graphical criteria implying ignorability of certain subsets.

In summary, the Pearlian framework requires the knowledge of the DAG and the PO framework assumes ignorability with respect to the observed features. Motivated by the limitations of both of these frameworks we ask: *can we significantly reduce the structural knowledge required about the DAG under non-ignorability of observed features and yet find valid adjustment sets?*

1.1 Our Contributions

We assume the following minimal expert-driven local structural knowledge: a known observed feature is a direct causal parent of the treatment. Given this, we propose a simple *invariance* test, and show that it is equivalent to testing if a subset not involving the known parent satisfies the back-door criterion (without requiring ignorability) when the features are pre-treatment. To design our invariance test, we use the known parent to create ‘fake environment variables’. We then test for invariance (across these environments) of the outcome conditioned on subsets of observed features (not containing the known parent) and the treatment. If a subset passes this invariance test, then it satisfies the back-door criterion (and therefore is a valid adjustment set) allowing for treatment effect estimation. Crucially, our result also goes in the other direction, i.e., if there exists a set (not containing the known parent) that satisfies the back-door criterion, then it will pass our invariance test.

We propose two algorithms based on this equivalence result to identify valid adjustments. In the first algorithm, we use a subset based search procedure that exploits conditional independence (CI) testing to check our invariance criterion. As is standard with

any subset based search approach, the application of our first algorithm is limited to small dimensional datasets. To overcome this, in our second algorithm, we leverage Invariant Risk Minimization (IRM) (Arjovsky et al., 2019), originally proposed to learn causal representations for out-of-distribution generalization, to act as a continuous optimization based scalable approximation for CI testing. We demonstrate the effectiveness of our algorithms in treatment effect estimation on both synthetic and benchmark datasets. In particular, we also show that IRM based algorithm scales well with dimensions in contrast to the subset search based approach. The source code of our implementation is available at <https://github.com/Abhin02/invariance-via-subsampling>.

1.2 Related Work

Next, we provide an overview of related work that directly concerns finding valid adjustment in treatment effect estimation. See Appendix B for an overview of prior work related to potential outcomes and usage of representation learning to debias treatment effect.

Finding valid adjustment with global knowledge. Finding valid adjustment sets for general interventional queries has been extensively studied in the Pearlian framework (Tian and Pearl, 2002). Given the complete knowledge of the DAG, a sound and complete algorithm to find valid adjustments was proposed by Shpitser and Pearl (2008). When only the observational equivalence class is known, i.e., partial ancestral graph or PAG (Zhang, 2008), Perkovic et al. (2018) provided a sound and complete algorithm for finding valid adjustments. VanderWeele and Shpitser (2011) showed that if a valid adjustment set exists amongst the observed features, then the union of all observed parents of outcome and all observed parents of treatment is also a valid adjustment set. However, they required global knowledge i.e., information about *every* observed feature while our work requires knowledge about *only one observed parent of treatment* i.e., local knowledge.

Finding valid adjustment with local knowledge. As opposed to the works described in the previous paragraph, another line of work (e.g., Entner et al. (2013); Cheng et al. (2020); Gultchin et al. (2020)) focused on finding valid adjustment sets by exploiting local knowledge of the DAG. In Entner et al. (2013), a two-step approach was proposed. First, an *anchor variable* is characterized by an observational criteria that is testable. Next, a conditional independence test is performed on the subsets not involving the anchor variable to find the valid adjustment set. In the reverse direction, if a valid adjustment set exists that does not contain the anchor variable, their test

is shown to succeed only if the anchor variable has no observed or unobserved parents. As a result, even if it were possible to carry out consistent treatment effect estimation based on adjustment sets not involving the anchor, their procedure need not necessarily enable it. In contrast, in these settings, under the assumption that the anchor variable (direct causal parent of the treatment) is specified by the expert, our invariance test enables consistent treatment effect estimation. On the other hand, in Cheng et al. (2020), the anchor variable is characterized by topological properties of the PAG. We provide examples where our procedure can correctly declare that consistent treatment effect is not possible but they cannot.

Following Entner et al. (2013), Gultchin et al. (2020) proposed a fully-differentiable optimization framework to find a representation of the features that passes the conditional independence criteria analogous to Entner et al. (2013). While their approach avoids the brute-force subset search required by Entner et al. (2013), their approach is as limited in the reverse direction as Entner et al. (2013). Further, their continuous optimization framework assumes the outcome is binary or the whole system (including the treatment) is linear Gaussian. Additionally, they use partial correlation as a proxy for conditional independence. This proxy is correct when the underlying distribution is Gaussian and in the worst-case constrains only the second moment. In other words, their framework doesn't provide formal guarantees even if one of the variables (e.g. treatment) isn't Gaussian. In contrast, our approach doesn't make these assumptions and is more general.

Invariance principle. The invariance principle (also known as modularity condition) is fundamental to causal bayesian networks (Bareinboim et al., 2012; Schölkopf, 2019). Arjovsky et al. (2019) proposed a continuous optimization framework called invariant risk minimization (IRM), to search for causal representations which satisfy invariance principle, that achieves out-of-distribution generalization. A recent line of work (e.g., Shi et al. (2020); Shah et al. (2021)) has focused on using IRM for treatment effect estimation. Shi et al. (2020) assumed (i) that there are no unmeasured confounders and (ii) access to interventional data is available (similar to IRM). We significantly differ from this as we allow unmeasured confounders and do not require interventional data – we create artificial environments by sub-sampling observational data – and leverage IRM to find valid adjustment sets that satisfy our criterion. On the other hand, while Shah et al. (2021) created environments artificially (similar to ours), their sub-sampling procedure lacks theoretical justification.

Further, they focus primarily on the setting where there is little support overlap between the control and the treatment group, and lack formal guarantees on finding valid adjustment sets.

2 PROBLEM FORMULATION

Notations. For a sequence of deterministic variables s_1, \dots, s_n , we let $\mathbf{s} := \{s_1, \dots, s_n\}$. For a sequence of random variables s_1, \dots, s_n , we let $\mathbf{s} := \{s_1, \dots, s_n\}$. Let $\mathbb{1}$ denote the indicator function.

2.1 Semi-Markovian Model, Effect Estimation, Valid Adjustment

Consider a causal effect estimation task with \mathbf{x} as the feature set, t as the observed treatment variable and y as the observed potential outcome. For the ease of exposition, we focus on binary t . However, our results apply to non-binary t as well. Further, while we consider discrete \mathbf{x} and y , our framework applies equally to continuous or mixed \mathbf{x} and y . Let \mathcal{G} denote the underlying DAG over the set of vertices $\mathcal{W} := \{\mathbf{x}, t, y\}$. For any variable $w \in \mathcal{W}$, let $\pi(w)$ denote the set of parents of w i.e., $\pi(w) = \{w_1 : w_1 \rightarrow w\}$.

To estimate the causal effect of treatment t on outcome y , a Markovian causal model requires the specification of the following three elements : (a) \mathcal{W} – the set of variables, (b) \mathcal{G} – the DAG over the set of vertices \mathcal{W} , and (c) $\mathbb{P}(w|\pi(w))$ – the conditional probability of w given its parents $\pi(w)$ for every $w \in \mathcal{W}$. Given the DAG \mathcal{G} , the causal effect of t on y can be estimated from observational data since $\mathbb{P}(w|\pi(w))$ is estimable from observational data whenever \mathcal{W} is observed.

Our ability to estimate the causal effect of t on y from observational data is severely curtailed when some variables in a Markovian causal model are unobserved. Let $\mathbf{x}^{(o)} \subseteq \mathbf{x}$ be the subset of features that are observed and $\mathbf{x}^{(u)} = \mathbf{x} \setminus \mathbf{x}^{(o)}$ be the subset of features that are unobserved. For any variable $w \in \mathcal{W}$, let $\pi^{(o)}(w) \subseteq \pi(w)$ denote the set of parents of w that are observed and let $\pi^{(u)}(w) := \pi(w) \setminus \pi^{(o)}(w)$ denote the set of parents of w that are unobserved. We focus on the semi-Markovian causal model (Tian and Pearl, 2002), defined below, since any causal model with unobserved variables can be mapped to a semi-Markovian causal model while preserving the dependencies between the variables (Verma and Pearl, 1990; Acharya et al., 2018).

Definition 1. (*Semi-Markovian Causal Model.*) A semi-Markovian causal model \mathcal{M} is a tuple $\langle \mathcal{V}, \mathcal{U}, \mathcal{G}, \mathbb{P}(v|\pi^{(o)}(v)), \pi^{(u)}(v), \mathbb{P}(\mathcal{U}) \rangle$ where:

1. \mathcal{V} is the set of observed variables, i.e. $\mathcal{V} = \{\mathbf{x}^{(o)}, t, y\}$,

2. \mathcal{U} is the set of unobserved (or exogenous) features, i.e. $\mathcal{U} := \mathcal{W} \setminus \mathcal{V} = \mathbf{x}^{(u)}$,
3. \mathcal{G} is the DAG over the set of vertices \mathcal{W} such that each member in \mathcal{U} has no parents and at-most two children.
4. $\mathbb{P}(\mathbf{v}|\pi^{(o)}(\mathbf{v}), \pi^{(u)}(\mathbf{v})) \forall \mathbf{v} \in \mathcal{V}$ is the set of unobserved conditional distributions of the observed variables, and
5. $\mathbb{P}(\mathcal{U})$ is the unobserved joint distribution over the unobserved features.

In a semi-Markovian model, unobserved variables with only one or no children are omitted entirely. See Figure 1 for a toy example of a semi-Markovian model with $\mathcal{V} = \{x_1, x_2, x_3, t, y\}$, $\mathcal{U} = \{u_1, u_2, u_3, u_4\}$, and $\mathcal{G} = \mathcal{G}^{toy}$.




Figure 1: The toy example \mathcal{G}^{toy} .

In observational data, we observe samples of \mathcal{V} from $\mathbb{P}(\mathcal{V})$ which is related to the semi-Markovian model by the following marginalization (Tian and Pearl, 2002): $\mathbb{P}(\mathcal{V}) = \mathbb{E}_{\mathbf{x}^{(u)}} [\prod_{\mathbf{v} \in \mathcal{V}} \mathbb{P}(\mathbf{v}|\pi^{(o)}(\mathbf{v}), \pi^{(u)}(\mathbf{v}))]$. Next, we define the notion of causal effect using the do-operator.

Definition 2. (Causal Effect.) The causal effect of the treatment t on the outcome y is defined as

$$\mathbb{P}(y|do(t = t)) = \sum_{t=t', \mathbf{x}^{(o)} = \mathbf{x}^{(o)}} \mathbb{1}_{t'=t} \mathbb{E}_{\mathbf{x}^{(u)}} \left[\prod_{\mathbf{v} \in \mathcal{V} \setminus \{t\}} \mathbb{P}(\mathbf{v}|\pi^{(o)}(\mathbf{v}), \pi^{(u)}(\mathbf{v})) \right]$$

The do-operator forces t to be t in the causal model \mathcal{M} , i.e the conditional factor $\mathbb{P}(t = t'|\pi^{(o)}(t), \pi^{(u)}(t))$ is replaced by the indicator $\mathbb{1}_{t=t'}$ and the resulting distribution is marginalized over all possible realizations of all observed variables except y . Next, we define average treatment effect and valid adjustment.

Definition 3. (Average Treatment Effect.) The average treatment effect (ATE) of a binary treatment t on the outcome y is defined as $ATE = \mathbb{E}[y|do(t = 1)] - \mathbb{E}[y|do(t = 0)]$.

Definition 4. (Valid Adjustment.) A set of variables $\mathbf{z} \subseteq \mathbf{x}$ is said to be a valid adjustment relative to the ordered pair of variables (t, y) in the DAG \mathcal{G} if $\mathbb{P}(y|do(t = t)) = \mathbb{E}_{\mathbf{z}}[\mathbb{P}(y|\mathbf{z} = \mathbf{z}, t = t)]$.

If $\mathbf{z} \subseteq \mathbf{x}^{(o)}$ is a valid adjustment relative to (t, y) , then the ATE can be estimated from observational data by regressing the factual outcomes for the treated and the

untreated sub-populations on \mathbf{z} i.e., $ATE = \mathbb{E}_{\mathbf{z}}[\mathbb{E}_y[y|t = 1, \mathbf{z}] - \mathbb{E}_y[y|t = 0, \mathbf{z}]]$.

For any variables $w_1, w_2 \in \mathcal{W}$, and a set $\mathbf{w} \subseteq \mathcal{W}$, (a) let $w_1 \perp_p w_2|\mathbf{z}$ denote that w_1 and w_2 are conditionally independent given \mathbf{z} and (b) let $w_1 \perp\!\!\!\perp_d w_2|\mathbf{w}$ denote that w_1 and w_2 are d-separated by \mathbf{w} in \mathcal{G} . For completeness, we provide the definition of d-separation in Appendix D as well as review potential outcomes (PO) framework (Imbens and Rubin (2015)), discuss ignorability and connect it with valid adjustment in Appendix C.

2.2 Back-door Criterion

We now discuss the *back-door* criterion (Pearl et al., 2016) – a popular sufficient graphical criterion for finding valid adjustments i.e., any set satisfying the back-door criterion is a valid adjustment (Pearl (1993)).

Definition 5. (Back-door criteria.) A set of variables $\mathbf{z} \subseteq \mathbf{x}$ satisfies the back-door criterion relative to the ordered pair of variables (t, y) in \mathcal{G} if no node in \mathbf{z} is a descendant of t and \mathbf{z} blocks every path between t and y in \mathcal{G} that contains an arrow into t .

Often, \mathcal{G} is represented without explicitly showing elements of \mathcal{U} but, instead, using bi-directed edges (Tian and Pearl (2002)) to represent confounding effects of \mathcal{U} . For example, Figure 2(a) uses bi-directed edges to represent unmeasured confounders (i.e., elements of \mathcal{U} that influence two variables in \mathcal{V}) in the DAG \mathcal{G}^{toy} .

Definition 6. (A Bi-directed Edge.) A bi-directed edge between nodes $v_1 \in \mathcal{V}$ and $v_2 \in \mathcal{V}$ (i.e., $v_1 \leftrightarrow v_2$) represents the presence (in \mathcal{G}) of a divergent path $v_1 \leftarrow u \rightarrow v_2$ where $u \in \mathcal{U}$.

In this work, we make the following structural assumption on the DAG \mathcal{G} under the semi-Markovian model \mathcal{M} . This assumption is analogous to the common assumption that all observed features are pre-treatment variables. As an example, consider the DAG \mathcal{G}^{toy} in Figure 2(a) that satisfies this assumption.

Assumption 1. Let the DAG \mathcal{G} be such that the treatment t has the outcome y as its only child. Further, the outcome y has no child.

3 MAIN RESULTS

In this section, we state our main results relating sub-sampling and invariance testing to the back-door criteria. First, we define the notions of sub-sampling and invariance. Next, we provide : (a) a sufficient d-separation condition (that can be realized by our invariance test under sub-sampling) for a class of back-door criteria (Theorem 3.1) and (b) a necessary d-separation condition (that can be realized by our




Figure 2: The toy example \mathcal{G}^{toy} : (a) with bi-directed edges; (b) where e has been sub-sampled using x_1 and t ; (c) where e has been sub-sampled using x_3 and t .

invariance test under sub-sampling) implied by a class of back-door criteria (Theorem 3.2). Combining these, we show equivalence between an invariance based d-separation condition and a class of back-door criteria (Corollary 1). Finally, we propose an algorithm to find all subsets of the observed features that satisfy the back-door criteria when all the parents of the treatment variable are known and observed (Appendix I).

Sub-sampling. We create a sub-sampling (or environment) variable e from the observed distribution $\mathbb{P}(\mathcal{V})$. Formally, we use a specific observed variable $x_t \in \mathbf{x}^{(o)}$ and a subset of the observed variables $\mathbf{v} \subseteq \mathcal{V} \setminus \{x_t, y\}$ to sub-sample e i.e., $e = f(x_t, \mathbf{v}, \eta)$ where η is a noise variable independent of \mathcal{W} , and f is a function of x_t, \mathbf{v} and η . The choices of x_t and \mathbf{v} , which differ for the sufficient condition (Theorem 3.1) and the necessary condition (Theorem 3.2), are made clear in the respective theorem statements. We let the sub-sampling variable e be discrete and think of the distinct values of e as identities of distinct artificial environments created via sub-sampling. While the case where e is continuous is similar in spirit, we postpone the nuances for a future work. Graphically, sub-sampling variable introduces a node e , an edge from x_t to e and edges from every $v \in \mathbf{v}$ to e in the DAG \mathcal{G} . For example, see Figure 2(b) where e is sub-sampled in toy example \mathcal{G}^{toy} with $x_t = x_1$ and $\mathbf{v} = \{t\}$.

Invariance testing. Our main results relate the back-door criterion to d-separation statements of the type $e \perp\!\!\!\perp_d y | \mathbf{z}$ for some $\mathbf{z} \subseteq \mathcal{V} \setminus \{y\}$. While our goal is to infer sets satisfying the back-door criterion from observational data, such d-separation statements cannot be tested for from observational data. To tackle this, we propose the notion of invariance testing. An invariance test is a conditional independence test of the form $e \perp_p y | \mathbf{z}$ for some $\mathbf{z} \subseteq \mathcal{V} \setminus \{y\}$ i.e., an invariance test tests if the sub-sampling variable is independent of the outcome conditioned on \mathbf{z} for some $\mathbf{z} \subseteq \mathcal{V} \setminus \{y\}$.

For our results involving invariance testing, we require the following limited set of faithfulness assumptions to ensure invariance testing with e is equivalent to d-separation statements involving e .

Assumption 2. (Sub-sampling Faithfulness) If $e \perp_p$

$y | \mathbf{z}$, then $e \perp\!\!\!\perp_d y | \mathbf{z}$, $\forall \mathbf{z} \subseteq \mathcal{V} \setminus \{y\}$.

Thus, in effect, we create *synthetic* environments and show that a class of back-door criterion either implies or is equivalent to a suitable invariance test. For our framework to work, we only require the knowledge of x_t from an expert. This is in contrast to any detailed knowledge of the structure of the DAG \mathcal{G} .

Sufficient condition. Suppose an expert provides us with an observed feature that has a direct edge or a bi-directed edge to the treatment. Let e be sub-sampled using this feature as x_t and any $\mathbf{v} \subseteq \mathcal{V} \setminus \{x_t, y\}$. The following result shows that any subset of the remaining observed features satisfying a d-separation involving e (or an invariance test under Assumption 2) also satisfies the back-door criterion. See Appendix F for a proof.

Theorem 3.1. Let Assumption 1 be satisfied. Consider any $x_t \in \mathbf{x}^{(o)}$ that has a direct edge or a bi-directed edge to t i.e., either $x_t \rightarrow t$, $x_t \leftrightarrow t$ or $x_t \leftarrow t$. Let e be sub-sampled using x_t and \mathbf{v} for any $\mathbf{v} \subseteq \mathcal{V} \setminus \{x_t, y\}$ i.e., $e = f(x_t, \mathbf{v}, \eta)$. Let $\mathbf{z} \subseteq \mathbf{x}^{(o)} \setminus \{x_t\}$. If e is d-separated from y by \mathbf{z} and t in \mathcal{G} i.e., $e \perp\!\!\!\perp_d y | \mathbf{z}, t$ in \mathcal{G} , then \mathbf{z} satisfies the back-door criterion relative to (t, y) in \mathcal{G} .

Remark 1. A stronger result that subsumes Theorem 3.1 was proven in Entner et al. (2013); we provide our theorem for clarity of exposition and completeness.

Necessary condition. Suppose an expert provides us with an observed feature that has a direct edge to the treatment. Let e be sub-sampled using this variable as x_t and any $\mathbf{v} \subseteq \{t\}$. The following result shows that any subset of the remaining observed features satisfying the back-door criterion satisfies a specific d-separation involving e (as well as an invariance test). See Appendix G for a proof.

Theorem 3.2. Let Assumption 1 be satisfied. Consider any $x_t \in \mathbf{x}^{(o)}$ that has a direct edge to t i.e., $x_t \rightarrow t$ or $x_t \leftarrow t$. Let e be sub-sampled using x_t and \mathbf{v} for any $\mathbf{v} \subseteq \{t\}$ i.e., $e = f(x_t, \mathbf{v}, \eta)$. Let $\mathbf{z} \subseteq \mathbf{x}^{(o)} \setminus \{x_t\}$. If \mathbf{z} satisfies the back-door criterion relative to (t, y) in \mathcal{G} , then e is d-separated from y by \mathbf{z} and t in \mathcal{G} i.e., $e \perp\!\!\!\perp_d y | \mathbf{z}, t$ in \mathcal{G} .

Remark 2. Theorem 3.2 is useful to find out (some)

sets that cannot be valid adjustments (see comparison with [Entner et al. \(2013\)](#) and [Gultchin et al. \(2020\)](#) as well as comparison with [Cheng et al. \(2020\)](#) below). Knowing whether a given set of features is valid for adjustment or not is crucial – especially in healthcare and social sciences – to avoid using decisions based on biased estimates from observational studies.

Remark 3. We note that [Theorem 3.2](#) requires x_t to be a parent of t (i.e., a direct edge to t) whereas [Theorem 3.1](#) requires x_t to be a parent of t or a spouse of t (i.e., a direct or a bi-directed edge to t).

Equivalence. Suppose an expert provides us with a feature that has a direct edge to the treatment. Let e be sub-sampled using this variable as x_t and any $\mathbf{v} \subseteq \{t\}$. Combining [Theorem 3.1](#) and [Theorem 3.2](#), we have the following Corollary showing equivalence of the back-door criterion and a specific d-separation involving e (as well as an invariance test under [Assumption 2](#)).

Corollary 1. Let [Assumption 1](#) be satisfied. Consider any $x_t \in \mathbf{x}^{(o)}$ that has a direct edge to t i.e., $x_t \rightarrow t$ or $x_t \leftrightarrow t$. Let e be sub-sampled using x_t and \mathbf{v} for any $\mathbf{v} \subseteq \{t\}$ i.e., $e = f(x_t, \mathbf{v}, \eta)$. Let $\mathbf{z} \subseteq \mathbf{x}^{(o)} \setminus \{x_t\}$. Then, \mathbf{z} satisfies back-door criterion relative to the ordered pair of variables (t, y) in \mathcal{G} if and only if e is d-separated from y by \mathbf{z} and t in \mathcal{G} i.e., $e \perp_d y | \mathbf{z}, t$ in \mathcal{G} .

Remark 4. While our framework captures a broad class of back-door criteria, it does not cover all the back-door criteria. For example, our method cannot capture that the M-bias problem ([Liu et al., 2012](#); [Imbens, 2020](#)) where no observed feature is a parent of the treatment. (see [Appendix H](#) for details).

Illustrative examples. First, we illustrate [Corollary 1](#) with our toy example \mathcal{G}^{toy} . We let $x_t = x_1$ and sub-sample e using x_1 and t (see [Figure 2\(b\)](#)). For this example, $\mathbf{z} \subseteq \{x_2, x_3\}$ i.e., $\mathbf{z} \in \{\emptyset, \{x_2\}, \{x_3\}, \{x_2, x_3\}\}$. It is easy to verify that $\mathbf{z} = \{x_2\}$ satisfies the back-door criterion relative to (t, y) in \mathcal{G}^{toy} but $\mathbf{z} = \emptyset$, $\mathbf{z} = \{x_3\}$, and $\mathbf{z} = \{x_2, x_3\}$ do not. Similarly, it is easy to verify that $e \perp_d y | x_2, t$ but $e \not\perp_d y | t$, $e \not\perp_d y | x_2, t$, and $e \not\perp_d y | x_2, x_3, t$ in \mathcal{G}^{toy} . See [Appendix F.2](#) for an illustration tailored to [Theorem 3.1](#).

Next, we illustrate the significance of the criteria that qualifies x_t to our results. In [Figure 2\(b\)](#), we let $x_t = x_1$ (x_1 has a direct or bi-directed edge to t). Here, the d-separation $e \perp_d y | x_2, t$ holds implying that x_2 satisfies back-door relative to (t, y) . In [Figure 2\(c\)](#), we let $x_t = x_3$ (x_3 does not have a direct or bi-directed edge to t). Here, the d-separation $e \perp_d y | x_3, t$ holds but x_3 does not satisfy the back-door relative to (t, y) .

Comparison with [Entner et al. \(2013\)](#) and [Gultchin et al. \(2020\)](#). In [Entner et al. \(2013\)](#), x_a is an anchor variable if it satisfies the observational

criterion $x_a \not\perp_d y | \mathbf{z}$ for some x_a and some \mathbf{z} not containing x_a . Further, if the CI test implied by (the d-separation condition) $x_a \perp_d y | \mathbf{z}, t$ is satisfied, then \mathbf{z} is shown to be a valid adjustment. While our sufficient condition in [Theorem 3.1](#) is implied by this result, we provide a proof tailored to our condition and notations in [Appendix F](#) for completeness.

However, the reverse direction in [Entner et al. \(2013\)](#) is as follows: if some \mathbf{z} (not containing x_a) is a valid adjustment, then $x_a \perp_d y | \mathbf{z}, t$, only when x_a does not have any (observed or unobserved) parent in addition to satisfying the criteria for the anchor variable. Under our criterion, if x_t is a direct parent of t , the reverse direction can be shown in generality (our [Theorem 3.2](#)).

As a concrete example, in \mathcal{G}^{toy} , [Entner et al. \(2013\)](#) cannot conclude that $\emptyset, \{x_3\}, \{x_2, x_3\}$ are not admissible i.e., not valid back-doors (because $x_a = x_1$ has an unobserved parent) while our [Theorem 3.2](#) can be used to conclude that. See the empirical comparison in [Appendix K.7](#). Likewise, [Gultchin et al. \(2020\)](#), which build on [Entner et al. \(2013\)](#), also cannot conclude that $\emptyset, \{x_3\}, \{x_2, x_3\}$ are not valid adjustment sets in \mathcal{G}^{toy} .

Comparison with [Cheng et al. \(2020\)](#). In [Cheng et al. \(2020\)](#), the anchor variable x_a is a *COSO variable* i.e., either a parent or a spouse of the treatment but neither a parent or a spouse of the outcome in the true maximal ancestral graph (MAG). Our criteria for x_t is different from this, and our result is neither implied by nor implies the result of [Cheng et al. \(2020\)](#).

Consider an example which is obtained by adding the edge $x_1 \rightarrow y$ to \mathcal{G}^{toy} in [Figure 2](#). The results of [Cheng et al. \(2020\)](#) are not applicable since the anchor variable x_t is a parent of the outcome in the true DAG (and thereby in the MAG). However, x_t is a parent of the treatment (i.e., it satisfies our criteria), and our [Theorem 3.2](#) is applicable. It can be used to conclude that $\emptyset, \{x_2\}, \{x_3\}$ and $\{x_2, x_3\}$ are not admissible sets. See the empirical comparison in [Appendix K.8](#).

Connections to Instrument Variable (IV). While our anchor (i.e., x_t) may look similar to IV, this is not the case: (i) An IV needs to satisfy the exclusion restriction i.e., it needs to be d-separated from y in \mathcal{G}_{-t} (i.e., the graph obtained by removing the edge from t to y in \mathcal{G}). However, we do not require x_t to be d-separated from y in \mathcal{G}_{-t} . (ii) Unlike our work, IVs can *only* provide bounds on ATE in non-parametric models; they provide perfect identifiability of ATE only in linear models ([Balke and Pearl, 1997](#)).

4 ALGORITHMS

Our invariance criterion in [Corollary 1](#) requires us to find a \mathbf{z} such that $e \perp_p y | \mathbf{z}, t$. In this section, given

n observational samples, we propose two algorithms that enable finding valid adjustment sets that pass our invariance criterion as well as use it to estimate ATE.

4.1 Invariance Testing and Subset Search

First, we propose an algorithm (Algorithm 1) based on conditional independence (CI) testing and it works as follows. The algorithm takes the sub-sampling variable e that is a function of \mathbf{x}_t (e could also be a function of both \mathbf{x}_t and t). The algorithm considers the set \mathcal{X} of all candidate adjustment sets that do not contain \mathbf{x}_t . For every candidate adjustment set \mathbf{z} in \mathcal{X} , our algorithm checks for CI between e and y conditioned on \mathbf{z} and t . If this CI holds, then \mathbf{z} satisfies the back-door criterion and is a valid adjustment set (see Corollary 1 and Assumption 2). The ATE estimated by our algorithm is the average of the ATE estimated by regressing on such valid adjustment sets. On actual datasets, we use the following criterion as acceptance for CI: a p-value threshold p_{value} is used to check if the p-value returned by the CI tester is greater than this threshold. We use the RCoT CI tester (see Appendix K.1).

Similar to Entner et al. (2013), the computational complexity of Algorithm 1 grows exponentially in the dimensionality of $\mathbf{x}^{(o)}$. This makes it impractical for high dimensional settings.

Algorithm 1: ATE estimation using subset search.

Input: $n, n_r, t, y, e, \mathcal{X}, p_{value}$

Output: ATE(\mathcal{X})

Initialization: ATE(\mathcal{X}) = 0, $c_1 = 0$

```

1 for  $r = 1, \dots, n_r$  do // Use a different
   train-test split in each run
2    $c_2 = 0$ ; ATEd = 0;
3   for  $\mathbf{z} \in \mathcal{X}$  do
4     if  $CI(e \perp_p y | \mathbf{z}, t) > p_{value}$  then
5        $c_2 = c_2 + 1$ ;
6       ATEd = ATEd +  $\frac{1}{n} \sum_{i=1}^n (\mathbb{E}[y | \mathbf{z} = \mathbf{z}^{(i)}, t = 1] - \mathbb{E}[y | \mathbf{z} = \mathbf{z}^{(i)}, t = 0])$ ;
7   if  $c_2 > 0$  then
8      $c_1 = c_1 + 1$ ;
9     ATE( $\mathcal{X}$ ) = ATE( $\mathcal{X}$ ) + ATEd/ $c_2$ ;
10 ATE( $\mathcal{X}$ ) = ATE( $\mathcal{X}$ )/ $c_1$ ;
```

4.2 IRM based Representation Learning

To alleviate these concerns, we propose a second algorithm based on invariant risk minimization (IRM). This leverages our use of the subsampling variable and creation of synthetic environments. IRM was proposed to address out-of-distribution generalization for supervised learning tasks and is aimed at learning

a predictor that relies only on the causal parents of the label y and ignore any other spurious variables. IRM takes data from different environments indexed as e and learns a representation Φ that transforms the features \mathbf{x} such that $e \perp y | \Phi(\mathbf{x})$. Given that our invariance criterion is of a similar form, and involves checking invariance of the outcome y conditioned on the feature set \mathbf{z} and the treatment t across environments e , IRM is a perfect fit to test this criterion.

Our IRM based procedure leverages IRMv1 from Arjovsky et al. (2019) with linear representation Φ . We take the data in treatment group $t = 1$ (or the control group $t = 0$) and divide it into different environments based on e and pass it as input to IRMv1. From the theory of IRM it follows that if the absolute value of some coefficient of Φ is low, then the corresponding component is unlikely to be a part of the subset that satisfies the invariance criterion. Following this observation, we define a vector of absolute values of Φ and denote it as $|\Phi|$. We divide the values in $|\Phi|$ into two clusters using k-means clustering with $k = 2$. We select the subset of the features that correspond to the cluster with a higher mean absolute value. We estimate the treatment effect by adjusting over this selected subset. Further details of the procedure can be found in Algorithm 2 (we describe the algorithm for treatment group and can run a similar procedure for control group). While the computational complexity of IRMv1 (and hence Algorithm 2) is unclear yet, in practice, Algorithm 2 is much faster and scales better (see Figure 3(c)) than Algorithm 1.

Algorithm 2: ATE estimation using IRM

Input: $n, n_r, t, y, e, \mathbf{x}^{(o)} \setminus x_t$

Output: ATE

Initialization: ATE = 0, $k = 2$

```

1 for  $r = 1, \dots, n_r$  do // Use a different
   train-test split in each run
2    $\Phi \leftarrow \text{IRMv1}(y, \mathbf{x}^{(o)} \setminus x_t, e, t = 1)$ 
3    $\mathbf{z}_{\text{irm}} \leftarrow \text{kmeans}(|\Phi|, k)$  //  $\mathbf{z}_{\text{irm}}$  is the subset of
   variables in the cluster with higher mean
   absolute value
4   ATE = ATE +  $\frac{1}{n} \sum_{i=1}^n (\mathbb{E}[y | \mathbf{z}_{\text{irm}} = \mathbf{z}^{(i)}, t = 1] - \mathbb{E}[y | \mathbf{z}_{\text{irm}} = \mathbf{z}^{(i)}, t = 0])$ 
5 ATE = ATE/ $n_r$ ;
```

5 EXPERIMENTS

ATE estimation and performance metrics: To test how successful our method is with respect to finding valid adjustments, we consider estimating the ATE of t on y . When the ground truth ATE is known, we report the absolute error in ATE prediction (averaged over n_r




Figure 3: Validating our theoretical results and our algorithms on the toy example \mathcal{G}^{toy} : (a) Sets not satisfying back-door ($\{x_1, x_2, x_3\}$, $\{x_2, x_3\}$) result in high ATE error; sets satisfying back-door ($\{x_1, x_2\}$, $\{x_2\}$) result in low ATE error. (b) Performance of Algorithm 1 and 2 on \mathcal{G}^{toy} . (c) Performance of Algorithm 2 in high dimensions.

runs). When the ground truth ATE is unknown, we report the estimated ATE (averaged over n_r runs).

As described in Section 2.1, ATE can be estimated from observational data by regressing y for the control and the treatment sub-populations on a valid adjustment set. We note that our work is complementary to works on ATE estimation as our focus is on finding valid adjustments. Once we select a valid adjustment, any of the available ATE estimation methods could be used. We use ridge regression with cross-validation as the regression model for baseline as well as our method.

Environment variable and parameters. For all of our experiments, we let $n_r = 100$ and $p_{value} = \{0.1, 0.2, 0.3, 0.4, 0.5\}$. For our experiments we create an environment variable as being a random function of x_t and t (i.e., $e = f(x_t, t)$). Exact details of their generation and alternate settings, such as the case of $e = f(x_t)$ (i.e., $\mathbf{v} = \emptyset$), are given in Appendix K.

Algorithms. We compare the following algorithms:

1. **Baseline:** This uses regression on all of the observed features i.e., $\mathbf{x}^{(o)}$ to estimate ATE. In other words, it assumes $\mathbf{x}^{(o)}$ is ignorable. See Appendix J for a pseudo-code of Baseline.
2. **Exhaustive:** Given x_t , this applies Algorithm 1 with \mathcal{X} being the set of all subsets of $\mathbf{x}^{(o)} \setminus x_t$.
3. **Sparse:** Given x_t , this applies Algorithm 1 with \mathcal{X} being the set of all subsets of $\mathbf{x}^{(o)} \setminus x_t$ of size at most k (which is determined in the context).
4. **IRM-t:** Given x_t , this applies Algorithm 2 to the samples from the treatment group.
5. **IRM-c:** Given x_t , this applies Algorithm 2 to the samples from the control group.

5.1 Synthetic Experiment

Description. Consider the toy example \mathcal{G}^{toy} from Figure 2 with unobserved features $u_1 \in \mathbb{R}$, $u_2 \in \mathbb{R}^{\tilde{d}}$, $u_3 \in \mathbb{R}^{\tilde{d}}$, $u_4 \in \mathbb{R}^{\tilde{d}}$ and observed features $x_1 \in \mathbb{R}$, $x_2 \in \mathbb{R}^{\tilde{d}}$, $x_3 \in \mathbb{R}^{\tilde{d}}$ i.e., $\mathbf{x}^{(u)} = \{u_1, u_2, u_3, u_4\} \in \mathbb{R}^{3\tilde{d}+1}$ and $\mathbf{x}^{(o)} =$

$\{x_1, x_2, x_3\} \in \mathbb{R}^{2\tilde{d}+1}$. Let $d = 2\tilde{d} + 1$ i.e., the dimension of the observed features. For dimension d , we generate a dataset (with $n = 50000$) using linear structural equation models for u 's, x 's and y and a logistic linear model for t and e . See Appendix K.2 for details

Results. First, we validate our theoretical results for $d = 5, 15, 25$ (see Figure 3(a)): (a) the ATE error for adjusting on $\{x_1, x_2, x_3\}$ is high since we are in a setting where $\mathbf{x}^{(o)}$ is *not ignorable*, (b) the ATE error for adjusting on $\{x_1, x_2\}$ is low since it satisfies the back-door criterion, (c) the ATE error for adjusting on $\{x_2, x_3\}$ is high since $e \not\perp_d y | x_2, x_3, t$, (d) the ATE error for adjusting on $\{x_2\}$ is low since $e \perp_d y | x_2, t$. Next, we validate our algorithms via Figure 3(b). With $x_t = x_1$, our algorithms Exhaustive, IRM-t, and IRM-c significantly outperform Baseline for $d = 3, 5, 7$ even for multiple p_{value} thresholds for Exhaustive. We note that IRM based algorithms significantly outperform the testing based algorithm even in moderately high dimensions ($d = 7$) and performs very well even for $d = 65$ as seen through in Figure 3(c).

5.2 Semi-synthetic Dataset : Infant Health and Development Program (IHDP)

Description. IHDP (Hill, 2011) is generated based on a RCT targeting low-birth-weight, premature infants. The 25-dimensional feature set (comprising of 17 different features) is pre-treatment i.e., it satisfies Assumption 1. The features measure various aspects about the children and their mothers e.g., child's birth-weight, the number of weeks pre-term that the child was born. See Appendix K.4 for details. In the treated group, the infants were provided with both intensive high-quality childcare and specialist home visits. A biased subset of the treated group is typically removed to create imbalance leaving 139 samples with $t = 1$ and 608 samples with $t = 0$. The outcome, typically simulated using setting "A" of the NPCI package (Dorie, 2016), is infants' cognitive test score.

Analysis. The outcome depends on all observed features. Therefore, the set of all observed features satisfies back-door (see Appendix K.4). To test our method, we drop 7 features and denote the resulting 16-dimensional feature set (comprising of 10 features) by $\mathbf{x}^{(o)}$ to create a challenging non-ignorable case. We use *child’s birth-weight* as x_t . Therefore, we keep this feature in $\mathbf{x}^{(o)}$. See Appendix K.4 for the choice of other features in $\mathbf{x}^{(o)}$.

Results. We compare **Baseline**, **Exhaustive**, **Sparse** with $k = 5$, **IRM-c** and **IRM-t**. All our algorithms except **IRM-c** significantly outperform **Baseline** (see Figure 4). The intuition behind $k = 5$ is the belief that valid adjustments of size 5 exist (see Appendix K.4)¹.

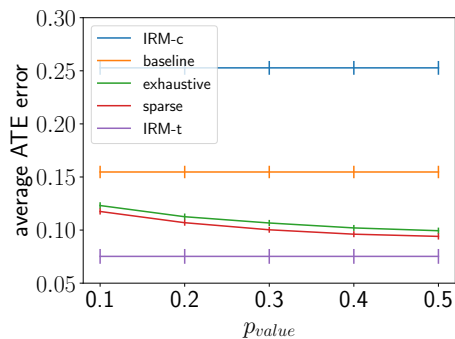


Figure 4: Performance on IHDP dataset.

5.3 Real Dataset : Cattaneo2

Description. Cattaneo2 (Cattaneo, 2010) studies the effect of maternal smoking on babies’ birth weight. The 20 observed features measure various attributes about the children, their mothers and their fathers. See Appendix K.5 for details. The dataset considers the maternal smoking habit during pregnancy as the treatment i.e., $t = 1$ if smoking (864 samples) and $t = 0$ if not smoking (3778 samples).

Analysis. Out of the features we have access to (see Appendix K.5), we pick mother’s age to be x_t .

Results. The ground truth ATE is unknown (because for every sample either y_0 or y_1 is observed). However, the authors in Almond et al. (2005) expect a strong negative effect of maternal smoking on the weights of babies – about 200 to 250 grams lighter for a baby with a mother smoking during pregnancy. We compare all the algorithms except **Exhaustive** with $x_t =$ mother’s age. For the **sparse** algorithm, we set $k = 5$ to ensure a reasonable run-time. As seen in Figure 5, the ATE estimated using all our algorithms fall in the desired interval (i.e., (-250,-200)) and suggest a larger negative

effect compared to the **Baseline**.




Figure 5: Performance on Cattaneo2 dataset.

6 CONCLUSION AND DISCUSSION

We showed that it is possible to find valid adjustment sets under non-ignorability with the knowledge of a single causal parent of the treatment. We achieved this by providing an invariance test that exactly identifies all the subsets of observed features (not involving this parent) that satisfy the back-door criterion.

Knowledge of a causal parent of the treatment.

Our invariance test depends on the causal parent of the treatment i.e., x_t via the environment variable i.e., e . Therefore, our approach works even when the expert knowledge of x_t is not available or samples of x_t are not observed so long as we have samples of e directly. Investigating the application of this insight is an interesting question for future research.

Assumption 1 and 2. Assumption 1 and faithfulness (a stronger version of Assumption 2) are commonly used in data-driven covariate selection works (Entner et al., 2013; Gultchin et al., 2020; Cheng et al., 2020). While settings beyond Assumption 1 are interesting for future research, finding valid adjustments under Assumption 1 is non-trivial and important in both PO and Pearlian framework (see the first paragraph in VanderWeele and Shpitser (2011)). Further, we note that Assumption 1 holds for some benchmark causal effect estimation datasets (e.g., IHDP, Twins). Lastly, while it is common to assume faithfulness with respect to conditional independencies involving the entire DAG, we assume faithfulness only with respect to conditional independencies involving the sub-sampling variable.

Alternate minimal DAG knowledge. As discussed in Remark 4, our method doesn’t cover all back-door criteria (e.g., the M-bias problem). Therefore, exploring alternate minimal DAG knowledge sufficient to test for a broader/different family of valid adjustments could be fruitful.

¹We note that **Sparse** still has to perform $\sum_{i=0}^5 \binom{9}{i} = 382$ tests to estimate ATE. Therefore, **Sparse** performs not very differently from **Exhaustive**.

Acknowledgements

We thank the anonymous reviewers of NeurIPS 2021 for bringing to our notice the works of [Entner et al. \(2013\)](#) and [Cheng et al. \(2020\)](#) as well as for several suggestions. We also thank the anonymous referees of AISTATS 2022 for their comments and feedback. Kartik Ahuja acknowledges the support provided by IVADO postdoctoral fellowship funding program.

References

- A. Abadie, D. Drukker, J. L. Herr, and G. W. Imbens. Implementing matching estimators for average treatment effects in stata. *The stata journal*, 4(3):290–311, 2004.
- A. Abadie, A. Diamond, and J. Hainmueller. Synthetic control methods for comparative case studies: Estimating the effect of california’s tobacco control program. *Journal of the American statistical Association*, 105(490):493–505, 2010.
- J. Acharya, A. Bhattacharyya, C. Daskalakis, and S. Kandasamy. Learning and testing causal models with interventions. *Advances in Neural Information Processing Systems*, 31, 2018.
- A. M. Alaa and M. van der Schaar. Bayesian inference of individualized treatment effects using multi-task gaussian processes. *arXiv preprint arXiv:1704.02801*, 2017.
- D. Almond, K. Y. Chay, and D. S. Lee. The costs of low birth weight. *The Quarterly Journal of Economics*, 120(3):1031–1083, 2005.
- M. Arjovsky, L. Bottou, I. Gulrajani, and D. Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.
- A. Balke and J. Pearl. Bounds on treatment effects from studies with imperfect compliance. *Journal of the American Statistical Association*, 92(439):1171–1176, 1997.
- E. Bareinboim, C. Brito, and J. Pearl. Local characterizations of causal bayesian networks. In *Graph Structures for Knowledge Representation and Reasoning*, pages 1–17. Springer, 2012.
- M. D. Cattaneo. Efficient semiparametric estimation of multi-valued treatment effects under ignorability. *Journal of Econometrics*, 155(2):138–154, 2010.
- D. Cheng, J. Li, L. Liu, K. Yu, T. D. Lee, and J. Liu. Towards unique and unbiased causal effect estimation from data with hidden variables. *arXiv preprint arXiv:2002.10091*, 2020.
- V. Dorie. Npci: Non-parametrics for causal inference. 2016. URL <https://github.com/vdorie/npci>.
- D. Entner, P. Hoyer, and P. Spirtes. Data-driven covariate selection for nonparametric estimation of causal effects. In *Artificial Intelligence and Statistics*, pages 256–264. PMLR, 2013.
- M. J. Funk, D. Westreich, C. Wiesen, T. Stürmer, M. A. Brookhart, and M. Davidian. Doubly robust estimation of causal effects. *American journal of epidemiology*, 173(7):761–767, 2011.
- L. Gultchin, M. Kusner, V. Kanade, and R. Silva. Differentiable causal backdoor discovery. In *International Conference on Artificial Intelligence and Statistics*, pages 3970–3979. PMLR, 2020.
- J. L. Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011.
- G. W. Imbens. Potential outcome and directed acyclic graph approaches to causality: Relevance for empirical practice in economics. *Journal of Economic Literature*, 58(4):1129–79, 2020.
- G. W. Imbens and D. B. Rubin. Rubin causal model. In *Microeconometrics*, pages 229–241. Springer, 2010.
- G. W. Imbens and D. B. Rubin. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press, 2015.
- F. Johansson, U. Shalit, and D. Sontag. Learning representations for counterfactual inference. In *International conference on machine learning*, pages 3020–3029. PMLR, 2016.
- N. Kallus. Deepmatch: Balancing deep covariate representations for causal inference using adversarial training. In *International Conference on Machine Learning*, pages 5067–5077. PMLR, 2020.
- S. R. Künzel, J. S. Sekhon, P. J. Bickel, and B. Yu. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the national academy of sciences*, 116(10):4156–4165, 2019.
- R. J. LaLonde. Evaluating the econometric evaluations of training programs with experimental data. *The American economic review*, pages 604–620, 1986.
- W. Liu, M. A. Brookhart, S. Schneeweiss, X. Mi, and S. Setoguchi. Implications of m bias in epidemiologic studies: a simulation study. *American journal of epidemiology*, 176(10):938–948, 2012.
- J. Pearl. [bayesian analysis in expert systems]: Comment: graphical models, causality and intervention. *Statistical Science*, 8(3):266–269, 1993.
- J. Pearl. Causal diagrams for empirical research. *Biometrika*, 82(4):669–688, 1995.
- J. Pearl. *Causality*. Cambridge university press, 2009.

- J. Pearl. Comment: understanding simpson’s paradox. *The American Statistician*, 68(1):8–13, 2014.
- J. Pearl, M. Glymour, and N. P. Jewell. *Causal inference in statistics: A primer*. John Wiley & Sons, 2016.
- E. Perkovic, J. Textor, M. Kalisch, and M. H. Maathuis. Complete graphical characterization and construction of adjustment sets in markov equivalence classes of ancestral graphs. 2018.
- P. R. Rosenbaum. Optimal matching for observational studies. *Journal of the American Statistical Association*, 84(408):1024–1032, 1989.
- P. R. Rosenbaum and D. B. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.
- P. R. Rosenbaum and D. B. Rubin. Constructing a control group using multivariate matched sampling methods that incorporate the propensity score. *The American Statistician*, 39(1):33–38, 1985.
- D. B. Rubin. Matching to remove bias in observational studies. *Biometrics*, pages 159–183, 1973.
- D. B. Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, 66(5):688, 1974.
- B. Schölkopf. Causality for machine learning. *arXiv preprint arXiv:1911.10500*, 2019.
- A. Shah, K. Ahuja, K. Shanmugam, D. Wei, K. R. Varshney, and A. Dhurandhar. Treatment effect estimation using invariant risk minimization. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5005–5009. IEEE, 2021.
- U. Shalit, F. D. Johansson, and D. Sontag. Estimating individual treatment effect: generalization bounds and algorithms. In *International Conference on Machine Learning*, pages 3076–3085. PMLR, 2017.
- C. Shi, D. M. Blei, and V. Veitch. Adapting neural networks for the estimation of treatment effects. *arXiv preprint arXiv:1906.02120*, 2019.
- C. Shi, V. Veitch, and D. Blei. Invariant representation learning for treatment effect estimation. *arXiv preprint arXiv:2011.12379*, 2020.
- Y. Shimoni, E. Karavani, S. Ravid, P. Bak, T. H. Ng, S. H. Alford, D. Meade, and Y. Goldschmidt. An evaluation toolkit to guide model selection and cohort definition in causal inference. *arXiv preprint arXiv:1906.00442*, 2019.
- I. Shpitser and J. Pearl. Complete identification methods for the causal hierarchy. *Journal of Machine Learning Research*, 9:1941–1979, 2008.
- J. A. Smith and P. E. Todd. Does matching overcome lalonde’s critique of nonexperimental estimators? *Journal of econometrics*, 125(1-2):305–353, 2005.
- E. V. Strobl, K. Zhang, and S. Visweswaran. Approximate kernel-based conditional independence tests for fast non-parametric causal discovery. *Journal of Causal Inference*, 7(1), 2019.
- A. Swaminathan, A. Krishnamurthy, A. Agarwal, M. Dudík, J. Langford, D. Jose, and I. Zitouni. Off-policy evaluation for slate recommendation. *arXiv preprint arXiv:1605.04812*, 2016.
- J. Tian and J. Pearl. A general identification condition for causal effects. In *Aaai/iaai*, pages 567–573, 2002.
- C. Uhler, G. Raskutti, P. Bühlmann, and B. Yu. Geometry of the faithfulness assumption in causal inference. *The Annals of Statistics*, pages 436–463, 2013.
- T. J. VanderWeele and I. Shpitser. A new criterion for confounder selection. *Biometrics*, 67(4):1406–1413, 2011.
- T. Verma and J. Pearl. Causal networks: Semantics and expressiveness. In *Machine intelligence and pattern recognition*, volume 9, pages 69–76. Elsevier, 1990.
- S. Wager and S. Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018.
- J. Yoon, J. Jordon, and M. Van Der Schaar. Ganite: Estimation of individualized treatment effects using generative adversarial nets. In *International Conference on Learning Representations*, 2018.
- J. Zhang. Causal reasoning with ancestral graphs. *Journal of Machine Learning Research*, 9:1437–1474, 2008.

Supplementary Material: Finding Valid Adjustments under Non-ignorability with Minimal DAG Knowledge

Organization. In Appendix A we briefly discuss any potential societal impacts of our work. In Appendix B, we discuss prior work related to potential outcomes and usage of representation learning to debias treatment effect. In Appendix C, we review potential outcomes (PO) framework, discuss ignorability and connect it with valid adjustment. In Appendix D, we provide the definition of d-separation as well as a few related definitions. In Appendix E, we provide a few additional notations. In Appendix F, we provide a proof of Theorem 3.1 and also provide an illustrative example for Theorem 3.1. In Appendix G, we provide a proof of Theorem 3.2. In Appendix H, we provide a discussion on the M-bias problem. In Appendix I, we provide an Algorithm (Algorithm 3) that, when all the parents of the treatment are observed and known, finds all subsets of the observed features satisfying the back-door criterion relative to (t, y) in \mathcal{G} as promised in Section 3. We also provide an example illustrating Algorithm 3 and the associated result via Corollary 2. In Appendix J, we provide an implementation of the **Baseline** ATE estimation routine considered in this work. In Appendix K, we discuss the usage of real-world CI testers in Algorithm 1, provide more discussions on experiments from Section 5, specify all the training details, as well as provide more details regarding the comparison of our method with Entner et al. (2013), Gultchin et al. (2020), and Cheng et al. (2020).

A SOCIETAL IMPACT

In health-care scenarios, since it is sometimes difficult/unethical to do randomized control trials (RCTs), sometimes the consensus treatment protocol is decided based on observational studies. Our algorithm could pick out a correct valid adjustment set when some existing methods assume ignorability due to lack of expert knowledge about the causal model.

On the flip side, due to lower testing power at finite samples or mis-identification of a feature as a direct parent of the treatment (a local causal knowledge required in our work), our algorithm could pick an incorrect valid adjustment set. This, in turn, could potentially result in miscalculation of the treatment effect. The consensus treatment protocols based on such observational conclusions could prove detrimental. However, we emphasize that this is a risk associated with most (if not all) observational studies and effect estimation algorithms.

B ADDITIONAL RELATED WORK

Potential Outcomes framework. Potential outcomes (PO) framework formalizes the notion of ignorability as a condition on the observed features that is sufficient (amongst others) for valid adjustment in treatment effect estimation (Imbens and Rubin, 2010). Various methods like propensity scoring (Rosenbaum and Rubin (1983)), *matching* (Rosenbaum and Rubin (1985)) of the treatment group and the control group based on features that satisfy ignorability, and synthetic control methods (Abadie et al. (2010)) have been used to debias effect estimation. In another line of work (Wager and Athey (2018); Künzel et al. (2019); Alaa and van der Schaar (2017)), treatment effect was estimated by regressing the outcome on the treated and the untreated sub-populations. While this list of works on the PO framework is by no means exhaustive, in a nutshell, these methods can be seen as techniques to estimate the treatment effect when a valid adjustment set is given.

Representation learning based techniques. Following the main idea behind matching (Rubin (1973); Abadie et al. (2004); Rosenbaum (1989)), recent methods inspired by deep learning and domain adaptation, used a neural network to transform the features and then carry out matching in the representation space (Shi et al. (2019); Shalit et al. (2017); Johansson et al. (2016); Yoon et al. (2018); Kallus (2020)). These methods aimed to correct the lack of overlap between the treated and the control groups while assuming that the representation learned is ignorable (i.e., a valid adjustment).

C REVIEW OF POTENTIAL OUTCOMES AND IGNORABILITY

We briefly review the potential outcomes (PO) framework in the context of treatment effect estimation (Imbens and Rubin, 2015). In the PO framework, there are exogenous variables called *units*. With a slight abuse of notation, we denote them by $\mathbf{x}^{(u)}$ as well. When $\mathbf{x}^{(u)}$ is fixed to say $\mathbf{x}^{(u)}$, the observed variables (including y) are deterministically fixed i.e., only the randomness in the units induces randomness in the observed variables. The PO framework typically studies the setup where the observed features $\mathbf{x}^{(o)}$ are pre-treatment (similar to semi-Markovian model under Assumption 1). Every observational sample $(\mathbf{x}^{(o)}, t, y)$ has an associated unit $\mathbf{x}^{(u)}$. For $t' \in \{0, 1\}$, the *potential outcome* $y_{t'}$ is the resulting outcome for the unit $\mathbf{x}^{(u)}$ when the treatment t is set (by an intervention) to t' .

Definition 7. (*Ignorability.*) Any $\mathbf{z} \subseteq \mathbf{x}^{(o)}$ satisfies the ignorability condition if $y_0, y_1 \perp_p t | \mathbf{z}$.

In the above definition, the potential outcomes y_0 and y_1 , the observed treatment t , and the features \mathbf{z} are all deterministic functions of the units $\mathbf{x}^{(u)}$. Therefore, the conditional independence criterion makes sense over the common probability measurable in the space of the units $\mathbf{x}^{(u)}$. As mentioned in Section 1, ignorability cannot be tested for from observational data since for every observational sample either y_0 or y_1 is observed (and not both).

In the PO framework, the ATE is defined as $\mathbb{E}_{\mathbf{x}^{(u)}}[y_1 - y_0]$. When $\mathbf{z} \subseteq \mathbf{x}^{(o)}$ is ignorable, it is also a valid adjustment relative to (t, y) in \mathcal{G} and therefore the ATE can be estimated by regressing on \mathbf{z} .

The Pearlian framework provides a generative model for this setup i.e., a semi-Markovian model (specifying a DAG that encodes causal assumptions relating exogenous and observed variables) as well as specifies graphical criteria that imply existence of valid adjustments relative to (t, y) in \mathcal{G} .

D D-SEPARATION

In this section, we define d-separation with respect to a semi-Markovian DAG \mathcal{G} . The d-separation or directed-separation is a commonly used graph separation criterion that characterizes conditional independencies in DAGs. First, we will define the notion of a path.

For any positive integer k , let $[k] := \{1, \dots, k\}$.

Definition 8. (*Path*) A path $\mathcal{P}(v_1, v_k)$ is an ordered sequence of distinct nodes $v_1 \dots v_k$ and the edges between these nodes such that for any $i \in [k]$, $v_i \in \mathcal{V}$ and for any $i \in [k-1]$, either $v_i \rightarrow v_{i+1}$, $v_i \leftarrow v_{i+1}$ or $v_i \leftrightarrow v_{i+1}$.

For example, in Figure 2(a), $\mathcal{P}(x_1, x_3) = \{x_1 \rightarrow x_2 \leftrightarrow x_3\}$ and $\mathcal{P}(t, y) = \{t \leftrightarrow x_1 \rightarrow x_2 \leftrightarrow x_3 \leftrightarrow y\}$ are two distinct paths. Next, we will define the notion of a collider.

Definition 9. (*Collider*) In a path $\mathcal{P}(v_1, v_k)$, for any $i \in \{2, \dots, k-1\}$, a collider at v_i mean that the arrows (or edges) meet head-to-head (collide) at v_i i.e. either $v_{i-1} \rightarrow v_i \leftarrow v_{i+1}$, $v_{i-1} \leftrightarrow v_i \leftarrow v_{i+1}$, $v_{i-1} \rightarrow v_i \leftrightarrow v_{i+1}$ or $v_{i-1} \leftrightarrow v_i \leftrightarrow v_{i+1}$.

For example, in $\mathcal{P}(x_1, x_3)$ defined above, there is a collider at x_2 . Next, we define the notion of a descendant path.

Definition 10. (*Descendant path*) A path $\mathcal{P}(v_1, v_k)$ is said to be an descendant path from v_1 to v_k if $\forall i \in [k-1]$, $v_i \rightarrow v_{i+1}$.

For example, in Figure 2(a), $\mathcal{P}(x_1, y) = \{x_1 \rightarrow x_2 \rightarrow y\}$ is an descendant path from x_1 to y . Next, we define the notion of a descendant.

Definition 11. (*Descendant*) A variable v_k is a descendant of a variable v_1 if there exists an descendant path $\mathcal{P}(v_1, v_k)$ from v_1 to v_k .

For example, in Figure 2(a), y is a descendant of x_1 .

Definition 12. (*Blocking path*) For any variables $v_1, v_2 \in \mathcal{W}$, a set $\mathbf{v} \subseteq \mathcal{W}$, and a path $\mathcal{P}(v_1, v_2)$, \mathbf{v} blocks the path $\mathcal{P}(v_1, v_2)$ if there exists a variable v in the path $\mathcal{P}(v_1, v_2)$ that satisfies either of the following two conditions:

- (1) $v \in \mathbf{v}$ and v is not a collider.

(2) neither variable v nor any of its descendant is in \mathbf{v} ; and v is a collider.

For example, in Figure 2(a), $\{x_2\}$ blocks the path $\mathcal{P}(x_1, y) = \{x_1 \rightarrow x_2 \rightarrow y\}$ because $x_2 \in \{x_2\}$ and x_2 is not a collider. Further, $\{x_2\}$ also blocks the path $\mathcal{P}(x_1, y) = \{x_1 \rightarrow x_2 \leftrightarrow x_3 \leftrightarrow y\}$ because $x_3 \notin \{x_2\}$ and x_3 is a collider..

Definition 13. (*D-separation*) For any variables $v_1, v_2 \in \mathcal{W}$, and a set $\mathbf{v} \subseteq \mathcal{W}$, v_1 and v_2 are d-separated by \mathbf{v} in \mathcal{G} if \mathbf{v} blocks every path between v_1 and v_2 in \mathcal{G} .

For example, in Figure 2(a), x_1 and y are d-separated by $\{x_2\}$.

E ADDITIONAL NOTATIONS

In this section, we will look at a few additional notations that will be used in the proofs of Theorem 3.1, Theorem 3.2, and Corollary 2.

E.1 \mathcal{G}_{-t}

Often it is favorable to think of the back-door criterion in terms of the graph obtained by removing the edge from t to y in \mathcal{G} . Let \mathcal{G}_{-t} denote this graph. The following (well-known) remark connects the back-door criterion to \mathcal{G}_{-t} .

Remark 5. Under Assumption 1, a set of variables $\mathbf{z} \subseteq \mathbf{x}$ satisfies the back-door criterion relative to the ordered pair of variables (t, y) in \mathcal{G} if and only if t and y are d-separated by \mathbf{z} in \mathcal{G}_{-t} .

Proof. Under Assumption 1, y is the only descendant of t i.e., no node in \mathbf{x} is a descendant of t . Therefore, from Definition 5, \mathbf{z} satisfying the back-door criterion relative to (t, y) in \mathcal{G} is equivalent to \mathbf{z} blocking every path between t and y in \mathcal{G} that contains an arrow into t . Further, under Assumption 1, there are no paths between t and y in \mathcal{G} that contain an arrow out of t apart from the direct path $t \rightarrow y$. However, this direct path $t \rightarrow y$ does not exist in \mathcal{G}_{-t} . Therefore, \mathbf{z} blocking every path between t and y in \mathcal{G} that contains an arrow into t is equivalent to \mathbf{z} blocking every path between t and y in \mathcal{G}_{-t} . Thus, \mathbf{z} satisfying the back-door criterion relative to (t, y) in \mathcal{G} is equivalent to \mathbf{z} blocking every path between t and y in \mathcal{G}_{-t} i.e., $t \perp\!\!\!\perp_d y | \mathbf{z}$ in \mathcal{G}_{-t} . \square

E.2 Subset of a path

Now, we will define the notion of a subset of a path.

Definition 14. (*Subset of a path*) A path $\mathcal{P}'(y_1, y_j)$ is said to be a subset of the path $\mathcal{P}(x_1, x_k)$ (denoted by $\mathcal{P}'(y_1, y_j) \subset \mathcal{P}(x_1, x_k)$) if $j < k$, $\exists i \in [k + 1 - j]$ such that $x_i = y_1$, $x_{i+1} = y_2, \dots, x_{i+j-1} = y_j$ and the edge between x_{i+l-1} and x_{i+l} is same as the edge between y_l and $y_{l+1} \forall l \in [j - 1]$.

For example, in Figure 2(a) $\mathcal{P}(x_1, x_3) = \{x_1 \rightarrow x_2 \leftrightarrow x_3\}$ is a subset of the path $\mathcal{P}(t, y) = \{t \leftrightarrow x_1 \rightarrow x_2 \leftrightarrow x_3 \leftrightarrow y\}$ i.e., $\mathcal{P}(x_1, x_3) \subset \mathcal{P}(t, y)$. For a path $\mathcal{P}(x_1, x_k)$, it is often convenient to represent the subset obtained by removing the nodes at each extreme and the corresponding edges by $\mathcal{P}(x_1, x_k) \setminus \{x_1, x_k\}$. For example, $\mathcal{P}(x_1, x_3) \setminus \{x_1, x_3\} = \mathcal{P}(t, y) \setminus \{t, y\}$.

F PROOF OF THEOREM 3.1 AND AN ILLUSTRATIVE EXAMPLE

In this section, we will prove Theorem 3.1 and also provide an illustrative example for Theorem 3.1. Recall the notions of *path*, *collider*, *descendant path*, *blocking path* and *d-separation* from Appendix D. Also, recall the notions of *subset of a path* and \mathcal{G}_{-t} as well as Remark 5 from Appendix E.

F.1 Proof of Theorem 3.1

We re-state the Theorem below and then provide the proof.²

²We say that \mathbf{z} satisfies the backdoor criterion if it blocks all the backdoor paths between t and y in \mathcal{G} i.e., paths between t and y in \mathcal{G} that contains an arrow into t . Please see Definition 5 in the main paper.

Theorem 3.1. *Let Assumption 1 be satisfied. Consider any $x_t \in \mathbf{x}^{(o)}$ that has a direct edge or a bi-directed edge to t i.e., either $x_t \rightarrow t$, $x_t \leftrightarrow t$ or $x_t \leftarrow t$. Let e be sub-sampled using x_t and \mathbf{v} for any $\mathbf{v} \subseteq \mathcal{V} \setminus \{x_t, y\}$ i.e., $e = f(x_t, \mathbf{v}, \eta)$. Let $\mathbf{z} \subseteq \mathbf{x}^{(o)} \setminus \{x_t\}$. If e is d-separated from y by \mathbf{z} and t in \mathcal{G} i.e., $e \perp_d y | \mathbf{z}, t$ in \mathcal{G} , then \mathbf{z} satisfies the back-door criterion relative to (t, y) in \mathcal{G} .*

Proof. We will prove this by contradiction. Suppose \mathbf{z} does not satisfy the back-door criterion relative to (t, y) in \mathcal{G} . From Remark 5, under Assumption 1, this is equivalent to t and y not being d-separated by \mathbf{z} in \mathcal{G}_{-t} . This is further equivalent to saying that there exists at least one unblocked path (not containing the edge $t \rightarrow y$) from t to y in \mathcal{G} when \mathbf{z} is conditioned on. Let $\mathcal{P}(t, y)$ denote the shortest of these unblocked paths. We have the following two scenarios depending on whether or not $\mathcal{P}(t, y)$ contains x_t . First, we will show that in both of these cases there exists an unblocked path³ $\mathcal{P}'(x_t, y)$ from x_t to y in \mathcal{G} when \mathbf{z}, t are conditioned on.

Note : All bi-directed edges in \mathcal{G} are unblocked because (a) none of the unobserved feature is conditioned on and (b) there is no collider at any of the unobserved feature.

- (i) $x_t \in \mathcal{P}(t, y)$: This implies that there is an unblocked path $\mathcal{P}''(x_t, y) \subset \mathcal{P}(t, y)$ from x_t to y in \mathcal{G} when \mathbf{z} is conditioned on. Suppose we now condition on t in addition to \mathbf{z} . The conditioning on t can affect the path $\mathcal{P}''(x_t, y)$ only⁴ if a) there is an unblocked descendant path from some $x_s \in \mathcal{P}''(x_t, y) \setminus \{x_t, y\}$ to t and b) x_s is a collider in the path $\mathcal{P}''(x_t, y) \setminus \{x_t, y\}$. However, conditioning on such a t cannot block the path $\mathcal{P}''(x_t, y)$. Thus, there exists an unblocked path $\mathcal{P}'(x_t, y) = \mathcal{P}''(x_t, y)$ in \mathcal{G} when \mathbf{z}, t are conditioned on.
- (ii) $x_t \notin \mathcal{P}(t, y)$: Under Assumption 1, \mathcal{G} cannot contain the edge $t \leftarrow y$ (because a DAG cannot have a cycle). Furthermore, under Assumption 1, t has no child other than y . Therefore, in this case, the path $\mathcal{P}(t, y)$ takes one of the following two forms : (a) $t \leftarrow x_s \cdots y$ or (b) $t \leftrightarrow x_s \cdots y$ for some $x_s \neq x_t$. In either case, there is a collider at t (i.e., either $x_t \rightarrow t \leftarrow x_s$, $x_t \rightarrow t \leftrightarrow x_s$, $x_t \leftrightarrow t \leftarrow x_s$ or $x_t \leftrightarrow t \leftrightarrow x_s$) in the path $\mathcal{P}'''(x_t, x_s)$ from x_t to x_s . Suppose we now condition on t in addition to \mathbf{z} . The conditioning on t unblocks the path $\mathcal{P}'''(x_t, x_s)$ because there is a collider at t . Also, similar to the previous case, the conditioning on t cannot block the path $\mathcal{P}(t, y)$ from t to y (passing through x_s). Therefore, we see that there is an unblocked path $\mathcal{P}'(x_t, y)$ from x_t to y (passing through t and x_s) in \mathcal{G} when \mathbf{z}, t are conditioned on (i.e., either $x_t \rightarrow t \leftarrow x_s \cdots y$, $x_t \rightarrow t \leftrightarrow x_s \cdots y$, $x_t \leftrightarrow t \leftarrow x_s \cdots y$ or $x_t \leftrightarrow t \leftrightarrow x_s \cdots y$).

Now, in each of the above cases, there is an edge from x_t to e because e is sub-sampled using x_t . Therefore, there exists an unblocked path $\mathcal{P}''''(e, y) \supset \mathcal{P}'(x_t, y)$ of the form $e \leftarrow x_t \cdots y$ in \mathcal{G} when \mathbf{z}, t are conditioned on because $x_t \notin \mathbf{z}$ i.e., x_t is not conditioned on. This is true regardless of whether x_t is an ancestor of \mathbf{z} or not since the edge $e \leftarrow x_t$ cannot create a collider at x_t . The existence of the path $\mathcal{P}''''(e, y)$ contradicts the fact that e is d-separated from y by \mathbf{z} and t in \mathcal{G} . This completes the proof. \square

F.2 An illustrative example for Theorem 3.1

Now, we will look into an example illustrating Theorem 3.1. Consider the DAG \mathcal{G}^{bi} in Figure 6. We let $x_t = x_1$ (because $x_1 \leftrightarrow t$) and sub-sample e using x_1 and t (see Figure 6). For this example, $\mathbf{z} \subseteq \{x_2, x_3\}$ i.e., $\mathbf{z} \in \{\emptyset, \{x_2\}, \{x_3\}, \{x_2, x_3\}\}$. It is easy to verify that $e \perp_d y | x_2, t$ but $e \not\perp_d y | t$, $e \not\perp_d y | x_3, t$, and $e \not\perp_d y | x_2, x_3, t$ in \mathcal{G}^{bi} . Given these, Theorem 3.1 implies that $\mathbf{z} = \{x_2\}$ should satisfy the back-door criterion relative to (t, y) in \mathcal{G}^{bi} . This is indeed the case and can be verified easily. Thus, we see that our framework has the potential to identify valid adjustment sets ($\{x_2\}$ for \mathcal{G}^{bi}) in the scenario where no causal parent of the treatment variable is known but a bi-directed neighbor of the treatment is known.

Note : Theorem 3.1 does not comment on whether $\emptyset, \{x_3\}$ and $\{x_2, x_3\}$ satisfy or do not satisfy the back-door criterion relative to (t, y) in \mathcal{G}^{bi} .

³Note: There is no possibility of an unblocked path from x_t to y in \mathcal{G} containing the edge $t \rightarrow y$ when \mathbf{z}, t are conditioned on. This is because t is conditioned on and any such path to y cannot form a collider at t .

⁴ $t \notin \mathcal{P}''(x_t, y)$ because $\mathcal{P}(t, y)$ is the shortest unblocked path (not containing the edge $t \rightarrow y$) from t to y in \mathcal{G} when \mathbf{z} is conditioned on.




Figure 6: The DAG \mathcal{G}^{bi} where e has been sub-sampled using x_1 and t .

G PROOF OF THEOREM 3.2

In this section, we will prove Theorem 3.2. Recall the notions of *path*, *collider*, *descendant path*, *descendant*, *blocking path* and *d-separation* from Appendix D. Also, recall the notions of *subset of a path* and \mathcal{G}_{-t} as well as Remark 5 from Appendix E.

We re-state the Theorem below and then provide the proof.⁵

Theorem 3.2. *Let Assumption 1 be satisfied. Consider any $x_t \in \mathbf{x}^{(o)}$ that has a direct edge to t i.e., $x_t \rightarrow t$ or $x_t \leftarrow t$. Let e be sub-sampled using x_t and \mathbf{v} for any $\mathbf{v} \subseteq \{t\}$ i.e., $e = f(x_t, \mathbf{v}, \eta)$. Let $\mathbf{z} \subseteq \mathbf{x}^{(o)} \setminus \{x_t\}$. If \mathbf{z} satisfies the back-door criterion relative to (t, y) in \mathcal{G} , then e is d-separated from y by \mathbf{z} and t in \mathcal{G} i.e., $e \perp\!\!\!\perp_d y | \mathbf{z}, t$ in \mathcal{G} .*

Proof. We will prove this by contradiction. Suppose $e \not\perp\!\!\!\perp_d y | \mathbf{z}, t$ in \mathcal{G} i.e., e and y are not d-separated by \mathbf{z}, t in \mathcal{G} . In other words, there exists at least one unblocked path from e to y in \mathcal{G} when \mathbf{z}, t are conditioned on. Let $\mathcal{P}(e, y)$ denote the shortest of these unblocked paths.

Depending on the choice of \mathbf{v} , we have the following two cases. In each of these cases, we will show that the path $\mathcal{P}(e, y)$ is of the form $e \leftarrow x_t \cdots y$.

- $\mathbf{v} = \{t\}$: e is sub-sampled using t and x_t . Therefore, the path $\mathcal{P}(e, y)$ can take one of the following two forms : (a) $e \leftarrow t \cdots y$ or (b) $e \leftarrow x_t \cdots y$. However, t is conditioned on and the path $e \leftarrow t \cdots y$ cannot form a collider at t (because of the edge $e \leftarrow t$). Therefore, the path $\mathcal{P}(e, y)$ cannot be of the form $e \leftarrow t \cdots y$ and has to be of the form $e \leftarrow x_t \cdots y$.
- $\mathbf{v} = \emptyset$: e is sub-sampled using only x_t . Therefore, the path $\mathcal{P}(e, y)$ has to be of the form $e \leftarrow x_t \cdots y$.

Now, observe that there is no collider at x_t in the path $e \leftarrow x_t \cdots y$ and x_t is not conditioned on (because $x_t \notin \mathbf{z}$). Therefore, there exists at least one unblocked path from x_t to y in \mathcal{G} when \mathbf{z}, t are conditioned on. Let $\mathcal{P}'(x_t, y) \subset \mathcal{P}(e, y)$ denote the shortest of these unblocked paths from x_t to y in \mathcal{G} when \mathbf{z}, t are conditioned on. The path $\mathcal{P}'(x_t, y)$ cannot contain the edge $t \rightarrow y$ since t is conditioned on and the path cannot form a collider at t (because of the edge $t \rightarrow y$).

We have the following two scenarios depending on whether or not $\mathcal{P}'(x_t, y)$ contains t . First, we will show that in both of these cases there exists an unblocked path $\mathcal{P}''(t, y)$ from t to y (that does not contain the edge $t \rightarrow y$) in \mathcal{G} when \mathbf{z} is conditioned on.

Note : All bi-directed edges in \mathcal{G} are unblocked because (a) none of the unobserved feature is conditioned on and (b) there is no collider at any of the unobserved feature.

- (1) $t \notin \mathcal{P}'(x_t, y)$: Suppose we now uncondition on t (but still condition on \mathbf{z}). We have the following two scenarios depending on whether or not unconditioning on t blocks the path $\mathcal{P}'(x_t, y)$ (while \mathbf{z} is still conditioned on).
 - (i) Unconditioning on t does not block the path $\mathcal{P}'(x_t, y)$: Consider the path $\mathcal{P}''(t, y) \supset \mathcal{P}'(x_t, y)$ from t to y of the form $t \leftarrow x_t \cdots y$. This path is unblocked in \mathcal{G} when \mathbf{z} is conditioned on because (a) by assumption the

⁵We say that \mathbf{z} satisfies the backdoor criterion if it blocks all the backdoor paths between t and y in \mathcal{G} i.e., paths between t and y in \mathcal{G} that contains an arrow into t . Please see Definition 5 in the main paper.

path $\mathcal{P}'(x_t, y)$ is unblocked in \mathcal{G} when \mathbf{z} is conditioned on and (b) there is no collider at x_t in this path (in addition to x_t not being conditioned on since $x_t \notin \mathbf{z}$). $\mathcal{P}''(t, y)$ does not contain the edge $t \rightarrow y$ because $\mathcal{P}'(x_t, y)$ does not contain the edge $t \rightarrow y$.

- (ii) Unconditioning on t blocks the path $\mathcal{P}'(x_t, y)$ (Refer Figure 7 for an illustration of this case): We will first create a set \mathbf{x}_S consisting of all the nodes at which the path $\mathcal{P}'(x_t, y)$ is blocked when t is unconditioned on (while \mathbf{z} is still conditioned on). Define the set $\mathbf{x}_S \subseteq \mathbf{x}^{(o)}$ such that for any $x_s \in \mathbf{x}_S$ the following are true: (a) $x_s \in \mathcal{P}'(x_t, y) \setminus \{x_t, y\}$, (b) the path $\mathcal{P}'(x_t, y)$ contains a collider at x_s , (c) there is a descendant path $\mathcal{P}^d(x_s, t)$ from x_s to t , (d) the descendant path $\mathcal{P}^d(x_s, t)$ is unblocked when \mathbf{z} is conditioned on, (e) $x_s \notin \mathbf{z}$, and (f) there is no unblocked descendant path from x_s to any $x_a \in \mathbf{z}$.

Since the path $\mathcal{P}'(x_t, y)$ is blocked when t is unconditioned on (while \mathbf{z} is still conditioned on), we must have that $\mathbf{x}_S \neq \emptyset$. Let $x_c \in \mathbf{x}_S$ be that node which is closest to y in the path $\mathcal{P}'(x_t, y)$. By the definition of \mathbf{x}_S and the choice of x_c , unconditioning on t cannot block the path $\mathcal{P}'''(x_c, y) \subset \mathcal{P}'(x_t, y)$ when \mathbf{z} is still conditioned on. Also, by the definition of \mathbf{x}_S , the descendant path $\mathcal{P}^d(x_c, t)$ from x_c to t is unblocked when \mathbf{z} is conditioned on.

Now consider the path $\mathcal{P}''(t, y)$ of the form $t \leftarrow \dots \leftarrow x_c \leftarrow \dots \rightarrow y$ i.e., $\mathcal{P}''(t, y) \supset \mathcal{P}'''(x_c, y)$ and $\mathcal{P}''(t, y) \supset \mathcal{P}^d(x_c, t)$. The path $\mathcal{P}''(t, y)$ is unblocked when \mathbf{z} is conditioned on since (a) $\mathcal{P}^d(x_c, t)$ is unblocked when \mathbf{z} is conditioned on, (b) $\mathcal{P}'''(x_c, y)$ is unblocked when \mathbf{z} is conditioned on, and (c) there is no collider at x_c and x_c is not conditioned on since $x_c \notin \mathbf{z}$. Furthermore, $\mathcal{P}''(t, y)$ does not contain the edge $t \rightarrow y$ because $\mathcal{P}'''(x_c, y) \subset \mathcal{P}'(x_t, y)$ does not contain the edge $t \rightarrow y$ and $\mathcal{P}^d(x_c, t)$ does not contain the edge $t \rightarrow y$.

- (2) $t \in \mathcal{P}'(x_t, y)$: In this case, there is an unblocked path $\mathcal{P}''''(t, y) \subset \mathcal{P}'(x_t, y)$ from t to y when \mathbf{z}, t are conditioned on. There are two sub-cases depending on whether or not unconditioning on t can block the path $\mathcal{P}''''(t, y)$ (while \mathbf{z} is still conditioned on).
- (A) Unconditioning on t does not block the path $\mathcal{P}''''(t, y)$: In this case, by assumption, the path $\mathcal{P}''(t, y) = \mathcal{P}''''(t, y)$ in \mathcal{G} is unblocked when \mathbf{z} is conditioned on. Furthermore, since $\mathcal{P}''(t, y) \subset \mathcal{P}'(x_t, y)$, $\mathcal{P}''(t, y)$ does not contain the edge $t \rightarrow y$.
- (B) Unconditioning on t blocks the path $\mathcal{P}''''(t, y)$: Let $x_{t'}$ be the node adjacent to t in the path $\mathcal{P}''''(t, y)$. Consider the path $\mathcal{P}''''(x_{t'}, y) \subset \mathcal{P}''''(t, y)$. Clearly, $t \notin \mathcal{P}''''(x_{t'}, y)$ since the path $\mathcal{P}'(x_t, y)$ was assumed to be the shortest unblocked path from x_t to y . Therefore, the only way unconditioning on t could block the path $\mathcal{P}''''(t, y)$ is if it blocked the path $\mathcal{P}''''(x_{t'}, y)$. Now, this sub-case is similar to the case (1)(ii) with $x_t = x_{t'}$ and $\mathcal{P}'(x_t, y) = \mathcal{P}''''(x_{t'}, y)$ ⁶. As in (1)(ii), it can be shown that there exists an unblocked path $\mathcal{P}''(t, y)$ in \mathcal{G} (that does not contain the edge $t \rightarrow y$) when \mathbf{z} is conditioned on.




Figure 7: Illustrating the case (1)(ii) in the proof of Theorem 3.2

Now, in each of the above cases, there exists an unblocked path $\mathcal{P}''(t, y)$ in \mathcal{G} when \mathbf{z} is conditioned on and this path does not contain the edge $t \rightarrow y$. Therefore, there exists an unblocked path $\mathcal{P}''(t, y)$ in \mathcal{G}_{-t} when \mathbf{z} is conditioned on (since $\mathcal{P}''(t, y)$ does not contain the edge $t \rightarrow y$) implying $t \not\perp_d y | \mathbf{z}$ in \mathcal{G}_{-t} . From Remark 5, under Assumption 1, this is equivalent to \mathbf{z} not satisfying the back-door criterion relative to (t, y) in \mathcal{G} leading to a contradiction. This completes the proof. \square

⁶The choice of edge (\rightarrow or \leftarrow) between x_t and t does not matter in (1)(ii).

H THE M-BIAS MODEL

In this section, we discuss the M-bias problem. It is a causal model under which although some observed features (that are pre-treatment) are provided, one must not adjust for any of it. This model has been widely discussed (Imbens, 2020; Liu et al., 2012) in the literature to underscore the need for algorithms that find valid adjustment sets.

We illustrate the M-bias problem using the semi-Markov model (with the corresponding DAG \mathcal{G}^M) in Figure 8.




Figure 8: The DAG \mathcal{G}^M illustrating the M-bias problem.

The DAG \mathcal{G}^M consists of the following edges: $t \rightarrow y, x_1 \leftrightarrow t, x_1 \leftrightarrow y$. It is easy to verify that $\{x_1\}$ does not satisfy the back-door criterion with respect to (t, y) in \mathcal{G}^M . Further, it is also easy to verify that the empty set i.e., \emptyset satisfies the back-door criterion with respect to (t, y) in \mathcal{G}^M . In what follows, we will see how our framework cannot be used to arrive at this conclusion.

There are no observed parents of t in \mathcal{G}^M . Therefore, Theorem 3.2 (i.e., the necessary condition) does not apply here. For Theorem 3.1 to be applicable, there is only one choice of x_t i.e., one must use $x_t = x_1$. Now, for any $\mathbf{v} \subseteq \{t\}$ such that e is sub-sampled according to $e = f(x_1, \mathbf{v}, \eta)$, e is not d-separated from y given only t . Therefore, one cannot conclude whether or not $\mathbf{z} = \emptyset$ satisfies the back-door criterion with respect to (t, y) in \mathcal{G}^M from Theorem 3.1 (i.e., the sufficiency condition). In summary, we see that our sufficient condition cannot identify the set satisfying the back-door criterion (i.e., the null set) and necessity condition does not apply in the case of the M-bias problem.

Therefore, there are models where sets satisfying the back-door criterion exist (for e.g., the empty set in the M-bias problem) and our results may not be able to identify them.

I FINDING ALL BACK-DOORS

Building on Corollary 1, we provide an Algorithm (Algorithm 3) that, when all the parents of the treatment are observed and known, finds the set of all the subsets of the observed features satisfying the back-door criterion relative to (t, y) in \mathcal{G} which we denote by \mathcal{Z} . We initialize Algorithm 3 with the set \mathcal{Z}_1 obtained by adding $\pi(t)$ to every element of the power set of $\mathbf{x}^{(o)} \setminus \pi(t)$. The set \mathcal{Z}_1 can be constructed easily with the knowledge of $\mathbf{x}^{(o)}$ and $\pi(t)$ provided to Algorithm 3. Then, we repeatedly apply Corollary 1 to each parent in turn to identify all back-doors. We state this result formally in Corollary 2 below.

Algorithm 3: Finding all back-doors

Input: $\pi(t), e, t, y, \mathbf{x}^{(o)}$

Output: \mathcal{Z}

Initialization: $\mathcal{Z} = \mathcal{Z}_1$

```

1 for  $x_t \in \pi(t)$  do
2   for  $\mathbf{z} \subseteq \mathbf{x}^{(o)} \setminus \{x_t\}$  do
3     if  $e \perp_p y | \mathbf{z}, t$  then
4        $\mathcal{Z} = \mathcal{Z} \cup \mathbf{z}$ 
    
```

Remark: Algorithm 3 is based on two key ideas : (1) Any subset of the observed features that contains all the parents of the treatment satisfies the back-door criterion relative to (t, y) in \mathcal{G} . Formally, consider the set \mathcal{Z}_1

obtained by adding $\pi(t)$ to every element of the power set of $\mathbf{x}^{(o)} \setminus \pi(t)$. Then, any $\mathbf{z} \in \mathcal{Z}_1$ satisfies the back-door criterion relative to (t, y) in \mathcal{G} . We use the set \mathcal{Z}_1 in the initialization step of Algorithm 3 as it can be constructed easily with the knowledge of $\mathbf{x}^{(o)}$ and $\pi(t)$. (2) For any $\mathbf{z} \notin \mathcal{Z}_1$ that satisfies the back-door criterion relative to (t, y) in \mathcal{G} , there exists $x_t \in \pi(t)$ such that $\mathbf{z} \subseteq \mathbf{x}^{(o)} \setminus x_t$. In this scenario, Algorithm 3 captures \mathbf{z} because $e \perp_p y | \mathbf{z}, t$ from Corollary 1 (under Assumption 2).

We now provide an example illustrating Algorithm 3, followed by Corollary 2 and its proof.

I.1 Example

We illustrate Algorithm 3 with an example. Consider the DAG \mathcal{G}_{bd} in Figure 9. It is easy to verify that, for \mathcal{G}_{bd} , $\mathcal{Z} = \{\{x_3\}, \{x_1, x_3\}, \{x_2, x_3\}, \{x_1, x_2\}, \{x_1, x_2, x_3\}, \{x_1, x_2, x_4\}, \{x_1, x_2, x_3, x_4\}\}$. Now, Algorithm 3 takes $\pi(t) = \{x_1, x_2\}$ and $\mathbf{x}^{(o)} = \{x_1, x_2, x_3, x_4\}$ as inputs. Therefore, $\mathcal{Z}_1 = \{x_1, x_2, \{x_1, x_2, x_3\}, \{x_1, x_2, x_4\}, \{x_1, x_2, x_3, x_4\}\}$ can be constructed by adding $\pi(t)$ to every element of the power set of $\mathbf{x}^{(o)} \setminus \pi(t)$ i.e., to the power set of $\{x_3, x_4\}$. Algorithm 3 is initialized with \mathcal{Z}_1 and the only remaining sets to be identified are $\{x_3\}, \{x_1, x_3\}$, and $\{x_2, x_3\}$. When $x_t = x_1$, Algorithm 3 will identify $\{x_3\}$ and $\{x_2, x_3\}$ as sets that satisfy the back-door criterion relative to (t, y) in \mathcal{G}_{bd} . Similarly, when $x_t = x_2$, Algorithm 3 will identify $\{x_3\}$ and $\{x_1, x_3\}$ as sets that satisfy the back-door criterion relative to (t, y) in \mathcal{G}_{bd} .




Figure 9: The DAG \mathcal{G}_{bd} for illustrating Algorithm 3

I.2 Corollary 2

Recall the notions of *path*, *collider*, *descendant path*, *descendant*, *blocking path* and *d-separation* from Appendix D. Also, recall the notions of *subset of a path* and \mathcal{G}_{-t} as well as Remark 5 from Appendix E.

Corollary 2. *Let Assumptions 1 and 2 be satisfied. Let \mathcal{Z} be the set of all sets $\mathbf{z} \subseteq \mathbf{x}^{(o)}$ that satisfy the back-door criterion relative to the ordered pair of variables (t, y) in \mathcal{G} . If all the parents of t are observed and known i.e., $\pi(t) = \pi^{(o)}(t)$ is known, then Algorithm 3 returns the set \mathcal{Z} .*

Proof. From Remark 5, under assumption 1, \mathbf{z} satisfying the back-door criterion relative to the ordered pair of variables (t, y) in \mathcal{G} is equivalent to t and y being d-separated by \mathbf{z} in \mathcal{G}_{-t} i.e., $t \perp_d y | \mathbf{z}$ in \mathcal{G}_{-t} . From Pearl et al. (2016), $\pi(t)$ always satisfies the back-door criterion relative to the ordered pair of variables (t, y) in \mathcal{G} i.e., $t \perp_d y | \pi(t)$ in \mathcal{G}_{-t} . Consider any $\mathbf{z} \subseteq \mathbf{x}^{(o)}$ such that $\pi(t) \subseteq \mathbf{z}$. First, we will show that $t \perp_d y | \mathbf{z}$ in \mathcal{G}_{-t} i.e., \mathbf{z} satisfies the back-door criterion relative to the ordered pair of variables (t, y) in \mathcal{G} .

Suppose $t \not\perp_d y | \mathbf{z}$ in \mathcal{G}_{-t} i.e., t and y are not d-separated by \mathbf{z} in \mathcal{G}_{-t} . This is equivalent to saying that there exists at least one unblocked path (not containing the edge $t \rightarrow y$) from t to y in \mathcal{G}_{-t} when \mathbf{z} is conditioned on. Without the loss of generality, let $\mathcal{P}(t, y)$ denote any one of these unblocked paths. The path $\mathcal{P}(t, y)$ has to be of the form $t \leftarrow x_t \cdots y$ where $x_t \in \pi(t)$ because (a) under Assumption 1, \mathcal{G} cannot contain the edge $t \leftarrow y$ (because a DAG cannot have a cycle) and (b) under Assumption 1, t has no child other than y . However, $x_t \in \pi(t) \subseteq \mathbf{z}$ i.e., x_t is conditioned on. Now since there is no collider at x_t in the path $\mathcal{P}(t, y)$, it cannot be unblocked and this leads to a contradiction. Therefore, \mathbf{z} satisfies the back-door criterion relative to the ordered pair of variables (t, y) in \mathcal{G} .

Now, consider the set \mathcal{Z}_1 obtained by adding $\pi(t)$ to every element of the power set of $\mathbf{x}^{(o)} \setminus \pi(t)$ i.e., $\mathcal{Z}_1 := \{\mathbf{z} \subseteq \mathbf{x}^{(o)} : \pi(t) \subseteq \mathbf{z}\}$. From the argument above, we have $\mathcal{Z}_1 \subseteq \mathcal{Z}$. From the knowledge of $\pi(t)$ and $\mathbf{x}^{(o)}$, one can easily construct the set \mathcal{Z}_1 and thus initialize \mathcal{Z} in Algorithm 3 with \mathcal{Z}_1 .

Now, consider the set $\mathcal{Z}_2 := \mathcal{Z} \setminus \mathcal{Z}_1$. Consider any set $\mathbf{z} \in \mathcal{Z}_2$ satisfying the back-door criterion relative to the ordered pair of variables (t, y) in \mathcal{G} . By the definition of \mathcal{Z}_1 (and \mathcal{Z}_2), there exists at least one parent of t not present in the set \mathbf{z} . In other words, there exists $x_t \in \pi(t)$ such that $\mathbf{z} \subseteq \mathbf{x}^{(o)} \setminus x_t$. From Corollary 1, under Assumption 2, this is equivalent to $e \perp_p y | \mathbf{z}, t$. Therefore, Algorithm 3 will capture the set \mathbf{z} . Since the choice of \mathbf{z} was random, Algorithm 3 will capture every $\mathbf{z} \in \mathcal{Z}_2$ and return $\mathcal{Z}_1 \cup \mathcal{Z}_2$. This completes the proof. □

J THE BASELINE

In this section, we provide an implementation of the **Baseline** considered in Section 5. This routine estimates the ATE from the observational data by regressing y for the treated and the untreated sub-populations on a given set \mathbf{z} . The **Baseline** we consider in this work is an instance of this routine. More specifically, for the **Baseline**, we set \mathbf{z} to be the set of all the observed features i.e., $\mathbf{z} = \mathbf{x}^{(o)}$. See Section 5 for details.

Algorithm 4: ATE estimation using \mathbf{z} as an adjustment set

Input: n, n_r, t, y, \mathbf{z}

Output: $\text{ATE}(\mathbf{z})$

Initialization: $\text{ATE}(\mathbf{z}) = 0$

- 1 **for** $r = 1, \dots, n_r$ **do** // Use a different train-test split in each run
 - 2 $\text{ATE}(\mathbf{z}) = \text{ATE}(\mathbf{z}) + \frac{1}{n} \sum_{i=1}^n (\mathbb{E}[y | \mathbf{z} = \mathbf{z}^{(i)}, t = 1] - \mathbb{E}[y | \mathbf{z} = \mathbf{z}^{(i)}, t = 0]);$
 - 3 $\text{ATE}(\mathbf{z}) = \text{ATE}(\mathbf{z}) / n_r;$
-

K ADDITIONAL EXPERIMENTS

In this section, we briefly discuss the usage of real-world CI testers in Algorithm 1. We also provide in-depth discussions on the synthetic experiment from Section 5.1, the experiments on IHDP from Section 5.2, and the experiments on Cattaneo from Section 5.3. Additionally, we specify all the training details, as well as provide more details regarding the comparison of our method with Entner et al. (2013), Gultchin et al. (2020), and Cheng et al. (2020).

K.1 Usage of CI testers in Algorithm 1

In this work we use the RCot real-world CI tester (Strobl et al., 2019).

The real-world CI testers produce a p-value close to zero if the CI does not hold and produce a p-value uniformly distributed between 0 and 1 if the CI holds. Since we use a non-zero p-value threshold, depending on the quality of the CI tester, the false positive rate for valid adjustment sets may be non-zero.

Suppose, for a CI tester and for an increasing sample size n , we find a sequence of Type-I error rate (α_n) and Type-II error rate (β_n) going to zero i.e., $\alpha_n, \beta_n \rightarrow 0$. Then, if there is a valid adjustment set, it is easy to see that our algorithm will have zero bias in the estimated effect when the significance threshold α_n is used as the p-value threshold in our algorithm.

K.2 Synthetic experiment

In this sub-section, we provide more details on the synthetic experiment in Section 5.1.

Let $\text{Uniform}(a, b)$ denote the uniform distribution over the interval $[a, b]$ for $a, b \in \mathbb{R}$ such that $a < b$. Let $\mathcal{N}(\mu, \sigma^2)$ denote the Gaussian distribution with mean μ and variance σ^2 . Let $\text{Bernoulli}(p)$ denote the Bernoulli distribution which takes the value 1 with probability p . Let $\text{Sigmoid}(\cdot)$ denote the sigmoid function i.e., for any $a \in \mathbb{R}$, $\text{Sigmoid}(a) = 1 / (1 + e^{-a})$. Let $\text{Softmax}(\cdot)$ denote the softmax function.

Dataset Description. We generate different variables as below:

- $u_i \sim \text{Uniform}(1, 2)$

- $x_1 \sim \theta_{11}u_1 + \theta_{12}u_2 + \mathcal{N}(0, 0.01)$ where $\theta_{11}, \theta_{12} \in \text{Uniform}(1, 2)$
- $x_2 \sim \theta_{21}x_1 + \theta_{22}u_2 + \theta_{23}u_3 + \mathcal{N}(0, 0.01)$ where $\theta_{21}, \theta_{22}, \theta_{23} \in \text{Uniform}(1, 2)$
- $x_3 \sim \theta_{31}u_3 + \theta_{32}u_4 + \mathcal{N}(0, 0.01)$ where $\theta_{31}, \theta_{12} \in \text{Uniform}(1, 2)$
- $t \sim \text{Bernoulli}(\text{Sigmoid}(\theta_{51}x_1 + \theta_{52}u_1))$ where $\theta_{51}, \theta_{52} \in \text{Uniform}(1, 2)$
- $y \sim \theta_{41}x_2 + \theta_{42}u_4 + \theta_{43}t + \mathcal{N}(0, 0.01)$ where $\theta_{41}, \theta_{42}, \theta_{43} \in \text{Uniform}(1, 2)$


We generate the weight vectors from $\text{Uniform}(1, 2)$ to ensure that the faithfulness assumption with respect to the sub-sampling variable is satisfied (i.e., Assumption 2). This is because for smaller weights, it is possible that conditionally dependent relations are declared as conditionally independent. See Uhler et al. (2013) for details.

For all our experiments, we use 3 environments i.e., $e \in \{0, 1, 2\}$ and generate the sub-sampling variable as below with $\hat{\mathbb{E}}$ denoting the empirical expectation. While other choices of sub-sampling function f could be explored, the natural choice (for discrete e) of softmax with random weights suffices.

- $e \sim \text{Softmax}(\theta_{61}(x_1 - \hat{\mathbb{E}}[x_1]) + \theta_{62}(t - \hat{\mathbb{E}}[t]))$ with $\theta_{61} := (\theta_{61}^{(1)}, \theta_{61}^{(2)}, \theta_{61}^{(3)}) \in \mathbb{R}^3$ and $\theta_{62} := (\theta_{62}^{(1)}, \theta_{62}^{(2)}, \theta_{62}^{(3)}) \in \mathbb{R}^3$ such that $\theta_{61}^{(1)}, \theta_{62}^{(1)} \in \text{Uniform}(1, 2)$, $\theta_{61}^{(2)} = \theta_{62}^{(2)} = 0$, and $\theta_{61}^{(3)}, \theta_{62}^{(3)} \in \text{Uniform}(-2, -1)$

In other words, we keep separation between the weight vectors associated with different environments to make sure that the environments look different from each other as expected by IRM.

Success Probability. For a given p_{value} threshold, we let the success probability of the set $\{x_2\}$ be the fraction of times (in n_r runs) the p-value of $\text{CI}(e \perp_p y | x_2, t)$ is more than p_{value} . In Figure 10a below, we show how the success probability of the set $\{x_2\}$ varies with different p_{value} thresholds i.e., $\{0.1, 0.2, 0.3, 0.4, 0.5\}$ for the dataset used in Section 5.1. As we can see in Figure 10a, the success probability of the set $\{x_2\}$, for the same p_{value} threshold, is much lower in high dimensions compared to low dimensions. We believe this happens (a) because of the non-ideal CI tester and (b) because the number of samples are finite. In contrast, our algorithms IRM-t and IRM-c always pick the set $\{x_2\}$ to adjust on i.e., $\mathbf{z}_{irm} = \{x_2\}$ for both IRM-t and IRM-c for $d = 3, 5, 7$.



(a) Success probability of the set $\{x_2\}$ in the toy example \mathcal{G}^{toy} for different p_{value} thresholds.

(b) Performance of Algorithm 1 on \mathcal{G}^{toy} when the candidate adjustment sets are \tilde{d} -dimensional.

Figure 10: Additional analysis on the toy example \mathcal{G}^{toy} .

Sparse subset search. In Section 5.1, we validated our algorithm by letting \mathcal{X} be the set of all subsets of $\mathbf{x}^{(o)} \setminus \{x_t\}$. However, for this synthetic experiment, we do know that only $x_2 \in \mathcal{X}$ satisfies the back-door criterion relative to (t, y) . Further, we know that x_2 is \tilde{d} -dimensional. Therefore, with this additional knowledge, we could instead let \mathcal{X} be the set of all \tilde{d} -dimensional subsets of $\mathbf{x}^{(o)} \setminus \{x_t\}$. In other words, we consider the Sparse algorithm from Section 5 with $k = \tilde{d}$ ⁷. We show the performance of this algorithm for this choice of \mathcal{X} , in comparison to the Baseline (i.e. using all observed features) as well as IRM-t and IRM-c, in Figure 10b for $d = 3, 5, 7$. With this restriction on the candidate adjustment sets, our algorithm performs better than it does in Figure 3(b) where there are no restrictions on the candidate adjustment sets.

⁷More precisely, the Sparse algorithm considers subsets of size at-most k . Here, we consider subsets of size exactly equal to k .


Performance with dimensions. The gains of our testing and subset search based algorithm over the **Baseline** are much more in the low dimensions compared to the high dimensions as seen in Figures 3(b) and 10b. We believe there are two primary reasons behind this : (a) The CI tester leaks more false positive in high dimensions compared to low dimensions (see Appendix K.1) and (b) The CI tester fails to consistently output a high p-value for the set $\{x_2\}$ in high dimensions (see Figure 10a). The gains of our IRM based algorithm remain consistent even in high dimensions as expected.

K.3 Generating the environment/sub-sampling variable


In all our experiments in Section 5, we let the sub-sampling variable depend on both x_t and t . Now, we will look into the case where the sub-sampling variable is generated as a function of only $x_t = x_1$ i.e., $e = f(x_1)$. More specifically, we generate the sub-sampling variable as below:

- $e \sim \text{Softmax}(\theta_{61}(x_1 - \hat{\mathbb{E}}[x_1]))$ with $\theta_{61} := (\theta_{61}^{(1)}, \theta_{61}^{(2)}, \theta_{61}^{(3)}) \in \mathbb{R}^3$ such that $\theta_{61}^{(1)} \in \text{Uniform}(1, 2)$, $\theta_{61}^{(2)} = 0$, and $\theta_{61}^{(3)} \in \text{Uniform}(-2, -1)$


For this setting, we show the plots analogous to those in Figure 3(a), Figure 3(b), Figure 10a and Figure 10b in Figure 11. As we can see in Figure 11a, Figure 11b, Figure 11c, and Figure 11d, the performance of our algorithm with $e = f(x_t)$ is similar to (at a high level) its performance with $e = f(x_t, t)$. This should not be surprising since Corollary 1 holds for any $\mathbf{v} \subseteq \{t\}$ i.e., for both $\mathbf{v} = \emptyset$ and $\mathbf{v} = \{t\}$. In other words, while theoretical tradeoff between the choice of \mathbf{v} i.e., \emptyset or $\{t\}$ is unclear, there is no major empirical difference. Note: We do not show the performance of IRM based algorithms for $e = f(x_1)$ since it is exactly the same as the performance for $e = f(x_1, t)$.




(a) Sets not satisfying back-door ($\{x_1, x_2, x_3\}$, $\{x_2, x_3\}$) result in high ATE error; sets satisfying back-door ($\{x_1, x_2\}$, $\{x_2\}$) result in low ATE error.



(c) Success probability of the set $\{x_2\}$ in the toy example \mathcal{G}^{toy} for different p_{value} thresholds.



(b) Performance of Algorithm 1 on \mathcal{G}^{toy} .



(d) Performance of Algorithm 1 on \mathcal{G}^{toy} when the candidate adjustment sets are \tilde{d} -dimensional.

Figure 11: Validating our theoretical results and our Algorithm 1 on \mathcal{G}^{toy} when $e = f(x_t)$.

K.4 IHDP


In this section, we provide more details on experiments in Section 5.2 on the IHDP⁸ dataset.

Dataset Description. First, we describe various aspects measured by the features available in this dataset. The feature set comprises of the following attributes (a) 1-dimensional : child’s birth-weight, child’s head circumference at birth, number of weeks pre-term that the child was born, birth order, neo-natal health index, mother’s age when she gave birth to the child, child’s gender, indicator for whether the child was a twin, indicator for whether the mother was married when the child born, indicator for whether the child was first born, indicator for whether the mother smoked cigarettes when she was pregnant, indicator for whether the mother consumed alcohol when she was pregnant, indicator for whether the mother used drugs when she was pregnant, indicator for whether the mother worked during her pregnancy, indicator for whether the mom received any prenatal care, (b) 3-dimensional : education level of the mother at the time the child was born, and (c) 7 -dimensional : site indicator.

The set of all observed features satisfies the back-door criterion for IHDP. As described in Section 5.2, the outcome simulated by the setting “A” of the NPCI package depends on all the observed features. In other words, there is a direct edge from each of the observed feature to the outcome y in this scenario. Also, recall from Section 5.2 that the feature set is pre-treatment (i.e., it satisfies Assumption 1). Therefore, from Remark 5, $\mathbf{z} \subseteq \mathbf{x}$ satisfies the back-door criterion relative to (t, y) in \mathcal{G} if and only if t and y are d-separated by \mathbf{z} in \mathcal{G}_{-t} . Here, when \mathbf{z} is the set of all observed features, it is easy to see that t and y are d-separated by \mathbf{z} in \mathcal{G}_{-t} . Therefore, the set of all observed features satisfies the back-door criterion.

Choices of features in $\mathbf{x}^{(o)}$. As mentioned in Section 5.2, we keep the feature child’s birth-weight in $\mathbf{x}^{(o)}$. In addition to these, we also keep the number of weeks pre-term that the child, child’s head circumference at birth, birth order, neo-natal health index, mother’s age when she gave birth to the child, child’s gender, indicator for whether the mother used drugs when she was pregnant, indicator for whether the mom received any prenatal care, and site indicator in $\mathbf{x}^{(o)}$.

Existence of valid adjustment sets of size 5. Since $\mathbf{x}^{(o)}$ comprises of only 10 different features, the set of all subsets of $\mathbf{x}^{(o)} \setminus \{x_t\}$ comprises of 512 elements for any x_t . Therefore, in principle, one could find the set with lowest ATE error amongst these 512 candidate adjustment sets instead of the averaging performed by our algorithm (Algorithm 1). In an attempt to do this for comparison with our algorithm, we accidentally came across the following subset of features : $\mathbf{x}^{(m)} = \{\text{child’s head circumference at birth, birth order, indicator for whether the mother used drugs when she was pregnant, indicator for whether the mom received any prenatal care, site indicator}\}$. The ATE estimated using $\mathbf{x}^{(m)}$ to adjust (termed as ‘the oracle’) significantly outperforms the ATE estimated using $\mathbf{x}^{(o)}$ to adjust (termed as ‘the baseline’ i.e., **Baseline**) as shown in Figure 12a.



(a) Comparison of **Baseline** (i.e., adjusting for $\mathbf{x}^{(o)}$) with oracle (i.e., adjusting for $\mathbf{x}^{(m)}$) on IHDP

(b) Success probability of the set $\mathbf{x}^{(m)}$ in IHDP for different p_{value} thresholds.

Figure 12: Additional analysis on IHDP.

Therefore, we believe that there exist valid adjustment sets of size 5 (or adjustment sets better than $\mathbf{x}^{(o)}$) for this dataset. Therefore, to curtail the run-time of **Exhaustive**, we consider **Sparse** with $\mathcal{X} = \text{subsets of } \mathbf{x}^{(o)} \setminus \{x_t\}$ with size at-most 5 in Section 5.2. However, as mentioned in Section 5.2, the performance of **Sparse** is similar to

⁸https://github.com/vdorie/npci/blob/master/examples/ihdp_sim/data/ihdp.RData

that of **Exhaustive** since (a) **Sparse** has to perform 382 tests and (b) there is no guarantee that $\mathbf{x}^{(m)}$ will be picked as a valid adjustment set (as explained below). Finally, we point out that the performance of **IRM-t** is closest to ‘the oracle’ as evident from Figure 4.

Success Probability. Similar to Section K.2, we consider the success probability of the set $\mathbf{x}^{(m)}$. For a given p_{value} threshold, we let the success probability of the set $\mathbf{x}^{(m)}$ be the fraction of times (in n_r runs) the p-value of $\text{CI}(e \perp_p y | \mathbf{x}^{(m)}, t)$ is more than p_{value} . In Figure 12b, we show how the success probability of the set $\mathbf{x}^{(m)}$ varies with different p_{value} thresholds i.e., $\{0.1, 0.2, 0.3, 0.4, 0.5\}$ for IHDP.

K.5 Cattaneo

In this section, we provide more details on experiments in Section 5.3 on the Cattaneo⁹ dataset.

Dataset Description. We describe various aspects measured by the features available in this dataset. The feature set comprises of the following attributes : mother’s marital status, indicator for whether the mother consumed alcohol when she was pregnant, indicator for whether the mother had any previous infant where the newborn died, mother’s age, mother’s education, mother’s race, father’s age, father’s education, father’s race, months since last birth by the mother, birth month, indicator for whether the baby is first-born, total number of prenatal care visits, number of prenatal care visits in the first trimester, and the number of trimesters the mother received any prenatal care. Apart from these, there are also a few other features available in this dataset for which we did not have access to their description.

K.6 Training details

For all of our experiments, we split the data randomly into train data and test data in the ratio 0.8 : 0.2. We use ridge regression with cross-validation and regularization strengths : 0.001, 0.01, 0.1, 1 as the regression model. We mainly relied on the following github repositories — (a) **causalib**¹⁰ (Shimoni et al., 2019), (b) **RCoT** (Strobl et al., 2019), (c) **ridgeCV**¹¹, and (d) **IRM**¹².

For IRM, we use 15000 iterations. We train the IRM framework using 2 environments and perform validation on the remaining environment. For validation, we vary the learning rate (of the Adam optimizer that IRM uses) between 0.01 and 0.001 and vary the IRM regularizer between 0.1 and 0.001. During training, we use a step learning rate scheduler which decays the initial learning rate by half after every 5000 iterations.

K.7 Comparison with Entner et al. (2013) and Gultchin et al. (2020)

As described in Section 3, Entner et al. (2013) and Gultchin et al. (2020) cannot be used to conclude that $\emptyset, \{x_3\}, \{x_2, x_3\}$ are not admissible i.e., not valid backdoors in \mathcal{G}^{toy} (because the variable $x_t = x_1$ has an unobserved parent) while our Theorem 3.2 can be used to conclude that. Here, we provide the p-values (averaged over 100 runs) corresponding to these in Table 1. As we can see, our invariance test results in a very small p-value for $\emptyset, \{x_3\}$, and $\{x_2, x_3\}$ leading to the conclusion that they are not valid backdoors in \mathcal{G}^{toy} .

Table 1: p-value of $\text{CI}(e \perp_p y | \mathbf{z}, t)$ for $\mathbf{z} = \emptyset, \{x_3\}$, or $\{x_2, x_3\}$ in \mathcal{G}^{toy} .

\mathbf{z}	$\mathbf{d} = 3$	$\mathbf{d} = 5$	$\mathbf{d} = 7$
\emptyset	$1.3 \times 10^{-15} \pm 2.9 \times 10^{-16}$	$1.1 \times 10^{-15} \pm 1.5 \times 10^{-16}$	$1.9 \times 10^{-15} \pm 6.7 \times 10^{-16}$
$\{x_3\}$	$1.4 \times 10^{-15} \pm 2.7 \times 10^{-16}$	$1.2 \times 10^{-15} \pm 3.1 \times 10^{-16}$	$1.0 \times 10^{-15} \pm 2.7 \times 10^{-16}$
$\{x_2, x_3\}$	$1.8 \times 10^{-4} \pm 1.8 \times 10^{-4}$	$5.1 \times 10^{-3} \pm 3.6 \times 10^{-3}$	$1.9 \times 10^{-4} \pm 1.3 \times 10^{-4}$

⁹www.stata-press.com/data/r13/cattaneo2.dta

¹⁰<https://github.com/ibm/causalib>

¹¹https://github.com/scikit-learn/scikit-learn/tree/15a949460/sklearn/linear_model/_ridge.py

¹²<https://github.com/facebookresearch/InvariantRiskMinimization>

K.8 Comparison with Cheng et al. (2020)

As described in Section 3, Cheng et al. (2020) cannot be used to conclude that $\emptyset, \{x_2\}, \{x_3\}, \{x_2, x_3\}$ are not admissible i.e., not valid backdoors in the DAG obtained by adding the edge $x_1 \rightarrow y$ to \mathcal{G}^{toy} (because there is no COSO variable) while our Theorem 3.2 can be used to conclude that. Here, we provide the p-values (averaged over 100 runs) corresponding to these in Table 2. As we can see, our invariance test results in a very small p-value for $\emptyset, \{x_2\}, \{x_3\}$, and $\{x_2, x_3\}$ leading to the conclusion that they are not valid backdoors.

Table 2: p-value of $CI(e \perp_p y | \mathbf{z}, t)$ for $\mathbf{z} = \emptyset, \{x_2\}, \{x_3\}$, or $\{x_2, x_3\}$ in the DAG obtained by adding the edge $x_1 \rightarrow y$ to \mathcal{G}^{toy} .

\mathbf{z}	$\mathbf{d} = 3$	$\mathbf{d} = 5$	$\mathbf{d} = 7$
\emptyset	$1.2 \times 10^{-15} \pm 2.0 \times 10^{-16}$	$1.6 \times 10^{-15} \pm 3.2 \times 10^{-16}$	$1.3 \times 10^{-15} \pm 2.1 \times 10^{-16}$
$\{x_2\}$	$8.8 \times 10^{-7} \pm 8.8 \times 10^{-7}$	$5.3 \times 10^{-4} \pm 4.2 \times 10^{-4}$	$8.1 \times 10^{-3} \pm 4.8 \times 10^{-3}$
$\{x_3\}$	$1.6 \times 10^{-15} \pm 3.2 \times 10^{-16}$	$9.7 \times 10^{-16} \pm 1.6 \times 10^{-16}$	$2.0 \times 10^{-15} \pm 3.8 \times 10^{-16}$
$\{x_2, x_3\}$	$3.5 \times 10^{-7} \pm 3.4 \times 10^{-7}$	$1.1 \times 10^{-3} \pm 7.8 \times 10^{-4}$	$1.4 \times 10^{-3} \pm 1.3 \times 10^{-3}$