

A Weaker Faithfulness Assumption based on Triple Interactions

Alexander Marx^{1,2}

Arthur Gretton³

Joris M. Mooij⁴

¹CISPA Helmholtz Center for Information Security, Saarland University, Saarbrücken, Germany

²Max Planck Institute for Informatics, Saarland University, Saarbrücken, Germany

³Gatsby Unit, University College London, London, United Kingdom

⁴Korteweg-de Vries Institute, University of Amsterdam, Amsterdam, The Netherlands

Abstract

One of the core assumptions in causal discovery is the faithfulness assumption—i.e. assuming that independencies found in the data are due to separations in the true causal graph. This assumption can, however, be violated in many ways, including xor connections, deterministic functions or cancelling paths. In this work, we propose a weaker assumption that we call 2-adjacency faithfulness. In contrast to adjacency faithfulness, which assumes that there is no conditional independence between each pair of variables that are connected in the causal graph, we only require no conditional independence between a node and a subset of its Markov blanket that can contain up to two nodes. Equivalently, we adapt orientation faithfulness to this setting. We further propose a sound orientation rule for causal discovery that applies under weaker assumptions. As a proof of concept, we derive a modified Grow and Shrink algorithm that recovers the Markov blanket of a target node and prove its correctness under strictly weaker assumptions than the standard faithfulness assumption.

1 INTRODUCTION

In this work, we focus on causal discovery from observational data, where we are given a sample from the joint distribution P of the observed variables and try to infer the true causal graph G between them. Two standard assumptions in this field are the causal Markov condition and the faithfulness assumption (Spirtes et al., 2000). While the causal Markov condition assumes that all separations in the true causal graph G imply independencies in P , the faithfulness assumption is its counterpart. That is, all independencies found in P are due to separations in G . Although both assumptions have great merit for causal discovery algorithms,




Figure 1: Failures of adjacency faithfulness: Assume in graph (a) X, Z are fair independent coins and $Y := (X \oplus Z) \oplus E$, where \oplus is the xor operator and E is a biased coin denoting a noise term. Then X is independent of Y (denoted as $X \perp_P Y$) and $Z \perp_P Y$. Graph (b) could correspond with a linear model where both directed paths from X to Y cancel out such that $X \perp_P Y$, but $X \not\perp_P Z$ and $Z \not\perp_P Y$.

especially the faithfulness assumption has been criticized in the past (Andersen, 2013; Zhang and Spirtes, 2016).

Despite it was proven that faithfulness violations in causally sufficient linear-Gaussian and discrete acyclic systems occur with Lebesgue measure zero (Meek, 1995b), it has also been shown that on a finite sample, empirical faithfulness violations do appear surprisingly often (Uhler et al., 2013). Even on population level, there exist simple generating mechanisms, as shown in Figure 1, that violate the faithfulness assumption. For instance, two independent random variables X and Z , that can be modelled by fair coins, together cause Y through a noisy xor relation. As a consequence, all three variables are marginally independent. Following the faithfulness assumption, there should be no edges connecting X, Y and Z in the causal graph—however, there are.

Faithfulness violations like the above have been intensively studied in the past (Ramsey et al., 2006; Zhang and Spirtes, 2008; Spirtes and Zhang, 2014) and several weaker assumptions such as adjacency faithfulness (Spirtes et al., 2000), P-minimality (Pearl, 2009), SGS-minimality (Spirtes et al., 2000) and frugality (Forster et al., 2017), which we review in Section 3.3, have been proposed. Although faithfulness violations induced by xor-type relations—i.e. both parents

are marginally independent of the child node—can be detected by most of the above approaches, they do not analyze under which conditions the DAG structure can be partially recovered, once such violations have been detected.

In this work, we propose a new assumption that we call *2-adjacency faithfulness*, which allows us to both detect such faithfulness violations and partially infer the underlying DAG structure under certain conditions. We start by explaining the standard concepts and notation in Section 2 and review failures of faithfulness as well as related work in Section 3. Then, we study the causal structure of xor-type connections in Section 4 and propose 2-adjacency faithfulness in Section 5. To partially infer causal DAGs that may contain such generating mechanisms, we introduce a sound orientation rule, in Section 6. Further, we show under which assumptions on the distribution this rule is applicable—which we formalize as *2-orientation faithfulness*—and analyze its failure cases. As a proof of concept, we provide a modified Grow and Shrink (GS) algorithm (Margaritis and Thrun, 2000) in Section 7 and show it correctly identifies the Markov blanket of a target node under strictly weaker assumptions than faithfulness. Besides, we give some intuition on how to extend well-known causal discovery algorithms based on our new assumptions.

2 DAGS AND INDEPENDENCE

In this section, we define the notation and provide definitions for separations on graphs and independence.

2.1 CAUSAL GRAPHS

A causal *directed acyclic graph* (DAG) G over a set of random variables V with joint distribution P is defined such that each pair of nodes that is adjacent in G is causally related. For simplicity, we will use the random variables V to also refer to the nodes of the graph. A directed edge $X \rightarrow Y$ in G between two nodes representing the random variables $X, Y \in V$ indicates that X is a *direct cause* or *parent* of Y and that Y is a *direct effect* or *child* of X . Accordingly, we denote the set of all parents of $X \in V$ with $\text{Pa}(X)$, the set of all children with $\text{Ch}(X)$ and the set of parents and children with $\text{PC}(X) := \text{Pa}(X) \cup \text{Ch}(X)$. Further, we write $\text{An}(X)$ for the set of *ancestors* and denote its *descendants* by $\text{De}(X)$, where X is an ancestor and descendant of itself. Respectively, we refer to the *non-descendants* of X as $\text{Nd}(X) := V \setminus \text{De}(X)$. Last, the *Markov blanket* of X is defined as $\text{MB}(X) := \text{PC}(X) \cup \text{Sp}(X)$, where $\text{Sp}(X)$ are the spouses of X , that is, nodes that share a child node with X . Importantly, X is *d-separated* of any other node in the graph given its Markov blanket and $\text{MB}(X)$ is the smallest such set.

DAGs are used to represent causal graphs under the assumption of acyclicity, no selection bias, and *causal sufficiency*,

that is, it is assumed that no two variables $X, Y \in V$ are caused by an unobserved confounder $Z \notin V$. This is also the setup on which we focus in this paper—i.e. assuming that all relevant variables are observed, that there are no causal cycles and that there has been no conditioning on selection variables. Further, as a short form to summarize a model as defined above, we write $\mathcal{M} = (G, V, P)$.

2.2 INDEPENDENCE AND SEPARATION

In the following, we define conditional independence in a probability distribution and *d-separation* in a graph.

Given three sets of random variables $X, Y, Z \subseteq V$, where P is the joint distribution over V , we denote that X is *probabilistically independent* of Y given Z in P as $X \perp_P Y \mid Z$.

D-separation (Pearl, 2009) is defined in terms of paths. A *path* p between X and Y , denoted $p = \langle X, \dots, Y \rangle$, is a sequence of distinct nodes X_1, \dots, X_n such that X_i is adjacent to X_{i+1} for $i = 1, \dots, n - 1$, $X_1 = X$ and $X_n = Y$. Further, we call a node C a *collider* on a path $\langle \dots, X, C, Y, \dots \rangle$, where C is adjacent to both X and Y , if two arrowheads point to it, that is $X \rightarrow C \leftarrow Y$.

Definition 1 (*d-Separation*) A path between two vertices X, Y in a DAG is *d-connecting* given a set Z , if

1. every non-collider on the path is not in Z , and
2. every collider on the path is an ancestor of Z .

If there is no path *d-connecting* X and Y given Z , then X and Y are *d-separated* given Z . Sets X and Y are *d-separated* given Z , if for every pair X, Y , with $X \in X$ and $Y \in Y$, X and Y are *d-separated* given Z .

As shorthand notation for separations on a DAG G , we write $X \perp_G Y \mid Z$ if X is *d-separated* from Y given Z . A useful set of tools for inferences on graphs and distributions are the graphoid axioms (Dawid, 1979; Spohn, 1980; Geiger et al., 1990). Since we use those axioms in our proofs, we provide them in Supplementary Material S.3.

We round up this section by defining the causal Markov condition (Spirtes et al., 2000) (CMC) for DAGs.

Definition 2 (Causal Markov Condition) Given the triple $\mathcal{M} = (G, V, P)$, the causal Markov condition holds, if every *d-separation* in G implies an independence in P .

The causal Markov condition is one of the most essential assumptions for causal discovery algorithms. On the other hand, assumptions about what properties of the graph can be inferred based on the given distribution have been weakened over time (Ramsey et al., 2006; Zhang and Spirtes, 2008; Forster et al., 2017). Most commonly known is the faithfulness assumption.

3 ADJACENCY FAITHFULNESS AND WHEN IT IS VIOLATED

To lay out the problem, we first explain faithfulness and adjacency faithfulness, then examine when those could fail and give a summary about the most relevant related approaches that use weaker assumptions.

The faithfulness assumption is one of the core assumptions made by most causal discovery algorithms (Spirtes et al., 2000) and it can be seen as the inverse assumption to CMC—i.e. assuming that all independencies found in P imply a d -separation in the causal graph. Adjacency faithfulness is a slightly weaker assumption.

Definition 3 (Adjacency Faithfulness) Let $\mathcal{M} = (G, V, P)$, if $X, Y \in V$ are adjacent in G , then they are probabilistically dependent given all $S \subseteq V \setminus \{X, Y\}$.

Alternatively, we could turn this definition around by stating that if we find a conditional independence in P , then we assume that there is no edge in the corresponding graph. Assuming adjacency faithfulness ensures that we recover the correct skeleton graph (i.e. the undirected graph). Correct detection of the skeleton together with the correct identification of all collider structures ensures that the detected graph is in the Markov equivalence class of the true graph (Verma and Pearl, 1991). The latter is ensured by additionally assuming orientation faithfulness (Zhang and Spirtes, 2008).

Definition 4 (Orientation-Faithfulness) Let $\mathcal{M} = (G, V, P)$ and let $\langle X, Y, Z \rangle$ be an unshielded path¹ in G .

1. If $X \rightarrow Y \leftarrow Z$, then X and Z are dependent given any subset in $V \setminus \{X, Z\}$ that contains Y ; otherwise
2. X and Z are dependent conditional on any subset of $V \setminus \{X, Z\}$ that does not contain Y .

The bottleneck here is the adjacency faithfulness assumption, as many causal discovery algorithms such as PC (Spirtes et al., 2000) or GES (Chickering, 2002) rely on finding adjacent nodes either by checking for marginal dependencies or adding single edges based on adjacency faithfulness and CMC. However, adjacency faithfulness can be violated in many ways, e.g. by xor-type connections, path cancellations, or deterministic relations. We briefly explain the first two below, as they are relevant for the remainder. For deterministic relations and finite sample failures, we refer to Lemeire et al. (2012).

3.1 XOR-TYPE RELATIONS

In this work, we focus on xor-type relations. That is, given a triple of nodes $X, Y, Z \in V$ s.t. $X \rightarrow Y \leftarrow Z$, where at least

¹For an unshielded path $\langle X, Y, Z \rangle$, X is adjacent to Y and Y is adjacent to Z , but X is not adjacent to Z .




Figure 2: Sample data for the collider graph $X \rightarrow Y \leftarrow Z$, where $X, Z \sim N(0, 1)$ are iid and $Y := \text{sign}(XZ) \cdot E$, with $E \sim \text{Exp}(\frac{1}{\sqrt{2}})$. The dependence is only detectable by considering all three variables jointly.

one of the causal edges cannot be detected by a marginal dependence, but only by looking at the joint distribution of X, Y and Z . The key here is that either parent of Y might not be dependent on Y , but by considering both parents, we can detect the dependence. To illustrate this, consider the following example where we describe a noisy xor with an unobserved noise variable modelled by a biased coin as it is common for binary causal structures (Inazumi et al., 2011).

Example 1 Let $\mathcal{M} = (G, V, P)$ be a causal model. Given variables $X, Y, Z \in V$ s.t. $X \rightarrow Y \leftarrow Z$ in G and there is no edge connecting X and Z , as in Figure 1(a), where X, Z are fair independent coins. Their common effect Y is generated as $Y := (X \oplus Z) \oplus E$, where \oplus denotes xor—i.e. $X \oplus Z := (X + Z) \bmod 2$ —and E is a biased coin with $P(E = 1) = p$, where $0 \leq p < \frac{1}{2}$ and $E \perp_P \{X, Z\}$. Hence, $X \not\perp_G Y$, $Z \not\perp_G Y$, however, due to the xor, we have that $X \perp_P Y$ and $Z \perp_P Y$. Both edges violate adjacency faithfulness. If we were to check the joint distribution, we can find that $Y \not\perp_P \{X, Z\}$, or $X \not\perp_P Z \mid Y$, since we get that $P(X = 1, Z = 1, Y = 1) = \frac{p}{4}$, where $P(X = 1, Z = 1) \cdot P(Y = 1) = \frac{1}{4} \cdot \frac{1}{2} = \frac{1}{8}$. Those terms are only equal if $p = \frac{1}{2}$, which we excluded by assumption.

Similar examples, where marginal dependencies might be hard to detect, can also be constructed for continuous data (Sejdinovic et al., 2013)—e.g. if X, Z are normally distributed with mean zero and variance one, and $Y := \text{sign}(XZ) \cdot E$, with exponentially distributed noise $E \sim \text{Exp}(\frac{1}{\sqrt{2}})$ (see Figure 2). The authors demonstrate that with a high probability no marginal dependence between X, Z as well as X, Y can be detected using a kernel dependence measure such as the Hilbert-Schmidt Independence Criterion (Gretton et al., 2005). They do, however, reliably detect a dependence between all three variables. Further,

Marx et al. (2021) show that dependencies generated by the above mechanism and generalizations of it to mixed-type data can be detected by a broad range of dependence measures such as kernel tests and tests based on conditional mutual information. Besides iid data, the work of Sejdinovic et al. (2013) has been extended to time series data (Rubenstein et al., 2016). In this setting, Rubenstein et al. (2016) showcase the efficiency of their approach on a Forex data set, which contains strong triple interactions.

Motivated by these positive results, we focus on the theoretical foundations that allow us to detect such triples in a causal setting and analyze under which conditions we can identify the collider.

3.2 CANCELLING PATHS

A minimal example of cancelling paths was given by Hesselow (1976) and is illustrated with the causal graph shown in Figure 1(b). In Hesselow’s example taking birth control pills (X) can influence the risk of getting thrombosis (Y) via two paths. It has a direct effect and also taking the pills reduces the chance of pregnancy (Z), which itself is a cause of thrombosis. However, the causal effects induced by those paths cancel such that $X \perp_P Y$. As an example mechanism that causes such a cancellation, consider a linear Gaussian system in which $Z := \alpha X$, $Y := \beta Z - \gamma X$ and $\gamma = \alpha\beta$. This failure of faithfulness was shown to be undetectable since X will be dependent on Y given Z and hence the graph $X \rightarrow Z \leftarrow Y$ is also a valid graph for those independencies—i.e. Markov equivalent (Zhang and Spirtes, 2008). There exist cancelling paths that consist of more than three variables, which are detectable, e.g. if Z is not adjacent to Y , but there is a path $Z \rightarrow W \rightarrow Y$ (Zhang and Spirtes, 2008).

3.3 WEAKER ASSUMPTIONS

In the following, we discuss different approaches on how to relax the faithfulness assumption.

Two well-studied assumptions are P-minimality (Pearl, 2009) and SGS-minimality (Spirtes et al., 2000). While the former states that from all DAGs that satisfy the causal Markov condition w.r.t. P , the DAG that entails most conditional independence statements is preferred. The latter assumes that no proper subgraph of the true DAG fulfils the causal Markov condition w.r.t. to P . From both assumptions, SGS-minimality is the weaker assumption (Zhang, 2013). In a different line of research, it was shown that SGS-minimality suffices for causal discovery approaches based on the additive noise assumption (Peters et al., 2014).

A more recent approach by Forster et al. (2017) introduces the concept of *frugality*, which is a stronger assumption than both minimality assumptions. The authors define a DAG G to be more frugal than G' , if G contains fewer edges than

G' . A maximally frugal DAG uses only as many edges as are necessary to satisfy the causal Markov condition. To determine maximally frugal graphs, one has to consider all causal orderings of the variables, which is rather costly, but can be solved using permutation algorithms (Raskutti and Uhler, 2018). Another approach to discover causal graphs based on frugality, or any of the above assumptions is based on SAT solvers (Zhalama et al., 2017).

In this paper, we introduce 2-adjacency faithfulness, which allows us to find xor-type relations, some faithfulness violations induced by cancelling paths and all relations that are detectable by assuming adjacency faithfulness. We conjecture that 2-adjacency faithfulness in combination with some minimality assumption, e.g. SGS minimality (see Section 7.1), is a slightly stronger assumption than frugality since frugality considers all permutations (Forster et al., 2017). Hence, frugality might be able to also detect structures for which it is necessary to observe more than three nodes to find a dependence. However, how often such structures occur in real data is unknown. Thus, resorting to only consider all triple structures might be more efficient than having to check all permutations. In addition, we extend existing work by providing a sound orientation rule that can be used to infer the edges within a 2-association, if they appear in a larger graph.

Next, we discuss xor-type relations in more detail. We use those structures as an example to illustrate one of the main properties of 2-associations, that we describe in Theorem 1.

4 UNFAITHFUL TRIPLES

We first define what we call an unfaithful triple² and its properties, and then argue why such a triple a) violates adjacency faithfulness and b) even if detected, the underlying DAG structure cannot be uniquely determined without further information.

Definition 5 (Unfaithful Triple) *Given $M = (G, V, P)$ and three distinct nodes $X, Y, Z \in V$: if X, Y and Z are marginally independent but not mutually independent in P , we call $\{X, Y, Z\}$ an unfaithful triple w.r.t. P .³ If further for each distinct pair of nodes $A, B \in \{X, Y, Z\}$:*

$$\forall S \subseteq V \setminus \{X, Y, Z\} : A \not\perp_P B \mid S \cup \{X, Y, Z\} \setminus \{A, B\},$$

we call $\{X, Y, Z\}$ a minimal unfaithful triple.

The first example for such a triple for three binary random variables was given by Bernstein (1927), which is equivalent

²Ramsey et al. (2006) used the term unfaithful triple for the non-detectable faithfulness violation explained in Section 3.2.

³Not mutually independent implies that $X \not\perp_P \{Y, Z\}$, $Y \not\perp_P \{X, Z\}$ or $Z \not\perp_P \{X, Y\}$.




Figure 3: Assume that $\{X, Y, Z\}$ form an unfaithful triple. Since X is d -separated from Y given U and Z , they do not form a minimal unfaithful triple. Neither do $\{U, Y, Z\}$, since Y can be d -separated from U given $\{W, Z\}$. Thus, only $\{U, W, Z\}$ can be a minimal unfaithful triple.

to Example 1. The minimality condition ensures that the three nodes are connected by a path of length two, as we will show below. This concept is also illustrated in Figure 3.

We start by showing that if three random variables $\{X, Y, Z\}$ are marginally independent, finding a dependence between all three variables, e.g. $X \not\perp_P \{Y, Z\}$, implies that also $Y \not\perp_P \{X, Z\}$ and $Z \not\perp_P \{X, Y\}$.

Lemma 1 Given $\mathcal{M} = (G, \mathbf{V}, P)$, let $\{X, Y, Z\} \subseteq \mathbf{V}$ form an unfaithful triple in P , then $X \not\perp_P \{Y, Z\}$, $Y \not\perp_P \{X, Z\}$ and $Z \not\perp_P \{X, Y\}$, which in addition implies that $X \not\perp_P Y \mid Z$, $X \not\perp_P Z \mid Y$ and $Y \not\perp_P Z \mid X$.

We provide the proofs for Lemma 1 and the following Lemma 2 and Theorem 1 in Supplementary Material S.3.

Consider Example 1. Since X, Y, Z form an unfaithful triple, we can infer from Lemma 1 that each pair is conditionally dependent given the third node. As there are no other nodes in the graph, X, Y, Z must form a minimal unfaithful triple.

Next, we show that (minimal) unfaithful triples must be connected in the causal graph.

Lemma 2 Given $\mathcal{M} = (G, \mathbf{V}, P)$, let $\{X, Y, Z\} \subseteq \mathbf{V}$ form an unfaithful triple in P . If CMC holds, each node in the triple is d -connected to at least one other node in the triple by a path in G .

Further, we show that a minimal unfaithful triple has to contain a collider on a path of length two that connects all three nodes in the triple, e.g. $X \rightarrow Y \leftarrow Z$. To do that, we first show a more general statement.

Theorem 1 Given $\mathcal{M} = (G, \mathbf{V}, P)$ with distinct $X, Y, Z \in \mathbf{V}$ and assume that CMC holds. If $\forall \mathbf{S} \subseteq \mathbf{V} \setminus \{X, Y, Z\}$ it holds that $X \not\perp_P Y \mid Z \cup \mathbf{S}$, $X \not\perp_P Z \mid Y \cup \mathbf{S}$ and $Y \not\perp_P Z \mid X \cup \mathbf{S}$, then one of the three nodes is a collider on a path of length two between the two other nodes, e.g. $X \rightarrow Y \leftarrow Z$ in G .

The theorem only states that there exists a collider, e.g. $X \rightarrow Y \leftarrow Z$, but not whether this path is shielded or not.

Since we do not assume any marginal dependence or independence in Theorem 1, we can derive that the same statement holds for a minimal unfaithful triple. Notice that for a minimal unfaithful triple each pair of nodes is marginally independent, which implies that there is no way to decide which of the three possible collider structures corresponds with the causal graph in the absence of further information.

Knowing that a minimal unfaithful triple has to contain a collider in G , it is obvious that such a structure violates adjacency faithfulness, as none of the edges is represented by a marginal dependence in P . The key point is that we can detect such interactions by taking multiple parents into account. In the following, we define a weaker assumption that allows us to detect and infer causal graphs that contain such faithfulness violations.

5 2-ADJACENCY FAITHFULNESS

To define our new assumption, we first need to define associations between a single node and a set of nodes.

Definition 6 (k -Association) Let P be the joint distribution of a set of observed random variables \mathbf{V} .

1. Given distinct $X, Y \in \mathbf{V}$, we say that X is 1-associated to Y , if $\forall \mathbf{S} \subseteq \mathbf{V} \setminus \{X, Y\} : X \not\perp_P Y \mid \mathbf{S}$.
2. Given distinct $X, Y_1, Y_2 \in \mathbf{V}$, X is 2-associated to $\{Y_1, Y_2\}$ if $\forall \mathbf{S} \subseteq \mathbf{V} \setminus \{X, Y_1, Y_2\}$ it holds that
 - i) $X \not\perp_P Y_1 \mid \mathbf{S} \cup Y_2$,
 - ii) $X \not\perp_P Y_2 \mid \mathbf{S} \cup Y_1$ and
 - iii) $Y_1 \not\perp_P Y_2 \mid \mathbf{S} \cup X$.

We call X strictly 2-associated to $\{Y_1, Y_2\}$, if X is 2-associated to $\{Y_1, Y_2\}$ and not 1-associated to Y_1 or Y_2 .

In other words, k -associations relate to two types of dependencies: certain conditional dependencies between pairs of variables (1-associations) and between triples (2-associations). For readability, we use a shorthand notation and write $X \text{--}_2 \{Y, Z\}$ if X is 2-associated to Y and Z resp. $X \text{--}_1 Y$ if X is 1-associated to Y . We denote a strict 2-association by “ $\overset{s}{\text{--}}_2$ ”. If we refer to a set \mathbf{Y} that contains at most two elements and we want to express that X is either 1- or 2-associated to this set, we write $X \text{--}_{\leq 2} \mathbf{Y}$. Similarly, we write $X \overset{s}{\text{--}}_{\leq 2} \mathbf{Y}$, if X is 1- or strictly 2-associated to \mathbf{Y} .

Pairwise dependencies can occur for example in a simple chain $X \rightarrow Y \rightarrow Z$, where no adjacency failure occurs. In this case, $X \text{--}_1 Y$ and $Y \text{--}_1 Z$. Triple interactions that match the definition of 2-associations, however, need to have a specific structure. As we saw in Theorem 1, 2-associations always contain a collider. Thus, a chain graph or a common cause structure does not induce a 2-association. On the other hand, the minimum unfaithful triple in Example 1 matches the definition, since $X \overset{s}{\text{--}}_2 \{Y, Z\}$, $Y \overset{s}{\text{--}}_2 \{X, Z\}$ and

$Z \overset{s}{\perp}_2 \{X, Y\}$. In general, strict 2-associations describe collider structures such as $X \rightarrow Y \leftarrow Z$ for which at least one of the edges violates adjacency faithfulness. If faithfulness holds, a collider structure induces a 2-association, but not a strict 2-association. We use this intuition for our new assumption.

Definition 7 (2-Adjacency Faithfulness) Given $M = (G, V, P)$, for all $X, Y \in V$, where X and Y are adjacent in G , there exists $Y \subseteq MB(X)$, with $Y \in \mathbf{Y}$, s.t. $X \overset{s}{\perp}_{\leq 2} Y$.

The main idea here is to weaken adjacency faithfulness such that if a marginal dependence is not present, i.e., adjacency faithfulness is violated, there will be a dependence in combination with a parent, child or spouse. If adjacency faithfulness is not violated, we will not find any strict 2-associations and our assumption reduces to adjacency faithfulness. By also considering strict 2-associations, however, we can discover a larger spectrum of causal mechanisms.

The textbook example for a mechanism that violates faithfulness but is detectable by assuming 2-adjacency faithfulness is the xor-connection described in Example 1. Here, $Y \overset{s}{\perp}_2 \{X, Z\}$, two parents, while $X \overset{s}{\perp}_2 \{Y, Z\}$ —i.e. a child and a spouse. We could even slightly adapt the mechanism and only model Z using an unbiased coin but use a biased coin for X . In this case, only X is marginally independent of Y and Z , while Z becomes dependent on Y .

Moreover, assuming 2-adjacency faithfulness could even allow us to detect some faithfulness violations that are due to cancelling paths. In particular, consider the two paths $X \rightarrow Y$ and $X \rightarrow Z \rightarrow W \rightarrow Y$ that cancel such that $X \perp_P Y$. Since $X \perp_P Y$, $X \perp_P W \mid Z$, X is not 1-associated to Y or W and hence could be strictly 2-associated to the set $\{W, Y\}$. Since we know that a 2-association contains a collider and we can neither find a 1-association to Y or W , we know that there has to be an edge violating adjacency faithfulness.

It is not possible to rely on orientation faithfulness when dealing with strict 2-associations. Although we know that a strict 2-association has to contain a collider, we do not know the skeleton structure within the triple and hence cannot apply orientation faithfulness.

6 ORIENTATION

So far, we showed how we can detect unfaithful triples from conditional (in)dependence statements under the weaker assumption of 2-adjacency faithfulness. Now imagine that we want to use this knowledge for causal discovery. If we observe an isolated triple that follows the dependence structure of the noisy xor, we can only tell that there is a collider. However, if we are given more information, we are able to break this symmetry.




Figure 4: In both distributions $Y \overset{s}{\perp}_2 \{X, Z\}$ and $Y \perp_1 W$. In the graph shown in (a) Y is a collider on all paths between $\{X, Z\}$ and W , whereas in (b) Y is a non-collider.

Example 2 Consider that X and Z are unbiased coins as in the noisy xor example. In addition, there is a binary variable W with $P(W = 1) = p$, where $0 < p < 1$ and an unobserved binary noise variable E with $P(E = 1) = q$, where $0 < q < \frac{1}{2}$. Now we generate Y as

$$Y := ((X \oplus Z) \wedge W) \oplus E,$$

where E, W, X and Z are drawn independently. The requirements for p ensure that W is dependent on Y and the requirements on E ensure that the dependencies are non-deterministic ($q \neq 0$) and evident without observing E ($q \neq \frac{1}{2}$). The corresponding causal graph is given in Figure 4(a). From the induced dependencies, that we derive in detail in Supplementary Material S.1, we can now obtain an asymmetry. In particular, $\{X, Y, Z\}$ form a minimal unfaithful triple, but only Y is dependent on W , whereas $\{X, Z\} \perp_P W$ and due to the xor, $X \perp_P W \mid Y$ as well as $Z \perp_P W \mid Y$. Thus, we can detect that there is no edge between X and W or Z and W since none of these pairs can be 2-associated. However, we do find that $X \not\perp_P W \mid \{Y, Z\}$ and $Z \not\perp_P W \mid \{Y, X\}$. As we will show in Theorem 2, we can use this information to identify that Y is the collider in the triple and that $W \rightarrow Y$.

To detect such an asymmetry, it is necessary that the collider in the triple is the effect of another node or pair of nodes. If, for example, X would be the collider in the triple and $W \rightarrow Y$ (see Figure 4(b)), we cannot find such an asymmetry. To generate that graph we could model Y as a noisy copy of W and construct X with a noisy xor from Y and Z . We still know that W is adjacent to Y , but we cannot direct any of the edges as for example we would find that $X \not\perp_P W \mid Z$, which we would also observe if Z would be the collider in the triple, or if we would flip the edge direction between Y and W —i.e. if W is a noisy copy of Y .

Based on this intuition, we propose an orientation rule that may include causal structures that induce strict 2-associations. To do so, we use a shorthand notation—i.e. write $Y \rightarrow X$, if for each element $Y \in \mathbf{Y}$ it holds that $Y \rightarrow X$ and vice versa write $X \rightarrow Y$ if X is a parent of each node $Y \in \mathbf{Y}$, that is, $\forall Y \in \mathbf{Y} : X \rightarrow Y$.

Definition 8 (Orientation Rule) Let $M := (G, V, P)$ and we are given two disjoint sets $X, Z \subseteq V$ and $Y \in V$, where

$Y \overset{s}{\perp\!\!\!\perp} X$ and $Y \overset{s}{\perp\!\!\!\perp} Z$, and no node $X \in \mathbf{X}$ is adjacent to some node $Z \in \mathbf{Z}$.

- i) If for each pair $X \in \mathbf{X}$ and $Z \in \mathbf{Z}$ it holds that X is dependent on Z given any subset of $V \setminus \{X, Z\}$ that contains $Y \cup (X \setminus \{X\}) \cup (Z \setminus \{Z\})$, then $X \rightarrow Y \leftarrow Z$,
- ii) otherwise, if for each pair $X \in \mathbf{X}$ and $Z \in \mathbf{Z}$ it holds that X is dependent on Z conditional on any subset of $V \setminus \{X, Z\}$ that contains $(X \setminus \{X\}) \cup (Z \setminus \{Z\})$ but does not contain Y , Y is a non-collider on at least one path $\langle X, Y, Z \rangle$ where $X \in \mathbf{X}$ and $Z \in \mathbf{Z}$.

Simply put, the above orientation rule relies on the fact that a (strict) 2-association contains a collider. Either Y is the collider on each path $\langle X, Y, Z \rangle$ between any variable $X \in \mathbf{X}$ and $Z \in \mathbf{Z}$ or Y is one of the parents in at least one of the triples and hence blocks at least one such path. If both sets \mathbf{X} and \mathbf{Z} only contain a single element, rule i) refers to a “normal” collider e.g. $X \rightarrow Y \leftarrow Z$ and rule ii) refers either to a chain like $X \rightarrow Y \rightarrow Z$ or to a common cause $X \leftarrow Y \rightarrow Z$. Let us consider Example 2 again, where we generated Y as a non-deterministic function of X, Z and W . First, we find that $Y \overset{s}{\perp\!\!\!\perp} \{X, Z\}$, $Y \dashv\!\!\!\perp W$ and W is not adjacent to X or Z (since W is not 1- or strictly 2-associated to X or Z), which is required to apply our rule. Further, we can apply rule i) since W is dependent on X given any set that includes $\{Y, Z\}$ and W is dependent on Z given any set that includes $\{Y, X\}$. Hence, we can infer the correct DAG structure $\{X, Z\} \rightarrow Y \leftarrow W$.

In the following we will first show that our orientation rule is sound—i.e. if rule i) or ii) can be applied, then we are sure we found the corresponding graph structure—and then analyze the inverse, that is, what assumptions need to hold s.t. the given graph implies the suggested dependence model.

Theorem 2 *Assuming that the causal Markov condition holds, the orientation rule in Definition 8 is sound.*

We provide the proof in Supplementary Material S.3.

The question that remains is: Does the inverse always hold? For example, if the true graph contains a non-collider structure such as $X \rightarrow Y \rightarrow Z$, will we always find that $X \not\perp\!\!\!\perp_p Z$? The short answer is no. Already when we only assume adjacency faithfulness, it can happen that $X \perp\!\!\!\perp_p Z$ although the true graph is $X \rightarrow Y \rightarrow Z$ and it holds that $X \not\perp\!\!\!\perp_p Y$ and $Y \not\perp\!\!\!\perp_p Z$, which is called failure of transitivity. More generally, assuming that orientation faithfulness holds, such failures will not occur. In the following, we extend this assumption to our setting.

Definition 9 (2-Orientation Faithfulness) *Let $M := (G, V, P)$ and we are given two disjoint sets $X, Z \subseteq V$ and $Y \in V$, where $Y \overset{s}{\perp\!\!\!\perp} X$ and $Y \overset{s}{\perp\!\!\!\perp} Z$, and no node $X \in \mathbf{X}$ is adjacent to some node $Z \in \mathbf{Z}$.*




Figure 5: In both figures, $Y \overset{s}{\perp\!\!\!\perp} X = \{X_1, X_2\}$, $Y \overset{s}{\perp\!\!\!\perp} Z = \{Z_1, Z_2\}$ (related nodes and edges are marked in black) and $X_2 \overset{s}{\perp\!\!\!\perp} \{U, Z_2\}$. Given only this information, we cannot tell if the path $\langle X_2, Y, Z_2 \rangle$ is unshielded (a) or shielded (b). While in graph (a), we could safely apply our orientation rule, the shielded graph (b) can be problematic. Due to the directed path from X_1 over U to Z_2 and the adjacency between X_2 and Z_2 , each pair $X, Z \in \mathbf{X} \times \mathbf{Z}$ is now d -connected given $\{Y\} \cup (X \setminus \{X\}) \cup (Z \setminus \{Z\})$. Thus, the condition for rule i) could hold, although $X \rightarrow Y \leftarrow Z$ is not in G .

- i) If $X \rightarrow Y \leftarrow Z$ is in G , then for each pair $X \in \mathbf{X}$ and $Z \in \mathbf{Z}$, X is dependent on Z given any subset of $V \setminus \{X, Z\}$ that contains $Y \cup (X \setminus \{X\}) \cup (Z \setminus \{Z\})$,
- ii) otherwise, for each pair $X \in \mathbf{X}$ and $Z \in \mathbf{Z}$, X is dependent on Z conditional on any subset of $V \setminus \{X, Z\}$ that contains $(X \setminus \{X\}) \cup (Z \setminus \{Z\})$, but not Y .

Equivalently to 2-adjacency faithfulness, 2-orientation faithfulness reduces to orientation faithfulness, if both sets \mathbf{X} and \mathbf{Z} only contain a single element. For orientation faithfulness, it has been shown that all failures can be detected under the assumption that adjacency faithfulness holds (Zhang and Spirtes, 2008). Sadly, an equally strong statement cannot be made for 2-adjacency faithfulness and 2-orientation faithfulness, as we discuss below.

6.1 FAILURES OF 2-ORIENTATION FAITHFULNESS

Without any assumptions, we can detect triples for which $Y \overset{s}{\perp\!\!\!\perp} X$ and $Y \overset{s}{\perp\!\!\!\perp} Z$, and know by assuming CMC that all 2-associations contain a collider. If further, all paths $\langle X, Y, Z \rangle$ with $(X, Z) \in \mathbf{X} \times \mathbf{Z}$ are unshielded, we can detect if any of the conditions in 2-orientation faithfulness fails. In particular, due to the soundness of our orientation rule, we would detect that none of the conditions in the orientation rule is satisfied if condition i) or ii) in 2-orientation faithfulness fails, as we show in Corollary 1.

Yet, we cannot detect all failures of 2-orientation faithfulness. That is, due to the fact that we might not always be able to detect whether all paths $\langle X, Y, Z \rangle$ are unshielded. If there is a direct edge between X and Z , we will always find that those nodes are either 1-associated or there exists

a third node U such that they are strictly 2-associated (if 2-adjacency faithfulness holds). However, if we find a strict 2-association between X and $\{Z, U\}$ there is no guarantee that the path is shielded. In particular, if U is the collider between X and Z , the triple is unshielded; but if Z is the collider between X and U , the triple is shielded (see Figure 5, in which X refers to X_2 and Z to Z_2). In a causal discovery algorithm, we could try to iteratively infer the DAG structure within such triples until we cannot apply the rule anymore. If we are lucky, we can first infer that $X \rightarrow U \leftarrow Z$ and after that also apply our rule for $\{X, Y, Z\}$. Keeping this exception in mind, we can derive the following corollary from Theorem 2.

Corollary 1 *Given $M := (G, \mathbf{V}, P)$ with $Y \in \mathbf{V}$ and $X, Z \subseteq \mathbf{V}$, where $X \cap Z = \emptyset$, $Y \overset{s}{\perp\!\!\!\perp} X$, $Y \overset{s}{\perp\!\!\!\perp} Z$ and no pair of nodes $(X, Z) \in X \times Z$ is adjacent. Assuming that CMC holds, we can detect if condition i) or ii) of 2-orientation faithfulness fails on the triple $\{X, Y, Z\}$.*

The proof is provided in Supplementary Material S.3. In general, 2-orientation faithfulness might be useful not only for constraint-based causal discovery methods, but also for algorithms that aim to discover the Markov blanket of a node or permutation-based discovery algorithms such as the Sparsest Permutation (SP) algorithm proposed by Raskutti and Uhler (2018). In Supplementary Material S.2, we provide a short discussion from which we conjecture that the SP algorithm can identify the collider pattern for strict 2-associations as in Figure 4, if 2-orientation faithfulness holds.

7 IMPLEMENTATION

As a proof of concept, we propose a simple modification of the Grow and Shrink (GS) algorithm (Margaritis and Thrun, 2000) to discover Markov blankets that may contain strict 2-associations. After that, we briefly discuss further challenges that need to be solved to propose a causal discovery algorithm based on our new assumptions.

The GS algorithm is a simple and theoretically sound causal discovery algorithm, that as a first step identifies the Markov blanket for each node (Margaritis and Thrun, 2000). This step of the algorithm consists of a grow phase, in which we iteratively discover a superset of the Markov blanket of a target node T , and a shrink phase, in which superfluous nodes are pruned.

To make sure that we can detect Markov blankets that contain strict 2-associations, we assume that 2-adjacency faithfulness holds and that we can detect all spouses. For the latter, there are two options. Either the target node, the spouse and the common child are connected via a strict 2-association, or the spouse node is only 1- or strictly 2-associated with the common child and not with the target node. For the first option, it suffices to assume 2-adjacency

faithfulness to detect the spouse, whereas for the second option we need to assume a variant of 2-orientation faithfulness. In particular, consider the graph $T \rightarrow C \leftarrow S$, in which $T \perp_1 Y$ and $C \perp_1 Y$. Assuming only 2-adjacency faithfulness will not guarantee that $T \not\perp\!\!\!\perp S \mid C$, additionally assuming 2-orientation faithfulness will. More generally, we need to assume that condition i) in 2-orientation faithfulness also holds for shielded triples, which boils down to assuming that the spouses of the target do not cancel each other out, as we explain below.

Assumption 1 *Let $M := (G, \mathbf{V}, P)$ and we are given two disjoint sets $X, Z \subseteq \mathbf{V}$ and $Y \in \mathbf{V}$, where $Y \overset{s}{\perp\!\!\!\perp} X$ and $Y \overset{s}{\perp\!\!\!\perp} Z$. If $X \rightarrow Y \leftarrow Z$ in G , then for each pair $X \in \mathbf{X}$ and $Z \in \mathbf{Z}$, X is dependent on Z given any subset of $\mathbf{V} \setminus \{X, Z\}$ that contains $Y \cup (X \setminus \{X\}) \cup (Z \setminus \{Z\})$.*

The above assumption is a relatively lightweight adaption of condition i) in 2-orientation faithfulness. In particular, let $X = \{X, T\}$, where T is the target node. Then all nodes in Z are spouses of T , which we can detect if Assumption 1 holds and even become part of $\text{PC}(T)$ if all paths $\langle T, Y, Z \rangle$ for $Z \in \mathbf{Z}$ are shielded. Thus, we would already add those nodes when looking for the parents and children of T . The only complication that may arise is if the second node $X \in \mathbf{X}$ is adjacent to a node in $Z \in \mathbf{Z}$ and this adjacency would lead to a cancellation such that Z is only dependent on T if we do not condition on X . The corresponding causal graph consists of the paths $T \rightarrow Y \leftarrow Z$ and $Y \leftarrow X \rightarrow Z$. Since X cannot block the path $\langle T, Y, Z \rangle$, such a scenario seems only possible if the causal mechanism that generates Z from X is deterministic. Based on this assumption, we introduce our adapted GS algorithm.

The generalized GS algorithm is shown in Algorithm 1, where we only modified the grow phase to also consider pairs of random variables. This allows us to find nodes to which the target node is strictly 2-associated or spouses to which a child node of T is strictly 2-associated using Assumption 1. The shrink phase is not modified and checks if singletons can be removed. Importantly, we will not remove single nodes of a true strict 2-association to T or a child of T , because we do not check for marginal dependencies. For example, assume that $T \overset{s}{\perp\!\!\!\perp} \{X, Z\}$ and both nodes were added in the grow phase, where X is a child of T and Z the corresponding spouse. If we try to remove X in the shrink phase, we find that $T \not\perp\!\!\!\perp_P X \mid S \setminus X$, since $Z \in S$. Hence, X remains in S , as well as Z .

In the following, we show that our proposed algorithm correctly identifies the Markov blanket of a target node assuming that 2-adjacency faithfulness, the causal Markov condition and Assumption 1 hold.

Theorem 3 *Given $M = (G, \mathbf{V}, P)$. Assuming that 2-adjacency faithfulness, Assumption 1 and CMC hold, Algorithm 1 correctly identifies $\text{MB}(T)$ for $T \in \mathbf{V}$.*

Algorithm 1: Modified GS for Markov Blankets

input : Random variables V with joint distribution P ,
Target $T \in V$

output : $MB(T)$

```
1  $V' \leftarrow V \setminus \{T\}$ ;  
2  $S \leftarrow \emptyset$ ;  
  // Grow Phase  
3 while  $(\exists X \in V' : T \not\perp_P X \mid S) \vee$   
4  $(\exists X, Z \in V' : T \not\perp_P X \mid S \cup \{Z\})$  do  
5    $S \leftarrow S \cup \{X\}$  ;  
  // Shrink Phase  
6 while  $\exists X \in S : T \perp_P X \mid S \setminus X$  do  
7    $S \leftarrow S \setminus X$  ;  
8 return  $S$ 
```

We provide the proof in Supplementary Material S.3. For discovering the Markov blanket, we do not need to know the collider of a strict 2-association since it only returns a set of nodes. The more challenging task is to implement our framework to discover causal networks, which we will briefly discuss below.

7.1 TOWARDS CAUSAL NETWORK INFERENCE

In this paper, we mainly focused on answering two questions: 1) How can we weaken faithfulness to detect xor-type structures, and 2) under which conditions can we identify the collider in such a triple? The assumptions that we derived are sufficient for Markov blanket discovery, as we showed above. However, in Markov blanket discovery, we are not concerned with finding a unique graph structure, our only goal is to detect the set of nodes, which contains the parents, children and spouses of a target node.

For causal discovery, assuming 2-adjacency faithfulness might be too inclusive. To illustrate this statement, consider Example 1 again, in which X and Z cause Y through a noisy xor. Besides the three possible collider structures $X \rightarrow Y \leftarrow Z$, $Y \rightarrow X \leftarrow Z$ and $X \rightarrow Z \leftarrow Y$, also a fully connected graph, e.g. $X \rightarrow Y \leftarrow Z$ and $X \rightarrow Z$, is compatible with 2-adjacency faithfulness since X is not 1-associated to Z or Y . To avoid finding such graphs with superfluous edges, we need to restrict the search space to those graphs G' for which no proper subgraph is compatible with 2-adjacency faithfulness and 2-orientation faithfulness. For future work, we want to investigate how this can be achieved. One possibility could be to combine our assumptions with SGS-minimality (Spirtes et al., 2000), which assumes that no proper subgraph of the true DAG G entails the causal Markov condition w.r.t. to P . In Example 1, assuming SGS-minimality in addition to 2-adjacency faithfulness would reduce the number of admissible graphs to the three possible collider structures, which include the true DAG.

To derive a causal discovery algorithm under the above assumptions, the most straightforward approach would be to further extend the GS algorithm. After detecting all Markov blankets, the GS algorithm distinguishes the spouses of a Markov blanket from the parents and children by detecting collider patterns. In this step, we could extend the existing rule with a modification of Definition 8. Similarly, we could extend well-known algorithms such as the PC algorithm (Spirtes et al., 2000) or the GES algorithm (Chickering, 2002) by modifying the skeleton phase, respectively the forward phase such that we can find triple interactions as we did for GS. The edge orientation could be done by first applying the orientation rule in Definition 8 and then applying a similar set of rules like Meek’s orientation rules (Meek, 1995a). Alternatively, it was shown that SAT-based causal discovery algorithms can be easily adapted to weaker assumptions than faithfulness (Zhalama et al., 2017), which might be an interesting direction for future work.

8 CONCLUSION

In this work, we proposed 2-adjacency faithfulness, which is a weaker version of adjacency faithfulness. Our new assumption is able to detect faithfulness violations caused by weak or non-existent marginal dependencies, which are detectable by considering a combination of parents, children or spouses. We provide an in-depth analysis of such dependencies and propose a sound orientation rule, which can infer part of the correct causal structure by detecting colliders. We complement this rule with 2-orientation faithfulness, which assumes that if a causal graph contains such collider structures, we will find that the corresponding conditional dependence statements hold in P . As a proof of concept, we extended the GS algorithm to find Markov blankets under strictly weaker assumptions than faithfulness.

For future work, we would like to develop a sound causal discovery algorithm based on 2-adjacency faithfulness and extend our theory to directed mixed graphs.

Acknowledgements

The authors would like to thank the anonymous reviewers for insightful and valuable comments. This work was supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement 639466). A. Marx is supported by the International Max Planck Research School for Computer Science (IMPRS-CS).

References

H. Andersen. When to Expect Violations of Causal Faithfulness and Why it Matters. *Philosophy of Science*, 80

- (5):672–683, 2013.
- S. Bernstein. *Theory of Probability*. Moscow, 1927.
- D. M. Chickering. Optimal Structure Identification With Greedy Search. *Journal of Machine Learning Research*, 3(Nov):507–554, 2002.
- A. P. Dawid. Conditional Independence in Statistical Theory. *Journal of the Royal Statistical Society: Series B (Methodological)*, 41(1):1–15, 1979.
- M. Forster, G. Raskutti, R. Stern, and N. Weinberger. The Frugal Inference of Causal Relations. *The British Journal for the Philosophy of Science*, 69(3):821–848, 2017.
- D. Geiger, T. Verma, and J. Pearl. Identifying Independence in Bayesian Networks. *Networks*, 20(5):507–534, 1990.
- A. Gretton, O. Bousquet, A. Smola, and B. Schölkopf. Measuring statistical dependence with hilbert-schmidt norms. In *International Conference on Algorithmic Learning Theory*, pages 63–77. Springer, 2005.
- G. Hesslow. Two Notes on the Probabilistic Approach to Causality. *Philosophy of science*, 43(2):290–292, 1976.
- T. Inazumi, T. Washio, S. Shimizu, J. Suzuki, A. Yamamoto, and Y. Kawahara. Discovering causal structures in binary exclusive-or skew acyclic models. In *Proceedings of the International Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 373–382, 2011.
- J. Lemeire, S. Meganck, F. Cartella, and T. Liu. Conservative independence-based causal structure learning in absence of adjacency faithfulness. *International Journal of Approximate Reasoning*, 53(9):1305–1325, 2012.
- D. Margaritis and S. Thrun. Bayesian Network Induction via Local Neighborhoods. In *Advances in Neural Information Processing Systems*, pages 505–511, 2000.
- A. Marx, L. Yang, and M. van Leeuwen. Estimating conditional mutual information for discrete-continuous mixtures using multi-dimensional adaptive histograms. In *Proceedings of the SIAM International Conference on Data Mining (SDM)*, pages 387–395, 2021.
- C. Meek. Causal inference and causal explanation with background knowledge. In *Proceedings of the International Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 403–410. Morgan Kaufmann Publishers Inc., 1995a.
- C. Meek. Strong completeness and faithfulness in Bayesian networks. In *Proceedings of the International Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 411–419, 1995b.
- J. Pearl. *Causality*. Cambridge university press, 2009.
- J. Peters, J. M. Mooij, D. Janzing, and B. Schölkopf. Causal discovery with continuous additive noise models. *Journal of Machine Learning Research*, 15(1):2009–2053, 2014.
- J. Ramsey, P. Spirtes, and J. Zhang. Adjacency-Faithfulness and Conservative Causal Inference. In *Proceedings of the International Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 401–408. AUAI Press, 2006.
- G. Raskutti and C. Uhler. Learning directed acyclic graph models based on sparsest permutations. *Stat*, 7(1), 2018.
- P. K. Rubenstein, K. P. Chwialkowski, and A. Gretton. A kernel test for three-variable interactions with random processes. In *Proceedings of the International Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 637–646, 2016.
- D. Sejdinovic, A. Gretton, and W. Bergsma. A Kernel Test for Three-Variable Interactions. In *Advances in Neural Information Processing Systems*, pages 1124–1132, 2013.
- P. Spirtes and J. Zhang. A Uniformly Consistent Estimator of Causal Effects under the k -Triangle-Faithfulness Assumption. *Statistical Science*, 29(4):662–678, 2014.
- P. Spirtes, C. N. Glymour, R. Scheines, D. Heckerman, C. Meek, G. Cooper, and T. Richardson. *Causation, Prediction, and Search*. MIT press, 2000.
- W. Spohn. Stochastic Independence, Causal Independence, and Shieldability. *Journal of Philosophical logic*, 9(1): 73–99, 1980.
- 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. Verma and J. Pearl. Equivalence and synthesis of causal models. In *Proceedings of the International Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 220–227, 1991.
- Zhalama, J. Zhang, F. Eberhardt, and W. Mayer. SAT-Based Causal Discovery under Weaker Assumptions. In *Proceedings of the International Conference on Uncertainty in Artificial Intelligence (UAI)*. AUAI, 2017.
- J. Zhang. A Comparison of Three Occam’s Razors for Markovian Causal Models. *The British journal for the philosophy of science*, 64(2):423–448, 2013.
- J. Zhang and P. Spirtes. Detection of Unfaithfulness and Robust Causal Inference. *Minds and Machines*, 18(2): 239–271, 2008.
- J. Zhang and P. Spirtes. The three faces of faithfulness. *Synthese*, 193(4):1011–1027, Apr 2016.