

An Introduction to Causal Discovery

Martin Huber

University of Fribourg, Department of Economics

1 Introduction

In social sciences, causal analysis (also known as impact evaluation) predominantly focuses on evaluating the causal effect of a predefined treatment variable, such as an education program, health treatment, or marketing intervention, on an outcome of interest, such as earnings, health, or sales. Causal discovery, on the other hand, which is more prominent in computer science, aims to learn the causal relationships among several or even many variables in a data-driven manner. In other words, causal discovery seeks to understand or unveil the causal associations within an entire system of variables, which can be depicted by a causal graph. The task of determining which variables influence others, based solely on statistical associations rather than presupposed causal structures (like assuming that a treatment may affect an outcome but not vice versa), poses a significant challenge. A growing number of studies demonstrate the circumstances and assumptions under which this endeavor is at least theoretically feasible, see for instance the literature reviews by Kalisch and Bühlmann (2014), Peters, Janzing, and Schölkopf (2017), and Glymour, Zhang, and Spirtes (2019) for further insights. This survey provides an introduction to concepts, algorithms, and examples of causal discovery (including empirical applications in the statistical software R) through the lens of social sciences. Section 2 discusses fundamental concepts like d-separation, causal faithfulness, and Markov equivalence. Section 3 sketches various algorithms for causal discovery. Section 4 discusses the back-door and front-door criteria for the identification of causal effects of predefined treatments. Section 5 concludes with examples of causal discovery for learning all variables that directly affect an outcome of interest and/or testing identification of causal effects in observational data.

2 d-Separation, Causal Faithfulness, and Markov equivalence

This section introduces essential concepts in causal discovery: d-separation, causal faithfulness, and Markov equivalence. D-separation, short for “dependency separation”, establishes a formal framework to analyze and comprehend the relationships between variables within causal graphs. Causal graphs describe a causal system of variables and consist of nodes, which represent in-

dividual variables or sets of variables, and arrows, which indicate causal effects between these variables. For instance, an arrow originating from a node (or variable) D and pointing to node Y signifies that variable D has a causal impact on variable Y , as illustrated in the left graph of figure 2.0.1. Furthermore, directed acyclic graphs (DAGs) rule out cyclic or simultaneous relations, like arrows going both from D to Y and Y to D . d-separation enables the identification of conditional independence relationships among variables based on the structure of a DAG.


The d-separation criterion of Pearl (1988) is founded on the concept of “blocking” or controlling for causal paths in a graph in a manner that establishes statistical (conditional) independence between two variables. This implies that after blocking specific paths connecting the two variables, they are no longer associated with each other. More concisely, a path between two (sets of) variables D and Y is blocked when conditioning on a (set of) control variable(s) C ,

1. if the path between D and Y is either a causal chain, implying that $D \rightarrow X \rightarrow Y$ or $D \leftarrow X \leftarrow Y$, or a confounding association, implying that $D \leftarrow X \rightarrow Y$, and variable (set) X is among control variables C (i.e. controlled for),
2. if the path between D and Y contains a collider, implying that $D \rightarrow S \leftarrow Y$, and variable (set) S or any variable (set) causally affected by S is not among control variables C (i.e. not controlled for).

With this definition of blocking in mind, the d-separation criterion states that variables D and Y are d-separated when conditioning on control variable(s) C if and only if C blocks all paths between D and Y . In other words, D and Y are d-separated (i) if we control for all (mediating) variables through which D affects Y or vice versa, as well as all (confounding) variables that jointly affect D and Y , and (ii) if we do not control for any (collider) variables that are jointly affected by D and Y , or any variables influenced by colliders. Since d-separation is a sufficient condition for the (conditional) independence of two variables, it is very useful for causal reasoning in complex causal models. For this reason, d-separation serves as a theoretical basis for various causal discovery algorithms designed to learn causal models from data. Reconsidering the left graph of figure 2.0.1, we observe that D and Y are not d-separated when we control

for variable X . Despite the fact that the confounding relationship $D \leftarrow X \rightarrow Y$ is blocked when conditioning on X , the causal effect $D \rightarrow Y$ remains unblocked such that D and Y are dependent conditional on X .

Figure 2.0.1: Selection on observables (left) and Y-learning (right)



To illustrate the usefulness of d-separation for causal discovery, consider the causal model depicted in the right graph of figure 2.0.1, which pertains to Y-learning, as e.g. discussed in Mani, Spirtes, and Cooper (2012). Variables X and Z are independent of each other when not controlling for D , while they are dependent conditional on D . According to the d-separation criterion, this necessarily implies that D is a collider ($X \rightarrow D \leftarrow Z$). Furthermore, both X and Z are associated with Y , as they both affect Y via D . However, when controlling for D , Y is conditionally independent of both X and Z . According to the d-separation criterion, this necessarily implies that D functions as a mediator, transmitting the causal effects of X and Z to Y ($X \rightarrow D \rightarrow Y$, $Z \rightarrow D \rightarrow Y$). In specific causal models, d-separation may enable the comprehensive learning of the entire causal model as in the Y-learning example, but this is not the case in general and unlikely to apply to most causal problems.

While d-separation is sufficient for the conditional independence of two variables, it is important to note that in special cases, two variables might be independent even if d-separation fails. To illustrate this, let us revisit the left graph in figure 2.0.1. If the confounding path of D and Y due to X ($D \leftarrow X \rightarrow Y$) perfectly offsets the causal effect of D on Y ($D \rightarrow Y$), then D and Y are statistically independent even though d-separation fails (due to the presence of both

confounding and the causal effect of D on Y). Causal discovery methods typically rely on the absence of different paths between variables that exactly cancel each other out. For this reason, we may need to explicitly rule out offsetting paths through the causal faithfulness (or stability) assumption, see for instance the discussions in Pearl (2000) and Spirtes, Glymour, and Scheines (2000). Causal faithfulness imposes that only variables that are d-separated in a DAG are conditionally independent, while any variables that are not d-separated are dependent. Consequently, the assumption mandates that d-separation is not only a sufficient, but also a necessary condition for conditional independence. In other words, two variables are statistically independent if and only if d-separation is met.

However, even with the application of causal faithfulness, d-separation may not (and in most observational data will not) lead to a unique determination of the causal model that underlies the observed associations of variables in the data. This implies that the conditional dependence and independence relationships detected through the d-separation criterion might be compatible with multiple causal models, resulting in ambiguity regarding the true causal structure. This issue motivates the concept of Markov equivalence classes for characterizing causal models that entail identical conditional dependencies and independencies. Distinct causal models belong to the same Markov equivalence class if they exhibit the exact same patterns of conditional dependence and independence in the data, such that the causal models are indistinguishable based on observational data alone.

To provide an example, let us consider two distinct causal models. In the first model, D affects Y exclusively through a mediator M ($D \rightarrow M \rightarrow Y$). In the second model, Y affects D exclusively through M ($D \leftarrow M \leftarrow Y$). In the absence of any further variables, the two models generate precisely the same patterns of conditional dependence and independence: D and Y are conditionally independent given M , while D and M are dependent given Y and Y and M are dependent given D , and any of D, M, Y are mutually dependent in the absence of conditioning on any variable. As a result, both models belong to the same Markov equivalence class, because we are unable to distinguish between these models based solely on the statistical associations in the data. Causal discovery becomes even more involved when allowing for unobserved variables in a causal model, e.g. an unobserved confounder U jointly affecting D and Y ($D \leftarrow U \rightarrow Y$).

Since the d-separation criterion cannot be applied to unobserved variables, their presence tends to exacerbate the uncertainty regarding the true causal model. Consequently, this can increase the number of causal models that align with the same Markov equivalence class.

In scenarios where the correct causal model cannot be learned solely from observational data, the challenge of model ambiguity can potentially be addressed through “external” sources of information concerning the causal structure. Such sources may include domain expertise, theoretical insights, past empirical findings, or knowledge about the temporal sequence of events. These external insights can provide valuable context that helps narrow down potential causal models. For example, if variable D (e.g., training participation) is measured at an earlier time point than variable Y (e.g., earnings after training participation), it becomes evident that a causal path from Y to D (such as $D \leftarrow Y$) can be ruled out. This is because future events cannot exert influence on past events, a principle discussed in works like Reichenbach (1991) concerning the connection between causality and time. Moreover, in order to investigate whether an observed association between training D and earnings Y is solely due to confounding ($D \leftarrow U \rightarrow Y$), or if it also involves a nonzero treatment effect ($D \rightarrow Y$), an experiment could be conducted in which the treatment is randomly assigned, as random assignment avoids confounding. Consequently, external sources of information might be helpful for clarifying some of the ambiguous causal associations in a DAG in order to reduce the number of plausible causal models within in a Markov equivalence class.

3 A Sketch of Algorithms for Causal Discovery

As discussed in the previous section, the objective of causal discovery is to recognize Markov equivalence classes encompassing all causal models that are indistinguishable (or statistically equivalent) in observed data because they exhibit the same d-separation patterns (i.e. patterns of conditional independence). For the practical application of the d-separation criterion in causal discovery, Verma and Pearl (1990) introduced the IC (or I-equivalence Class) algorithm for observational data. The IC algorithm operates under the assumptions that causal faithfulness holds and that there are no unobserved confounders that jointly affect any pair of observed

variables for which the causal associations are to be estimated. The IC algorithm consists of the following steps:

1. For all pairs of variables A and B in the data, search for a set of variables X such that A and B are conditionally independent when controlling for X . Link A and B by an edge, which represents an undirected association in a causal graph, if and only if no X exists that satisfies conditional independence.
2. For all pairs of variables A and B that do not share an edge with each other, but share both an edge with a variable S , verify whether S is in set X of step 1. If this is not the case, S is a collider such that the causal association is $A \rightarrow S \leftarrow B$.
3. In the resulting graph with partially determined causal associations (directed causal arrows) and partially undetermined associations (edges), orient the direction of as many edges as possible subject to two conditions: (i) Any alternative orientation would yield a new collider structure. (ii) Any alternative orientation would yield a directed cycle (i.e. a circular causal relation between variables).

Bluntly speaking, step 1 of the algorithm finds those pairs of variables that are dependent conditional on any feasible set of control variables. Variables in such a pair are then connected by an undirected edge since the specific causal path remains unknown at this stage. Step 2 pinpoints collider paths which unveil the causal directions (arrows) between variables. Step 3 finds further causal associations that adhere to the constraints of not creating circular causal relations (which are prohibited in DAGs) or unwarranted further collider paths (as correct collider paths were already detected in step 2). As highlighted by Pearl (2000), step 3 can be systematized in several ways, e.g. by applying the rules provided in Verma and Pearl (1992) for orienting edges into causal arrows. These rules are sufficient to identify the maximum potential number of causal arrows in a causal graph based on observed data.

Spirtes and Glymour (1991) introduced a refined method known as the PC algorithm. It reduces computation time by limiting step 1 of searching for sets X that entail the conditional independence of A and B to variables that share edges with (i.e., are adjacent to) either A or

B. Glymour, Zhang, and Spirtes (2019) offer a comprehensive discussion and illustration of the PC algorithm based on the causal scenario of Y-learning considered in section 2. Numerous further enhancements have been proposed in causal discovery algorithms, e.g. for integrating external information regarding causal relationships between specific variables. One important contribution is the Fast Causal Inference (FCI) algorithm introduced by Spirtes, Glymour, and Scheines (2000), which allows for and in some causal models even detects the presence of unobserved variables. Glymour, Zhang, and Spirtes (2019) discuss an example for the application of the FCI algorithm in the presence of unobserved confounders.

An important question in causal discovery is how to perform statistical inference, e.g. to obtain p-values or confidence intervals for the estimated causal effects between variables. To test the various conditional independence assumptions between pairs of variables in the data, algorithms typically adopt either pairwise test statistics like t-tests or global goodness-of-fit statistics that are computed collectively for all variables (rather than pairs of variables), such as the Bayes Information Criterion (BIC). In this context, it is important to acknowledge that verifying conditional independencies among multiple variables amounts to jointly testing several hypotheses and therefore introduces multiple hypothesis testing issues. Consequently, the critical values of pairwise tests must be adjusted to account for the number of tested causal associations (which grow exponentially in the number of variables in the causal model), e.g. by a Bonferroni-type adjustment, see e.g. Holm (1979). Otherwise, there is an increased risk that testing might (by random chance) erroneously indicate the presence of specific causal associations that are actually absent, implying an increased type I error rate in testing. However, on the flip side, such corrections for multiple hypothesis testing may decrease the power of the tests, implying a reduced chance (or probability) to detect non-zero causal effects, implying an increased type II error rate in testing. This issue becomes particularly pronounced when when there are many variables in the causal model, which implies more rigorous corrections.

The algorithms discussed so far explore conditional independencies and dependencies to estimate the set of causal models that fit within the Markov equivalent class of the true (albeit unknown) causal model. However, there are alternative avenues to causal discovery that rely on specific assumptions about the functional nature of causal relationships between variables. Under

these assumptions, it becomes possible to uniquely determine the true causal model without relying on d-separation or causal faithfulness. As discussed in Hoyer, Janzing, Mooij, Peters, and Schölkopf (2008), Zhang and Hyvärinen (2009a), and Peters, Mooij, Janzing, and Schölkopf (2014), causal relationships can be learned from the data under the following conditions:

1. The causal effect of any variable (say, D) on another variable (say, Y) is defined by a nonlinear function.
2. The effects of unobserved variables on Y can be expressed by a single error term that is independent of D and does not interact with D (such that the effects of unobserved variables on Y are homogeneous across values of D).

In other words, if causal associations between observed variables are nonlinear and causal effects of unobserved variables can be represented by additively separable (i.e., non-interacting) and independent error terms, then we can deduce the direction of causality between observed variables from the data. Formally, this requires that Y is characterized by the following model:

$$Y = \mu(D) + \varepsilon, \tag{1}$$

where μ is a nonlinear function of D and ε is an additive error term that is independent of D .

The nonlinearity of μ carries an important implication. When erroneously assuming that Y affects D and estimating a reverse (or “anticausal”) association of D as a function of Y , the resulting error term will not be independent of Y . This holds true except for highly specific cases, as elaborated in Peters, Janzing, and Schölkopf (2017). Notably, in scenarios beyond these special cases, this phenomenon enables the identification of the correct causal model, wherein D affects Y . In this correct model, the error term is independent of D . In practice, testing the causal direction may be based on running nonlinear regressions of both Y on D and D on Y and verifying in which of these two cases the estimated errors (or residuals) are independent of the regressors. To this end, we can apply a statistical test for the independence of the estimated errors and regressors in either regression and choose that causal model as the supposedly correct one for which the test yields a higher (and statistically insignificant) p-value. To avoid overfitting

bias when computing the p-values, sample splitting is advisable, which implies estimating the regression models and conducting the independence tests in distinct subsamples obtained from randomly splitting the full sample. Sample splitting ensures that the regression and testing stages are not statistically associated with each other.

The model in equation (1) can be extended or modified in various ways, all while preserving the ability to identify causality. For instance, Breunig and Burauel (2021) consider testing the direction of causality between two variables D and Y when controlling for observed covariates, denoted by X , while allowing for heteroskedasticity of ε in X . As the independence between ε and D is only required to hold conditional on X , this implies a type of selection-on-observables assumption. Formally, the setup is defined by the following nonlinear model:

$$Y = \mu(D, X) + \varepsilon, \tag{2}$$

where the variance of error term ε may differ across values of X .

Furthermore, Zhang and Chan (2006) and Zhang and Hyvärinen (2009b) consider a generalization of equation (1), known as post-nonlinear (PNL) model, which consists of two nested non-linear functions:

$$Y = q(\mu(D) + \varepsilon), \tag{3}$$

where both q and μ are non-linear functions and q is assumed to be invertible. It is evident that equation (1) is a special case of equation (3) when q is defined as the identity function. For this reason, equation (3) is more general, because it permits Y to be a complex (rather than additive) function of the impact of the treatment, characterized by $\mu(D)$, and of unobserved variables, characterized by ε . Another special case of both equations (1) and (3) is a model in which the causal associations between observed variables D and Y are linear rather than nonlinear:

$$Y = \alpha + \beta D + \varepsilon, \tag{4}$$

where α is the constant term and β is the causal effect of D on Y . Shimizu, Hoyer, Hyvärinen, Kerminen, and Jordan (2006) show that the correct causal model is identified if the additively separable error term ε is non-normally distributed. The latter assumption is crucial, as unique identification in linear models is generally not feasible if errors are normally distributed, unless additional assumptions are imposed. For instance, Peters and Bühlmann (2013) show that in linear models with normally distributed errors, causality can still be learned under the (strong) assumption that the error terms in the various equations of a causal model all have the same variance.

To demonstrate the practical implementation of causal discovery algorithms in R, we load the *bnlearn* package created by Scutari (2010) using the *library* command. Additionally, we load the *causalweight* package and utilize the *data* command to analyze the *coffeeleaflet* data set, which stems from an experimental study aimed at evaluating the impact of a leaflet discussing coffee production implications on the environmental awareness of Bulgarian students. We create a new frame set named *data*, which only contains the pretreatment characteristic “mother’s education” (the variable *mumedu* in the 8th column of the *coffeeleaflet* data), the randomly assigned leaflet treatment (the variable *treatment* in the 32nd column), and the outcome “awareness of waste production due to coffee production” measured on a five-point scale (the variable *awarewaste* in the 38th column): `data=coffeeleaflet[,c(8,32,38)]`. As the variable *treatment* is randomly assigned, it should be independent of *mumedu*, while both variables may affect the outcome *awarewaste*. In this case, we have the causal relation $treatment \rightarrow awarewaste \leftarrow mumedu$, such that *awarewaste* is a collider. We feed our data set into the *pc.stable* function to conduct causal discovery based on the PC algorithm and store the results in an R object named *output*. Finally, we employ the *plot* command to visualize the causal graph derived from the PC algorithm. The box below provides the R code for each of the steps.

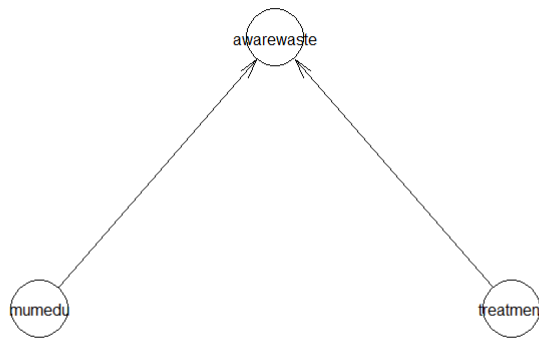
```

library(bnlearn)           # load bnlearn package
library(causalweight)     # load causalweight package
data(coffeeleaflet)      # load coffeeleaflet data
data=coffeeleaflet[,c(8,32,38)] # select variables nr. 8, 32, and 38
output=pc.stable(data)   # run the PC algorithm
plot(output)             # plot the learned causal model

```

Executing the code produces the DAG depicted in figure 3.0.1. We see that the algorithm detects the previously mentioned collider structure, implying that mother’s education and treatment assignment are independent (as expected in a well conducted experiment), while they both affect the awareness outcome. As mentioned before, such algorithms might fail to uniquely determine the direction of the causal arrows in more complicated models with more variables and more involved dependence structures than in our toy model containing a single collider. And even in our toy model, we acknowledge that the causal arrow from mother’s education to the awareness outcome might be confounded, as mother’s education is not randomly assigned.

Figure 3.0.1: Collider structure



As a further R example, we consider learning the direction of causality in nonlinear models characterized by equation (2) by testing the independence of the estimated nonlinear regression

function and the error terms (or residuals). Using the *library* command, we load the *dHSIC* package for independence testing, the *mgcv* package for running nonlinear regressions, and the *datarium* package. The latter contains the *marketing* data set consisting of 200 observations with information on sales and advertising budgets, which we load using the *data* command. We are interested in learning the causal relationship between the variable *newspaper*, which measures the budget of advertisement in newspapers and which we define as treatment D , and the variable *sales*, which we define as outcome Y .

We set a seed (*set.seed(1)*) for the replicability of our results to follow and use the *gam* function to estimate two different regression models. In the first regression, we estimate the outcome Y as a nonlinear function of the treatment D : $model1 = gam(Y \sim s(D))$, where the wrapper $s(D)$ implies that the association is measured by series regression. We store the output in an R object named *model1*. In the second (reverse) regression, we estimate the treatment D as a nonlinear function of the outcome Y and store the output in the object *model2*. We then feed the residuals of the first regression model, *model1\$residuals*, and D into the *dhsic.test* command to test the independence between both variables and save the results in object *test1*. Using the same command, we test the independence between the residuals of the second regression model and Y and save the results in object *test2*. Finally, we investigate the p-values of the tests stored in *test1\$p.value* and *test2\$p.value*.

```

library(dHSIC)           # load dHSIC package
library(mgcv)           # load mgcv package
library(datarium)       # load datarium package
data(marketing)         # load marketing data
D=marketing$newspaper   # define treatment (newspaper advertising)
Y=marketing$sales       # define outcome (sales)
model1=gam(Y ~ s(D))   # estimate Y as nonlinear function of D
model2=gam(D ~ s(Y))   # estimate D as nonlinear function of Y
set.seed(1)             # set seed
test1=dhsic.test(model1$residuals, D) # independence test for first model
test2=dhsic.test(model2$residuals, Y) # independence test for second model
test1$p.value; test2$p.value         # show p-values

```

The tests yield p-values of 0.136 (or roughly 14%) and 0.012 (or roughly 1%) for the regressions of Y on D and of D on Y , respectively. Consequently, we find that a causal effect of the supposed treatment D on outcome Y cannot be rejected at the 10% level of statistical significance. In contrast, a reverse causal effect from Y to D is rejected at the 5% level of significance.

4 The Back-Door and Front-Door Criteria


A concept closely related to d-separation discussed in section 2 is the back-door criterion, see e.g. Pearl (2000). It provides a causal graph-based framework for determining variables that must be controlled for to prevent confounding bias when measuring the causal effect of a treatment variable D , on an outcome variable Y . More concisely, the back-door criterion employs d-separation to detect an appropriate set of control variables such that any confounding association between D and Y , also known as back-door path, is eliminated or blocked. More formally, a set of covariates X satisfies the back-door criterion in a DAG where the effect of D on Y is of interest,

1. if X blocks any path between D and Y that contains an arrow (causal effect) into D ,
2. if no variable in X is causally affected by D .

To provide an example illustrating the concept of a back-door path and the back-door criterion, let us reconsider the left graph in figure 2.0.1. In this case, the presence of covariates X that jointly affect D and Y forms a back-door path ($D \leftarrow X \rightarrow Y$) which confounds the causal relation of D and Y . Controlling for X satisfies the back-door criterion, as it closes the back-door path (i.e. controls for confounding), while not including any variable that is causally affected by D . This implies that the so-called selection-on-observables or unconfoundedness assumption is satisfied with respect to D , meaning that the treatment is as good as randomly assigned conditional on X , see for instance Imbens (2004) for a discussion of this assumption. Consequently, the treatment effect of D on Y is identified when controlling for X .

In contrast to the previous example, figure 4.0.1 presents a causal model where controlling

Figure 4.0.1: M-bias



for X does not satisfy the back-door criterion. Although the back-door condition that X is not causally affected by D holds, the condition that X blocks any path between D and Y containing an arrow into D is not met. This situation arises because X is affected by both V and U and thus a collider, such that controlling for X introduces dependence between V and U . As V affects D and U affects Y , this collider bias generates a dependence or back-door path between D and Y , leading to the failure of the back-door criterion. This specific form of collider bias is known as M-bias. However, the back-door criterion holds when not controlling for X , thus preventing collider bias, as V (which influences D) is not associated with Y when X is not controlled for. Alternatively, the back-door criterion can be met by controlling for X and either V or U , or both. The rationale here is that controlling for either V or U blocks the confounding bias between D and Y that is introduced by controlling for X .

It is worth noting that a web-based application called “DAGitty” (Directed Acyclic Graph Interactive Tool) for creating, analyzing, and visualizing causal graphs is available at <https://dagitty.net/>. It makes use of the back-door criterion to determine if and under which conditions the effect of a treatment D on an outcome Y is identified in a causal model. DAGitty provides a visual representation of causal graphs, allows differentiation between observed and unobserved variables, and highlights all backdoor paths between D and Y that need to be controlled for even in complex models with numerous variables. This tool proves useful for researchers and analysts in comprehending the underlying causal structure and making correct decisions about the fea-

sibility of identifying treatment effects based on observed variables. If identification is possible, DAGitty also indicates the covariates that need to be controlled for. DAGitty is also available as R package, which is provided by Textor, van der Zander, Gilthorpe, Liskiewicz, and Ellison (2017).


A further criterion for identifying the causal effect of D on Y is the so-called front-door criterion, see Pearl (2000), which does not rely on blocking all back-door paths between D and Y . Yet, it builds on the back-door criterion by applying it sequentially with respect to the effect of D on a mediating variable M , through which presumably any effect of D on Y operates, and the effect of M on Y . Formally, a set of variables M satisfies the front-door criterion for identifying the causal effect of D on Y ,

1. if any effect of D on Y operates via M ,
2. if there is no unblocked back-door path from D to M ,
3. if all back-door paths from M to Y are blocked when controlling for D .

The first assumption states that M fully mediates the effect of D on Y . The second assumption imposes that no variables jointly affect D and M , ensuring that the treatment-mediator relation is unconfounded and any potential mediator, $M(d)$, is independent of D . We note that d is either 0 or 1 if the treatment is binary, but the front-door criterion also applies to treatments with multiple values. The third assumption asserts that conditional on D , no variables jointly affect M and Y , such that any potential outcome, $Y(d)$, is independent of M conditional on D . Alongside these assumptions, the identification of causal effects based on the front-door criterion requires a specific common support restriction: For a binary treatment D , it must hold that $0 < \Pr(D = 1|M) < 1$, ensuring that both treated and untreated observations exist for each possible value of M in the population.

The front-door criterion is related to both causal mediation models and instrumental variables. The assumptions imply “complete” mediation in the sense that there only exists an indirect effect of D on Y via M but no direct effect, similar to the exclusion restriction in instrumental variable methods. The left graph in figure 4.0.2 provides an illustration of the

Figure 4.0.2: Front-door criterion without and with controlling for covariates (left and right)



front-door criterion. U denotes unobserved confounders of D and Y , where the dashed arrows imply that the causal paths from U cannot be observed. In contrast, confounders of the causal association between treatment D and mediator X (which is on the causal path between D and Y) or between mediator M and outcome Y conditional on D must not exist. This still permits unobservables affecting M (which are omitted in the graph), as long as those unobservables do not affect D or Y , too, and are not associated with U . If the front-door criterion holds, the mean potential outcome under treatment state d (with d being either 0 or 1 for a binary treatment) corresponds to the following expression, see for instance Frölich and Sperlich (2019) for a derivation of this result:

$$E[Y(d)] = E[\nu(M)|D = d], \quad (5)$$

where $\nu(m) = E[E[Y|D, M = m]]$ is a nested mean outcome, in which the conditional mean outcome $E[Y|D, M = m]$ is averaged over values of the treatment D , while keeping the mediator fixed value $M = m$. For a binary treatment, $\nu(m) = E[Y|D = 1, M = m] \cdot \Pr(D = 1) + E[Y|D = 0, M = m] \cdot \Pr(D = 0)$, with \Pr denoting a probability. In a second step, we obtain the mean potential outcome $E[Y(d)]$ by averaging $\nu(m)$ across all values of M in the treatment group with $D = d$, as shown in equation (5). For a binary mediator M , for instance, we have that $E[\nu(M)|D = d] = \nu(1) \cdot \Pr(M = 1|D = d) + \nu(0) \cdot \Pr(M = 0|D = d)$.

The average treatment effect (ATE) of a binary treatment corresponds to the difference in the expressions for the mean potential outcomes under treatment and nontreatment:

$$\Delta = E[\nu(M)|D = 1] - E[\nu(M)|D = 0]. \quad (6)$$

Equation (6) calculates ATE as the average change in outcome Y due to a shift in mediator M induced by switching the treatment from $D = 1$ to $D = 0$, which utilizes the first assumption that any effect of D on Y operates through M . Furthermore, the third assumption ensures that D blocks all back-door (or confounding) paths from M to Y , implying that by controlling for D in the nested mean outcome $\nu(M)$, we can assess the effect of M on Y by varying the values of M in $\nu(M)$. Finally, given the second assumption that there is no back-door path between D and M , the effect of D on M is identified as well. As a result, we can compute the mean potential outcomes and the ATE by simply averaging $\nu(M)$ within the treated and nontreated groups.

While equation (5) identifies mean potential outcomes as required for the evaluation of the ATE, the front-door criterion was originally formulated in terms of the probability that a potential outcome takes a specific value y , denoted as $\Pr(Y(d) = y)$. As discussed in Pearl (1995), this potential outcome probability is identified through the following equation:

$$\Pr(Y(d) = y) = \sum_m \Pr(M = m|D = d) \cdot \left(\sum_{d'} \Pr(Y = y|D = d', M = m) \cdot \Pr(D = d') \right), \quad (7)$$

where we assume that all of Y, D, M are discrete variables such that we can measure their probabilities.

The front-door criterion has been relatively infrequently applied in empirical research. One potential application discussed in Pearl (2000) is the assessment of the causal effect of smoking (D) on lung cancer (Y) mediated by tar deposits (M). If unobserved genetic factors (U) simultaneously influence smoking behavior and lung cancer, then the back-door criterion fails for D and Y due to confounding. Nevertheless, the effect of smoking on lung cancer can po-

tentially be identified if the front-door criterion is met. This criterion necessitates that there are no unobservable variables affecting both smoking behavior and tar deposits, or tar deposits and lung cancer, and that smoking solely affects lung cancer through tar deposits (exclusion restriction). As noted e.g. by Imbens (2020), the restrictions on the unobservables imposed by the front-door criterion fail in this application if genetic factors influence smoking behavior, the inclination to accumulate tar, and lung cancer altogether, or if unobservables (like a hazardous working environment) jointly affect tar deposits and lung cancer. Additionally, if there exist alternative causal pathways through which smoking behavior affects lung cancer beyond tar deposits, then the exclusion restriction assumption fails. In numerous empirical contexts, meeting the requirements of the front-door criterion can be challenging, potentially explaining its limited prevalence in practical applications.

However, it is worth mentioning that the assumptions underlying the front-door criterion can be relaxed to some extent. For example, the assumptions might appear more plausible in empirical applications when allowing observed covariates X to affect all of D , M , and Y , as illustrated in the right graph of figure 4.0.2. In this case, the ATE is identified when controlling for X . Consequently, this entails modifying expression (6) to

$$\Delta = E_X[E_M[\nu(M, X)|D = 1, X] - E_M[\nu(M, X)|D = 0, X]], \quad (8)$$

where $\nu(m, x) = E[E[Y|D, M = m, X = x]|X = x]$ is a nested conditional mean outcome given X . Intuitively, we perform similar calculations as before, but this time within groups that share the same values of the covariates X to obtain the conditional ATE (CATE), assuming that the front-door criterion holds conditional on X . Finally, we average the CATE across all values of X (as indicated by the subscript X in the expectation operator E_X) to obtain the ATE for the entire population.

Another relaxation of the front-door criterion is to allow for a direct effect of D on Y that does not operate through M , thereby violating the exclusion restriction. As demonstrated in Fulcher, Shpitser, Marealle, and Tchetgen Tchetgen (2019), it is then still feasible to assess the

indirect treatment effect of D on Y that operates via M , whereas the (total) ATE of D on Y and the direct effect (not operating via M) are not recovered. Furthermore, the authors propose a doubly robust (DR) method, which provides an estimate of the indirect effect when the exclusion restriction is violated (but the other front-door assumptions hold) and an estimate of the ATE when all front-door assumptions hold, see also the discussion in Gorbach, de Luna, Karvanen, and Waernbaum (2023). Yet another DR estimator permits assessing causal effects when the front-door criterion is satisfied conditional on covariates X , as considered in equation (8). This estimator may be combined with machine learning to control for covariates in a data-driven way.

5 Examples of Causal Discovery


As discussed in section 3, attempting to uncover all causal relationships between all observed (or even unobserved) variables in a causal model can be very challenging or practically infeasible in many empirical scenarios. This section delves into a few examples where causal discovery is applied to somewhat more modest objectives than learning an entire causal structure. Yet, these objectives appear interesting in causal analysis and may be more feasible from a practical standpoint.

For example, researchers or analysts might have an interest in identifying all observed variables that directly influence a specific outcome variable Y . In other words, one aims at detecting all treatment variables exerting a direct causal impact on Y . For instance, we may want to determine all observed variables that have an effect on earnings (which could include education, profession, work experience, labor market conditions, among others) or affect health (which could include health behavior, genetic factors, living/working environment, among others). Learning treatments requires imposing specific identifying assumptions, such as the condition that any variables jointly affecting any treatment and the outcome are observed in the data. This implies that the selection-on-observables assumption holds for each of the potentially numerous treatments under consideration, which is a strong assumption. Furthermore, we need to rule out reverse causal effects of the outcome Y on the treatments. This appears plausible if the observed variables are measured at an earlier point in time than the outcome, as variables measured later

cannot impact the values of variables measured earlier.

The graph on the left in figure 5.0.1 provides a visual illustration, in which the observed variables D and X causally affect the outcome Y . In contrast, Z is not part of the treatments directly affecting Y , as it only exerts an indirect effect on the outcome through D . For both D and X , the selection-on-observables assumption holds in the left graph, as the unobserved term U affecting the outcome (where the non-observability of the effect is indicated by the dashed line) neither affects D nor X . Conversely, the selection-on-observables assumption fails for X in the right graph of figure 5.0.1, where U jointly affects X and Y .

Figure 5.0.1: Learning treatments and/or the identification of treatment effects



Soleymani, Raj, Bauer, Schölkopf, and Besserve (2022) provide an algorithm in this context that selects treatment variables in a data-driven way, while also controlling for observed covariates in a data-driven way based on the double machine learning (DML) framework, see Chernozhukov, Chetverikov, Demirer, Duflo, Hansen, Newey, and Robins (2018). The idea is to sequentially consider each of the observed variables that may affect Y as treatment D , while considering all remaining observed variables as covariates X to estimate the direct effect of each candidate treatment on the outcome Y by DML. Finally, the algorithm retains only those observed variables as treatment variables that exhibit statistically significant effects on Y (where judging statistical significance should account for issues related to testing multiple hypotheses based on multiple candidate treatments). Another machine learning-based algorithm proposed by Quinzan, Soleymani, Jaillet, Rojas, and Bauer (2023) adopts a somewhat distinct testing

approach. In a first step, the algorithm estimates the outcome as a function of all observed variables by DML. In a second step, it repeats the estimation of the outcome when excluding one candidate treatment D . If the omission of D statistically significantly alters the outcome estimate compared to using all observed variables (as done in the first step), then D is retained as treatment variable, otherwise it is dropped. By sequentially applying this process to all candidate treatments, the algorithm provides an estimated set of treatments that have a direct effect on Y .

Another example of applying causal discovery concepts involves learning whether sufficient conditions for identifying the causal effect of a treatment on interest (D) on an outcome of interest (Y) hold in observational data when imposing causal faithfulness (see section 2). In contrast to the previous example, the objective here is not to learn all treatments influencing an outcome, but rather to assess whether the effect of a predefined treatment is identifiable within the available data. As outlined by Huber and Kueck (2023), if there exists an instrument Z that affects D and is conditionally independent of Y given X , then this is a sufficient condition for the selection-on-observables assumption. Consequently, the causal effect of D on Y is identified conditional on X . This result follows from an application of the d-separation criterion: If covariates X permit controlling for all backdoor paths (or confounders) between Z and Y as well as D and Y and if Z is a valid instrument in the sense that it does not directly affect the outcome Y other than through D , then controlling for D in addition to X leads to the independence of Z and Y . In particular, conditioning on D does not introduce collider bias (see section 2) between Z and Y given X , implying that D is as good as randomly assigned conditional on X , such that the selection-on-observables assumption holds. In both graphs of figure 5.0.1, the testable implication holds for treatment D , implying that the effect of D on Y is identified conditional on X . Instrument Z affects outcome Y only through treatment D and the unobservable U does not jointly affect Z and Y or D and Y conditional on X . Therefore, Z and Y are conditionally independent given X and D .

These insights can be used in a causal discovery approach to learn sets of instruments Z and covariates X that satisfy the conditional independence of Z and Y given X and D (if such sets exist at all), as outlined in Apfel, Hatamyar, Huber, and Kueck (2023). The method is

based on considering all observed variables which are not the treatment D or the outcome Y and assuming one of them to be a candidate instrument Z , while the remaining ones (or subsets thereof) constitute the covariates X . In a first step, the first-stage effect of Z on D given X is estimated by DML to verify whether the conditional association between Z and D is strong enough, which is a precondition for testing. Only if this is the case, the second step consists of testing the conditional independence of Z and Y given X and D by DML. This two-step approach is iteratively repeated to each variable that is neither D nor Y , effectively cycling through all candidates for the role of the instrument Z .

For instance, assume that D is college education and Y is later life earnings, while the remaining variables are socio-economic characteristics like parental education, place of residence, among others. The algorithm would initially designate parental education as Z and the remaining variables (or their subsets) as X , followed by treating place of residence as Z with the remaining variables as X , and so forth. If for one or even several definitions of Z and X the conditional independence assumption is not rejected by the conditional independence tests, this signals the possible fulfillment of the selection-on-observables assumption, particularly in large samples with a reduced uncertainty in testing. Consequently, one can opt for the definition of control variables X (and Z) that yields the highest p-value when running the conditional independence test, thereby implying the lowest probability of violating conditional independence.

Another causal discovery method that combines elements of the previous two examples is suggested in Peters, Bühlmann, and Meinshausen (2015). This approach aims at the identification of multiple treatments (like in the first example), which directly affect an outcome Y and satisfy a selection-on-observables assumption, by means of instruments (like in the second example), also referred to as “environments” in the computer science literature. If an instrument Z affects outcome Y only through one or several other observed variables and is independent of Y conditional on a set of candidate treatment variables (rather than a single treatment), then this set will include variables that have a direct causal effect on Y . Another way of putting this is that if e.g. the conditional mean outcome does not depend on (or is invariant across) variations in instrument Z conditional on a set of candidate treatments \tilde{D} , such that $E[Y|\tilde{D}, Z] = E[Y|\tilde{D}]$, then \tilde{D} contains variables affecting Y because this set intercepts any (average) effect of Z on Y .

In light of this, Peters, Bühlmann, and Meinshausen (2015) propose an algorithm that aims at finding sets of \tilde{D} for which the estimation of $E[Y|\tilde{D}]$ remains invariant across different values of Z . The method’s capability of detecting treatments directly affecting the outcome hinges on the strength of the association between Z and the treatments. If the first-stage association is weak or absent, the estimation of $E[Y|\tilde{D}]$ will be (almost) invariant across values of Z even if important treatments are missing in set \tilde{D} . For this reason, more and stronger instruments increase the chances of finding true treatment that directly affect Y .

A final example concerns testing the identification of the causal effect of D on Y (as in the second example) when substituting instrumental variable assumptions on Z , which rule out a direct effect of Z on Y or an association between Z and unobservables U affecting Y , with a different set of assumptions. Karlsson and Krijthe (2023) present a method under the assumption that Z , representing specific environments like distinct observational datasets, directly influences all of the covariates X , treatment D , outcome Y , and unobservables U (which impact Y). However, the mechanisms through which Z affects each of X , D , U , and Y must be independent of one another. In other words, Z is presumably characterized by a set of variables that are mutually independent, with each variable only influencing one of X , D , U , and Y . When represented as $Z = (Z_X, Z_D, Z_U, Z_Y)$, the assumption of independent causal mechanisms implies, for example, that the effect of Z_D on D must not be associated with the effect of Z_U on U .

Under these conditions, the selection-on-observables assumption for treatment D can be tested using the following algorithm. First, an observation with index i is randomly selected (e.g. $i = 1$ for the first randomly sampled observation) in each environment $z \in 1, 2, \dots, \mathcal{Z}$, where \mathcal{Z} is the number of different environments. Denoting $D_i(z)$ as the treatment state of the observation with index i in environment z , the next step involves aggregating the treatments of index i across all environments into a treatment vector, denoted as $\mathcal{D}_i = (D_i(1), D_i(2), \dots, D_i(\mathcal{Z}))$. Proceeding analogously for covariates X and outcome Y , the vectors \mathcal{X}_i and \mathcal{Y}_i are obtained. The final step is to test whether, for any distinct indices $i \neq j$ (e.g., $i = 1$ and $j = 2$), \mathcal{D}_j is conditionally independent of \mathcal{Y}_i when controlling for \mathcal{D}_i , \mathcal{X}_i , and \mathcal{X}_j . A rejection of conditional independence indicates a violation of the selection-on-observables assumption. It is worth noting that the choice of observations assigned indices i and j within an environment is arbitrary if the samples

in each environment are drawn randomly; the only requirement is that i and j cannot be assigned to the same observation within an environment.

As an empirical example in R, let us consider testing whether the causal effect of a predefined treatment is identified in observational data based on a predefined instrument, using the method of Huber and Kueck (2023). We load the *causalweight* package and use the *data* command to load the *JC* data for the evaluation of the Job Corps program. The goal is to test whether the average effect of training participation in the first year after program assignment (D) on the health state four years after assignment (Y) is identified when controlling for the baseline covariates (X) measured prior to training. To this end, we test whether random assignment to Job Corps (Z), which has a strong first-stage effect on D , is conditionally mean independent of Y given D and X . To perform this test, we define the variables $Z=JC[,1]$ (as the first column in the *JC* data contains the binary assignment variable), $D=JC[,37]$, $X=JC[,2:29]$ (as columns 2 to 29 contain the pretreatment covariates), and $Y=JC[,46]$. Next, we feed these variables into the *identificationDML* command, a DML-based test that by default applies lasso regression to control for X and D in a data-driven way. We store the results in an R object named *output* and inspect the p-value of the test by calling *output\$pval*. The box below provides the R code for each step.

```

library(causalweight)           # load causalweight package
data(JC)                       # load JC data
Z=JC[,1]                        # define instrument (assignment)
D=JC[,37]                       # define treatment (training)
X=JC[,2:29]                     # define covariates
Y=JC[,46]                       # define outcome (health state)
output=identificationDML(y=Y, d=D, x=X, z=Z) # run identification test
output$pval                      # p-value of the test

```

The test yields a p-value of 0.26 (or 26%), which indicates that the null hypothesis of conditional mean independence between Z and Y given D and X cannot be rejected at conventional levels of statistical significance. For this reason, we do not find compelling statistical evidence for a violation of the selection-on-observables assumption concerning D or the instrument valid-

ity of Z when controlling for X . This suggests that the baseline covariates might be sufficiently rich to control for confounders jointly affecting D and Y .

In a next step, we load the *InvariantCausalPrediction* package developed by Meinshausen (2019) and implement the algorithm of Peters, Bühlmann, and Meinshausen (2015) to identify sets of observed variables that entail invariant estimations of the conditional mean outcome when assuming a linear outcome model. We now consider all variables in the *JC* data that were measured either before or during the first year after the random assignment to Job Corps as candidate treatment variables, namely, columns 2 to 37 of the data set. To this end, we define $X=as.matrix(JC[,c(2:37)])$, where the *as.matrix* command converts our data frame into a numeric matrix, as required by the algorithm. We note that X includes pretreatment covariates as well as training participation and covariates measured in the first year after program assignment. As before, random assignment to Job Corps is used as the instrument (Z), which presumably does not directly affect the health outcome Y , but only indirectly through elements in X (in particular, training participation). We feed X , Y , Z into the *ICP* command to execute the algorithm and set the argument *alpha=0.05* for testing the impact of candidate treatments at the 5% level of significance (including a correction for multiple hypothesis testing). We save the results in an object named *output* and wrap the latter by the *plot* command to plot the results. See the box below for the R code of the various steps.

```

library(InvariantCausalPrediction)      # load InvariantCausalPrediction package
X=as.matrix(JC[,c(2:37)])                # observables (candidate treatments)
output=ICP(X=X,Y=Y,ExpInd=Z,alpha=0.05) # algorithm invariant causal prediction
plot(output)                            # plot results


```

Running the code generates the graph in figure 5.0.2, which conveys information regarding the selection of variables in X as regressors in various outcome regression models. Each of these regression models adheres to the criterion that the estimation of the conditional mean outcome ($E[Y|\tilde{D}]$) based on the respective set of selected regressors (\tilde{D}) is rather invariant across values of the instrument Z . In this graph, the dots correspond to the estimates of the regression coefficients (y-axis) for different variables coming from X (x-axis) for each of the

models that meet the invariance criterion. The slim vertical bars are the 95% confidence intervals of the estimated coefficients. In contrast, the fatter vertical bars correspond to the union (or combination) of the confidence intervals across all models satisfying the invariance condition, i.e., they are derived from the maximum upper and minimum lower bounds of any confidence interval across all relevant models.

Variables without any coefficient estimates and confidence intervals are never selected to be in the set \tilde{D} . Among the variables having nonzero coefficients (and thus, nonzero effects on the health outcome) in models satisfying the invariance criterion are gender, education, health at baseline and after the first year, a missing dummy for smoking behavior at baseline, and training participation in the first year. However, for any of those variables, the union of the confidence intervals includes a zero effect. Nevertheless, we can observe, for instance, that the coefficients on training participation (*trainy1*) are often negative and never positive, suggesting a health-increasing effect (due to the inverse coding of the outcome). As already mentioned, more instruments (particularly those with large first-stage effects on elements in X) may be helpful in reducing estimation uncertainty and thus the length of the confidence intervals.

Figure 5.0.2: Invariant causal prediction



References

- Apfel, N., J. Hatamyar, M. Huber, and J. Kueck (2023). Learning control variables and instruments for causal analysis in observational data. *working paper, University of Fribourg*.
- Breunig, C. and P. Buraue (2021). Testability of reverse causality without exogeneous variation. *arXiv preprint 2107.05936*.
- Chernozhukov, V., D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins (2018). Double/de-biased machine learning for treatment and structural parameters. *The Econometrics Journal* 21, C1–C68.
- Frölich, M. and S. Sperlich (2019). *Impact Evaluation: Treatment Effects and Causal Analysis*. Cambridge University Press.
- Fulcher, I. R., I. Shpitser, S. Marealle, and E. J. Tchetgen Tchetgen (2019). Robust Inference on Population Indirect Causal Effects: The Generalized Front Door Criterion. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 82, 199–214.
- Glymour, C., K. Zhang, and P. Spirtes (2019). Review of causal discovery methods based on graphical models. *Frontiers in Genetics* 10, 1–15.

- Gorbach, T., X. de Luna, J. Karvanen, and I. Waernbaum (2023). Contrasting Identifying Assumptions of Average Causal Effects: Robustness and Semiparametric Efficiency. *Journal of Machine Learning Research* 24, 1–65.
- Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian journal of statistics*, 65–70.
- Hoyer, P., D. Janzing, J. M. Mooij, J. Peters, and B. Schölkopf (2008). Nonlinear causal discovery with additive noise models. *Advances in neural information processing systems* 21.
- Huber, M. and J. Kueck (2023). Testing the identification of causal effects in observational data. *arXiv preprint 2203.15890*.
- Imbens, G. W. (2004, Feb.). Nonparametric estimation of average treatment effects under exogeneity: a review. *The Review of Economics and Statistics* 86, 4–29.
- Imbens, G. W. (2020). Potential outcome and directed acyclic graph approaches to causality: Relevance for empirical practice in economics. *Journal of Economic Literature* 58, 1129–1179.
- Kalisch, M. and P. Bühlmann (2014). Causal structure learning and inference: A selective review. *Quality Technology & Quantitative Management* 11, 3–21.
- Karlsson, R. K. A. and J. H. Krijthe (2023). Detecting hidden confounding in observational data using multiple environments. *arXiv preprint 2205.13935*.
- Mani, S., P. L. Spirtes, and G. F. Cooper (2012). A theoretical study of y structures for causal discovery. *arXiv preprint 1206.6853*.
- Meinshausen, N. (2019). Invariantcausalprediction: Invariant causal prediction. *R package*.
- Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: networks of plausible inference*. San Mateo: Morgan Kaufmann.
- Pearl, J. (1995). Causal diagrams for empirical research. *Biometrika* 82, 669–710 (with discussion).
- Pearl, J. (2000). *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press.
- Peters, J. and P. Bühlmann (2013, 11). Identifiability of Gaussian structural equation models with equal error variances. *Biometrika* 101, 219–228.
- Peters, J., P. Bühlmann, and N. Meinshausen (2015). Causal inference using invariant prediction: identification and confidence intervals. *arXiv preprint 1501.01332*.
- Peters, J., D. Janzing, and B. Schölkopf (2017). *Elements of causal inference: foundations and learning algorithms*. The MIT Press.

- Peters, J., J. M. Mooij, D. Janzing, and B. Schölkopf (2014). Causal discovery with continuous additive noise models. *Journal of Machine Learning Research* 15, 2009–2053.
- Quinzan, F., A. Soleymani, P. Jaillet, C. R. Rojas, and S. Bauer (2023). Drcfs: Doubly robust causal feature selection. In *International Conference on Machine Learning*, pp. 28468–28491. PMLR.
- Reichenbach, H. (1991). *The direction of time*, Volume 65. Univ of California Press.
- Scutari, M. (2010). Learning bayesian networks with the bnlearn r package. *Journal of Statistical Software* 035(i03).
- Shimizu, S., P. O. Hoyer, A. Hyvärinen, A. Kerminen, and M. Jordan (2006). A linear non-gaussian acyclic model for causal discovery. *Journal of Machine Learning Research* 7.
- Soleymani, A., A. Raj, S. Bauer, B. Schölkopf, and M. Besserve (2022). Causal feature selection via orthogonal search. *Transactions on Machine Learning Research*.
- Spirtes, P. and C. Glymour (1991). An algorithm for fast recovery of sparse causal graphs. *Social science computer review* 9, 62–72.
- Spirtes, P., C. N. Glymour, and R. Scheines (2000). *Causation, prediction, and search*. MIT press.
- Textor, J., B. van der Zander, M. S. Gilthorpe, M. Liskiewicz, and G. T. Ellison (2017, 01). Robust causal inference using directed acyclic graphs: the R package ‘dagitty’. *International Journal of Epidemiology* 45, 1887–1894.
- Verma, T. and J. Pearl (1990). Equivalence and synthesis of causal models. In *Proceedings of the Sixth Annual Conference on Uncertainty in Artificial Intelligence*, pp. 255–270.
- Verma, T. and J. Pearl (1992). An algorithm for deciding if a set of observed independencies has a causal explanation. In *Uncertainty in artificial intelligence*, pp. 323–330. Elsevier.
- Zhang, K. and L.-W. Chan (2006). Extensions of ica for causality discovery in the hong kong stock market. In *International Conference on Neural Information Processing*, pp. 400–409. Springer.
- Zhang, K. and A. Hyvärinen (2009a). Causality discovery with additive disturbances: An information-theoretical perspective. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2009, Bled, Slovenia, September 7-11, 2009, Proceedings, Part II 20*, pp. 570–585. Springer.
- Zhang, K. and A. Hyvärinen (2009b). On the identifiability of the post-nonlinear causal model. In *25th Conference on Uncertainty in Artificial Intelligence (UAI 2009)*, pp. 647–655. AUAI Press.