

Toward Falsifying Causal Graphs Using a Permutation-Based Test

Elias Eulig^{1,2,*,†}, Atalanti A. Mastakouri³, Patrick Blöbaum³,
Michaela Hardt^{4,*}, Dominik Janzing³

¹German Cancer Research Center (DKFZ)

²Heidelberg University

³Amazon Research Tübingen

⁴University Hospital Tübingen

elias.eulig@dkfz.de, {atalanti,janzind}@amazon.de, bloebp@amazon.com, michaela.hardt@uni-tuebingen.de

Abstract

Understanding causal relationships among the variables of a system is paramount to explain and control its behavior. For many real-world systems, however, the true causal graph is not readily available and one must resort to predictions made by algorithms or domain experts. Therefore, metrics that quantitatively assess the goodness of a causal graph provide helpful checks before using it in downstream tasks. Existing metrics provide an *absolute* number of inconsistencies between the graph and the observed data, and without a baseline, practitioners are left to answer the hard question of how many such inconsistencies are acceptable or expected. Here, we propose a novel consistency metric by constructing a baseline through node permutations. By comparing the number of inconsistencies with those on the baseline, we derive an interpretable metric that captures whether the graph is significantly better than random. Evaluating on both simulated and real data sets from various domains, including biology and cloud monitoring, we demonstrate that the true graph is not falsified by our metric, whereas the wrong graphs given by a hypothetical user are likely to be falsified.

Project — <https://eulig.github.io/dag-falsification>

1 Introduction

Directed Acyclic Graphs (DAGs) are a core concept of causal reasoning as they form the basis of structural causal models (SCMs) which can be used to predict the effect of interventions in a causal system or even to answer counterfactual questions. For that reason, they have numerous applications, including biology (Sachs et al. 2005; Huber et al. 2007; Pingault et al. 2018), medicine (Shrier and Platt 2008) and computer vision (Wang and Sahbi 2013; Chalupka, Perona, and Eberhardt 2015; Wang et al. 2020).

Nevertheless, for many real-world systems the true causal relationships, represented in the form of a DAG, are often not readily available. If randomized controlled trials are not possible, inferring the DAG from passive observational data alone is a hard problem that rests on strong assumptions on statistical properties (e.g. causal faithfulness), functional relationships, noise distributions, or graph constraints

*Work done at Amazon Research Tübingen.

†Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

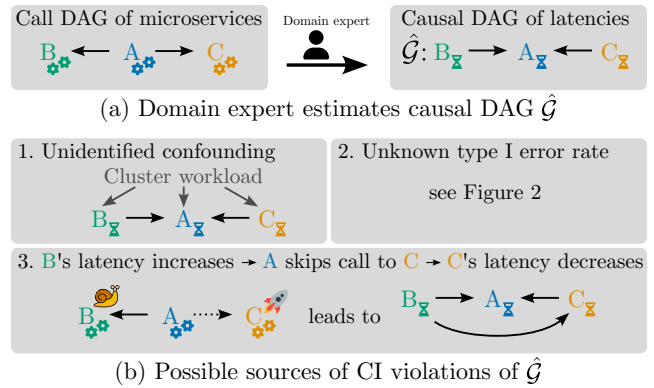


Figure 1: (a) For monitoring micro-service architectures, domain experts may invert the call graph to obtain a causal graph of latencies and error rates (Budhathoki et al. 2022); (b) However, there are multiple reasons why this graph might violate independence statements on observed data.

(Shimizu et al. 2006; Peters and Bühlmann 2014; Chickering 1996; Chickering, Heckerman, and Meek 2004; Claassen, Mooij, and Heskes 2013; Mooij et al. 2016). These assumptions are often violated in practice. Instead of relying on discovery from observational data, domain experts can describe known dependencies in a system (Fig. 1; (a)). This approach, however, is subject to human error and incomplete knowledge. This is particularly problematic for applications in causal inference, where a wrong graph structure will lead to wrong conclusions about the effect of interventions.

Efforts towards quantitatively evaluating the consistency of a given graph (either originating from a domain expert or a causal discovery algorithm) using observational data alone are of utmost importance. However, existing attempts towards this direction report the raw number of inconsistencies between DAG and observed data without a baseline of how many such inconsistencies are to be expected in the first place (Textor et al. 2016; Reynolds et al. 2022). Striving for zero is unrealistic and the acceptable number of violations depends on many factors including the size and complexity of the graph. On real-world data, the fraction of conditional independence violations for expert-elicited graphs




Figure 2: Type I error rate at $\alpha = 5\%$ for different sizes D of the conditioning set for one parametric (partial correlation) and two nonparametric CI tests, KCI (Zhang et al. 2011) and GCM (Shah and Peters 2020). Data (solid: $N = 100$, dashed: $N = 500$) were sampled from gaussian-linear conditionals. More details are given in Supp. A.1.

can be surprisingly high due to unidentified confounding, high type I error rates of the conditional independence tests (Shah and Peters 2020), or complex interactions between the variables (Fig. 1; (b)). Other methods require the existence and knowledge of DAGs from a related system (Pitchforth and Mengersen 2013), or are similarly difficult to interpret due to a missing baseline (Madigan, York, and Allard 1995; Friedman and Koller 2003).

To overcome these drawbacks, in this work we develop a novel metric¹ to evaluate a given graph using observational data alone. Aiming to understand if a number of violations of a user-specified graph is high or low, we compare it against a baseline that we construct by randomly permuting its nodes. Through this comparison we can shed light onto the question if the violations of a user-specified graph are false-positives or point to real deficiencies of the graph.

2 Related Work

In this section, we review existing works that aim to quantify the consistency of a given DAG with observed data.

The R package *dagitty* (Textor et al. 2016) implements functions to evaluate DAG-dataset consistency by testing conditional independence (CI) relations implied by the graph structure via the d-separation (Pearl 1988) and global Markov condition. However, directly using the number of violations of graph-implied CIs as a metric to evaluate a given graph is not suitable for real-world applications because there exists no nonparametric CI test with valid level² over all distributions and thus the probability of type I errors remains unknown in practice (Shah and Peters 2020, Th. 2), c.f. Fig. 2. Therefore, without a baseline comparison, the raw number of violations (absolute or fraction) does not provide

¹An implementation of our metric is available in the Python package *DoWhy* (Blöbaum et al. 2024).

²Valid test would mean that for some CI test and sample size N , any distribution from the null is rejected at most with probability of the prespecified significance level α (Shah and Peters 2020).

the user with a meaningful measure of whether or not the observed inconsistencies of the given graph are significant.

Reynolds et al. (2022) validate a given DAG (6 nodes, 8 edges) that relates exposure to spaceflight environment to the performance and health outcomes of rats and mice. Together with testing CIs implied by d-separations using *dagitty*, they test whether dependencies implied by the graph lead to marginal dependencies in the observations (faithfulness). Similar to (Textor et al. 2016) interpreting the results of those tests is challenging, without a baseline comparison. Our metric overcomes this drawback by providing such a baseline comparison to estimate whether the observed inconsistencies are significant.

Pitchforth and Mengersen (2013) suggest a number of questions to validate expert-elicited Bayesian Networks (BNs). In particular, the authors propose to validate a given BN by verifying that it is similar to BNs from the same domain already established in the literature. Nevertheless, the framework does not provide a quantified measure and many of the questions assume the existence of comparison models in the literature, which may not be available in many domains. In contrast, we propose a quantitative metric, constructed via a surrogate baseline and not reliant on the existence of similar models from the same domain.

Finally, numerous works exist on Bayesian structure learning (e.g., Madigan, York, and Allard 1995; Friedman and Koller 2003; Giudici and Castelo 2003) where the posterior $p(\mathcal{G} \mid \mathcal{D}) \propto p(\mathcal{D} \mid \mathcal{G}) p(\mathcal{G})$ of a graph \mathcal{G} is estimated given some data \mathcal{D} . While the likelihood $p(\mathcal{D} \mid \mathcal{G})$ could be used as a measure for the goodness of \mathcal{G} , the problem of a missing baseline remains.

3 Background

3.1 Notation

This work aims to evaluate the consistency of a given DAG $\hat{\mathcal{G}} = (\mathbf{V}, \hat{\mathcal{E}})$ with vertices $\mathbf{V} = \{1, \dots, n\}$, random variables $\mathbf{X} = \{X_i : i \in \mathbf{V}\}$, and edges $\hat{\mathcal{E}} \subseteq \mathbf{V}^2$ using observations $\mathcal{D} = \{\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(N)}\}$ i.i.d. sampled from the joint distribution³ $P(\mathbf{X})$. We assume the existence of some unknown true causal DAG \mathcal{G}^* and that $P(\mathbf{X})$ satisfies the causal Markov condition relative to \mathcal{G}^* .

In the following, we refer to $X \perp\!\!\!\perp_{\mathcal{D}} Y \mid Z$ as the outcome of a CI test using \mathcal{D} , where X and Y are conditionally independent given Z (and likewise $X \not\perp\!\!\!\perp_{\mathcal{D}} Y \mid Z$ to denote conditional dependency). Note that whether independence is rejected is both a property of \mathcal{D} and the choice of a particular CI test. In this work, to denote that a set of nodes Z d-separates (Pearl 1988) X from Y , we write: $X \perp\!\!\!\perp_{\mathcal{G}} Y \mid Z$.

For some graph $\mathcal{G} = (\mathbf{V}, \mathcal{E})$, we call a node i a parent of node j if $(i, j) \in \mathcal{E}$ and denote with $\text{Pa}_j^{\mathcal{G}}$ the set of all parents of $j \in \mathcal{G}$ ($\text{Pa}_j^{\mathcal{G}} = \emptyset$ iff j is a root node). Furthermore, for nodes i, j if there exists a direct path $i \rightarrow \dots \rightarrow j$ we call i an ancestor of node j and j a descendant of i . We denote with $\text{Anc}_j^{\mathcal{G}}$ the set of all ancestors of $j \in \mathcal{G}$ and with $\text{ND}_i^{\mathcal{G}}$ the set of all non-descendants of $i \in \mathcal{G}$.

³In the following, we will assume that this P has a density (w.r.t. a product measure).

3.2 Validating Local Markov Conditions

One common approach to evaluate a DAG using observed data is by means of statistical testing of independence relations implied by $\hat{\mathcal{G}}$ (e.g. Textor et al. 2016; Reynolds et al. 2022). In the following, we will formally introduce this test and elaborate on why it is unsuitable to use as a metric directly. In Sec. 4 we will then present a baseline to overcome said drawbacks.

One of the standard assumptions in causal inference is the causal Markov condition, which allows us to factorize a joint probability distribution $P(\mathbf{X})$ over the variables \mathbf{X} of \mathcal{G} into the product

$$P(\mathbf{X}) = P(X_1, X_2, \dots, X_n) = \prod_{i=1}^n P(X_i | \text{Pa}_i^{\mathcal{G}}). \quad (1)$$

We can equivalently formulate this as the parental Markov condition:

Theorem 1 (Parental Markov condition (Pearl 2009)). *A probability distribution P is Markov relative to a DAG $\mathcal{G} = (\mathbf{V}, \mathcal{E})$, iff $X_i \perp\!\!\!\perp_P ND_i^{\mathcal{G}} \setminus \text{Pa}_i^{\mathcal{G}} \mid \text{Pa}_i^{\mathcal{G}}$.*

For the remainder of this work we will use the terms parental Markov condition and local Markov condition (LMC) interchangeably. To test whether a distribution satisfies the parental Markov condition relative to a DAG we can thus list the CIs entailed by the DAG via the graphical d-separation criterion (Pearl 1988) and test whether those are satisfied or violated by the distribution at hand.

Definition 1 (Parental triples). For some graph \mathcal{G} , we refer to the ordered triple $(i, j \in ND_i^{\mathcal{G}} \setminus \text{Pa}_i^{\mathcal{G}}, Z = \text{Pa}_i^{\mathcal{G}})$ as parental triple and denote the set of all such triples implied by \mathcal{G} as $T_{\text{Pa}}^{\mathcal{G}}$.

According to Theorem 1 every parental triple $(i, j, Z) \in T_{\text{Pa}}^{\mathcal{G}}$ implies the CI $X_i \perp\!\!\!\perp_P X_j \mid Z$. Similar to Textor et al. (2016) we will now introduce the notion of violations of LMC:

Definition 2 (Violations of LMCs). We denote with $V_{\text{LMC}}^{\hat{\mathcal{G}}, \mathcal{D}}$ the set of triples $(i, j, Z) \in T_{\text{Pa}}^{\hat{\mathcal{G}}}$, for which (i, j) is an ordered pair, and we observe LMC violations on data \mathcal{D} , i.e.

$$V_{\text{LMC}}^{\hat{\mathcal{G}}, \mathcal{D}} = \left\{ (i, j, Z) \in T_{\text{Pa}}^{\hat{\mathcal{G}}} : X_i \not\perp\!\!\!\perp_{\mathcal{D}} X_j \mid Z \right\}. \quad (2)$$

Furthermore, we denote the fraction of LMC violations with

$$\phi_{\text{LMC}}^{\hat{\mathcal{G}}, \mathcal{D}} = \frac{|V_{\text{LMC}}^{\hat{\mathcal{G}}, \mathcal{D}}|}{|T_{\text{Pa}}^{\hat{\mathcal{G}}}|}. \quad (3)$$

Using the two metrics from Definition 2 directly to measure the goodness of a user-given graph is not suitable for real-world applications due to the reasons detailed in Sec. 2, particularly the unknown type I error rate of a CI test (c.f. Shah and Peters (2020) & Fig. 2).

4 A Baseline for Violations of Local Markov Conditions

As discussed in the previous sections, the metrics $V_{\text{LMC}}^{\hat{\mathcal{G}}, \mathcal{D}}$ and $\phi_{\text{LMC}}^{\hat{\mathcal{G}}, \mathcal{D}}$ are insufficient to measure the consistency of a graph. In the following, we therefore derive a baseline to compare the number of LMC violations $V_{\text{LMC}}^{\hat{\mathcal{G}}, \mathcal{D}}$ to.

4.1 Finding a Suitable Baseline

Motivated by a very skeptical view whether the user specified graph is related to the observed independence structure at all, we are interested in a baseline that is a *random* draw of a set of conditional independence statements. For example, consider the dependence structure of microservices in a distributed system, where multiple effects can lead to violations of independence statements on observed data (Fig. 1).

In general, there can be different reasons why the pattern of observed conditional independences appears unrelated to $\hat{\mathcal{G}}$. On the one hand, the domain expert who provided $\hat{\mathcal{G}}$ may have messed up causal links and directions entirely. But even if all the links of $\hat{\mathcal{G}}$ are correct, additional confounding and violations of faithfulness⁴ can mess up the independence structure. In both cases, DAG and independences appear random *relative to each other* regardless of whether we think the DAG or the pattern of independences to be random. Our experiments show that ‘better than random’ is a surprisingly high bar, and both domain experts and causal discovery algorithms fail to meet it in many settings.

We now introduce a set of properties that draws from the random baseline should satisfy:

- P1: They should infer the same number m of conditional independences as the given graph $\hat{\mathcal{G}}$.
- P2: They should partition m into m_r conditional independences with conditioning sets of size r with the same numbers m_r as $\hat{\mathcal{G}}$ does.
- P3: Conditional independences inferred by draws from the baseline should be closed under the semi-graphoid axioms (Pearl and Paz 1986; Geiger, Verma, and Pearl 1990).

P1 ensures that the number of observed LMC violations on the baseline are comparable to those observed on $\hat{\mathcal{G}}$. E.g. a random baseline inferring fewer CIs than $\hat{\mathcal{G}}$ will also result in much fewer LMC violations.⁵ If $\hat{\mathcal{G}}$ is sparse because it has been drawn with the implicit intention of explaining only *highly* significant dependences, we should not reject it just because it shows many LMC violations with respect to *our* significance level. Instead, we should rather benchmark $\hat{\mathcal{G}}$ against a random guess that infers equally many independences. P2 ensures that type I and type II errors, when testing implied CIs on the baseline, are comparable to those on $\hat{\mathcal{G}}$.

⁴Note, that while faithfulness violations cannot lead to CI violations for $\hat{\mathcal{G}}$, it can make a random set of CI statements equally data consistent.

⁵This is analogous to a classifier (predicting CIs), that classifies only few samples as positive, hence exhibiting few false positives (LMC violations).

E.g. a random baseline inferring mostly CIs with small conditioning sets may exhibit fewer LMC violations than $\hat{\mathcal{G}}$ by virtue of a smaller type I error rate (c.f. Fig. 2). Lastly, P3 ensures that the CI statements from the baseline do not imply additional CIs (which are not already in the set) via the semi-graphoid axioms. For example, due to decomposition, $X \perp\!\!\!\perp Y \cup W | Z \Rightarrow X \perp\!\!\!\perp Y | Z \ \& \ X \perp\!\!\!\perp W | Z$.

A natural choice for a baseline that satisfies all the requirements above can be constructed by sampling node-permutations of $\hat{\mathcal{G}}$. More formally, let S_n denote the set of permutations π on the vertices $\{1, \dots, n\}$ of some graph \mathcal{G} . For any permutation $\pi \in S_n$ we denote with $\sigma_\pi(\mathcal{G})$ the graph for which the edge $i \rightarrow j$ exists iff $\pi(i) \rightarrow \pi(j)$ exists in \mathcal{G} . Because of the one-to-one correspondence between π and σ we will drop the subscript and in the following refer to $\sigma \in S_{\mathcal{G}}$ as one node-permutation of \mathcal{G} .

We can then construct our baseline by sampling random node-permutations $\sigma \in S_{\hat{\mathcal{G}}}$ of the given graph $\hat{\mathcal{G}}$. Let $O(\hat{\mathcal{G}})$ define the orbit of $\hat{\mathcal{G}}$ under $S_{\hat{\mathcal{G}}}$, i.e., the set of DAGs obtained via permutations. Since all DAGs in $O(\hat{\mathcal{G}})$ imply the same CIs as $\hat{\mathcal{G}}$ up to renaming of variables, they satisfy P1 and P2. Furthermore, they satisfy P3 since the set of CIs entailed by a graph are closed under the semi-graphoid axioms.^{6,7}

Note that the mapping $S_{\hat{\mathcal{G}}} \rightarrow O(\hat{\mathcal{G}})$ defined by $\sigma \mapsto \sigma(\hat{\mathcal{G}})$ is in general not one-to-one because there will often be a non-trivial subgroup that leaves $\hat{\mathcal{G}}$ invariant (the stabilizer subgroup $\text{Stab}(\hat{\mathcal{G}})$ of $S_{\hat{\mathcal{G}}}$).⁸

Proposition 1. *Uniform sampling of permutations from the set of all node permutations $S_{\hat{\mathcal{G}}}$ results in uniform sampling from the DAGs in the orbit $O(\hat{\mathcal{G}})$.*

A proof is provided in Supp. D.1. Further, $O(\hat{\mathcal{G}})$ decomposes into Markov equivalence classes of equal size. This can easily be seen by the same argument when we consider the action of $S_{\hat{\mathcal{G}}}$ on the set of Markov equivalence classes and introduce the corresponding (larger) stabilizer subgroup. Thus, uniform sampling of permutations from the set of all node permutations $S_{\hat{\mathcal{G}}}$ also results in uniform sampling from the Markov equivalence classes in $O(\hat{\mathcal{G}})$.

4.2 A Permutation Test to Evaluate DAGs

Using the above baseline we state the following null hypothesis:

Hypothesis H_0 : The DAG $\hat{\mathcal{G}}$ is drawn uniformly at random from some distribution Q on the set of DAGs that is invariant under permutations of nodes, that is $Q(\mathcal{G}) = Q(\sigma(\mathcal{G}))$ for all $\sigma \in S_{\mathcal{G}}$.

To test this hypothesis we consider the number of violations $|V_{\text{LMC}}^{\hat{\mathcal{G}}, \mathcal{D}}|$ as test statistics and build the null via

⁶This is because faithful distributions exist, which satisfy only CIs corresponding to d-separation.

⁷This choice of baseline comes with another feature, namely the implicit testing of violations of faithfulness (c.f. Supp. E).

⁸E.g. the permutation (2, 1, 3) leaves the graph $X_1 \rightarrow X_3 \leftarrow X_2$ invariant.

$|V_{\text{LMC}}^{\sigma(\hat{\mathcal{G}}), \mathcal{D}}|$:

$$p_{\text{LMC}}^{\hat{\mathcal{G}}, \mathcal{D}} = \Pr \left(|V_{\text{LMC}}^{\sigma(\hat{\mathcal{G}}), \mathcal{D}}| \leq |V_{\text{LMC}}^{\hat{\mathcal{G}}, \mathcal{D}}| \right). \quad (4)$$

Proposition 2. *p_{LMC} is a valid p-value, i.e. if H_0 is true, then $\Pr \left(p_{\text{LMC}}^{\hat{\mathcal{G}}, \mathcal{D}} \leq \alpha \right) \leq \alpha$.*

A proof is provided in Supp. D.2. Computing the quantity in (4) using all $n!$ permutations is infeasible for large n . Therefore, we approximate it via Monte Carlo sampling with T random permutations: $\{\sigma_i \sim S_{\hat{\mathcal{G}}}\}_{i=0}^T$. For an estimated p-value we can also report binomial proportion confidence intervals.

Another quantity that proves to be useful in practice is the fraction of DAGs in $O(\hat{\mathcal{G}})$ that are Markov equivalent to $\hat{\mathcal{G}}$. To this end we first define

$$V_{\text{TPa}}^{\mathcal{G}'} = \{(X_i, X_j, Z) \in \mathbf{T}_{\text{Pa}}^{\mathcal{G}'} : X_i \not\perp\!\!\!\perp X_j | Z\} \quad (5)$$

as the set of all ordered triples that are parentally d-separated in \mathcal{G}' but not d-separated in \mathcal{G} .

Proposition 3. *Suppose some given graph $\hat{\mathcal{G}}$. If $V_{\text{TPa}}^{\sigma(\hat{\mathcal{G}}), \hat{\mathcal{G}}} = \emptyset$, then $\sigma(\hat{\mathcal{G}})$ and $\hat{\mathcal{G}}$ are Markov equivalent.*

The proof is provided in Supp. D.3. Using this graphical criterion we can define a second metric

$$\begin{aligned} p_{\text{TPa}}^{\hat{\mathcal{G}}} &:= \Pr \left(|V_{\text{TPa}}^{\sigma(\hat{\mathcal{G}}), \hat{\mathcal{G}}} \leq |V_{\text{TPa}}^{\hat{\mathcal{G}}, \hat{\mathcal{G}}}| \right) \\ &= \Pr \left(|V_{\text{TPa}}^{\sigma(\hat{\mathcal{G}}), \hat{\mathcal{G}}} = 0 \right), \end{aligned} \quad (6)$$

which can be used to measure how informative the CI structure of $\hat{\mathcal{G}}$ is about the possible causal orderings. If $p_{\text{TPa}}^{\hat{\mathcal{G}}} > \alpha$, for some prespecified threshold α , then the number of Markov equivalent DAGs in $O(\hat{\mathcal{G}})$ is large and consequently $\hat{\mathcal{G}}$ provides us with limited information about the true graph in the sense of testing of CIs. For an information-theoretic interpretation of our test see Supp. F.

4.3 Interpretation of p_{LMC} and p_{TPa}

We note that p-values are often misinterpreted and misused in practice (Vidgen and Yasseri 2016; Wasserstein and Lazar 2016; Amrhein, Greenland, and McShane 2019) and therefore provide the following interpretation based on Popper's theory of falsification (Popper 2005). According to this, every scientific theory must admit potential falsifiers, i.e. measurable observations that would falsify the theory. Increasing confidence in such theory can then only come from observations that it permits, i.e. as it withstands attempts to falsify it.

The test $p_{\text{TPa}}^{\hat{\mathcal{G}}}$ provides us with a measure of the falsifiability of a given graph $\hat{\mathcal{G}}$. If the number of random DAGs that are Markov equivalent to $\hat{\mathcal{G}}$ is large (and consequently $p_{\text{TPa}}^{\hat{\mathcal{G}}}$ is large) this limits the falsifiability of $\hat{\mathcal{G}}$ via CI testing. Contrary, if $p_{\text{TPa}}^{\hat{\mathcal{G}}}$ is small, the CIs entailed by $\hat{\mathcal{G}}$ are 'characteristic' and it admits many potential falsifiers. Based on these considerations, we propose the following interpretation of our tests for practitioners:

(a) If $p_{\text{TPa}}^{\hat{\mathcal{G}}} \leq \alpha$, $\hat{\mathcal{G}}$ is falsifiable by testing implied CIs.

- (b) If (a) and further $p_{\text{LMC}}^{\hat{\mathcal{G}}, \mathcal{D}} > \alpha$, then $\hat{\mathcal{G}}$ is falsified.
- (c) If (a) and further $p_{\text{LMC}}^{\hat{\mathcal{G}}, \mathcal{D}} \leq \alpha$, then there is no CI-based evidence against $\hat{\mathcal{G}}$ and we cannot falsify $\hat{\mathcal{G}}$ using our test. Consequently, $\hat{\mathcal{G}}$ is corroborated (but not verified).

4.4 Relating p_{LMC} to the Identification of Cause-Effect Pairs

We now want to present an interpretation of p_{LMC} which is more closely related to *causal* questions. To this end, consider the following task.

Task 1 (Identification of unconfounded cause-effect pairs). Identify all ordered pairs (X_i, X_j) such that there is a directed path from X_i to X_j and no X_k with $k \neq i, j$ that has a directed path to X_i and a directed path to X_j that does not go through X_i .

We can then ask whether $\hat{\mathcal{G}}$ performed better at Task 1 than a random baseline. However, in analogy to P1 above, the baseline should satisfy the following property:

P4: Draws from the random baseline should infer the same number of unconfounded cause-effect pairs as the given graph $\hat{\mathcal{G}}$.

Likewise to P1 & P2, all DAGs in $O(\hat{\mathcal{G}})$ satisfy P4 since they preserve $\hat{\mathcal{G}}$ up to a renaming of variables. By defining the number of wrongly inferred causal effects as test statistics and building the null via the node-permutation baseline, we can compute a p-value p_{CE} .

We will now see that under the assumption of faithfulness and knowledge of a single non-effect node, our metric p_{LMC} is loosely related to p_{CE} .

Theorem 2. *Suppose an additional node X_0 for which we know (e.g. through strong domain knowledge) that it is not an effect of any of the $1, \dots, n$ nodes, but satisfies $X_0 \not\perp\!\!\!\perp X_l, \forall l = 1, \dots, n$. Further, let the joint distribution $P(X_0, X_1, \dots, X_n)$ be Markov and faithful relative to some DAG \mathcal{G} with these $n + 1$ nodes. Then (X_i, X_j) with $i, j = 1, \dots, n$ is an unconfounded cause-effect pair if and only if the following two conditions hold:*

$$X_i \not\perp\!\!\!\perp_P X_j \quad (7)$$

$$X_0 \perp\!\!\!\perp_P X_j | X_i. \quad (8)$$

The proof is provided in Supp. D.4. We conclude that in our idealized scenario, testing whether $\hat{\mathcal{G}}$ is better than random at Task 1, is similar to testing whether $\hat{\mathcal{G}}$ performs better than random in identifying all pairs (X_i, X_j) for which both conditions (7) and (8) hold. Here, the baseline is generated by the permutation group S_{n-1} , that is, the stabilizer group of X_0 in S_n .

5 Experiments

In the following we evaluate our proposed metric on synthetic and real data for which the true DAG is known or a consensus graph is established in the literature. Additionally, we introduce a novel dataset from cloud monitoring where the reversed call graph provides an estimate of the true causal graph (Fig. 1; (a)), thus providing a useful test beyond synthetic and existing real-world datasets.

5.1 Experimental Setup

We conduct experiments with two different sources of given graphs: Emulated domain experts with partial knowledge of either a subset of nodes or a subset of edges of the true DAG $\mathcal{G}^* = (\mathbf{V}, \mathcal{E}^*)$ and causal discovery algorithms, which are a popular choice to estimate a DAG in the absence of domain expertise. Note that all $\hat{\mathcal{G}}$ in our experiments differ from \mathcal{G}^* in a structured, i.e. non-random, way.

Node Domain Expert (DE-V) This model mimics the situation where a domain expert knows all causal edges for some subset $K \subseteq \mathbf{V}$ of the nodes in the system and knows the overall sparsity of the DAG. For all the other nodes, however, the expert has no domain knowledge and thus assigns edges randomly between pairs of nodes not both in K . We define different levels of DE-V to emulate the fraction of nodes for which there exists domain knowledge, i.e. $|K|/|\mathbf{V}|$, where $|K|/|\mathbf{V}| = 0$ corresponds to the situation where the domain expert has no knowledge and $|K|/|\mathbf{V}| = 1$ corresponds to $\hat{\mathcal{G}} = \mathcal{G}^*$. More details are given in Supp. A.2.

Edge Domain Expert (DE-E) This model mimics a domain expert with edge-specific knowledge about \mathcal{G}^* . To construct $\hat{\mathcal{G}}$, we randomly remove and flip some of the true edges and add some new ones. By construction $\hat{\mathcal{G}}$ of a DE-E is related to the Structural Hamming Distance (SHD) (Acid and de Campos 2003; Tsamardinos, Brown, and Aliferis 2006) and thus the desired similarity of a given graph can be controlled by means of the SHD. We characterize different DE-E by the $\text{SHD}(\hat{\mathcal{G}}, \mathcal{G}^*)$ they entail (or $\text{SHD}(\hat{\mathcal{G}}, \mathcal{G}^*)/|\mathcal{E}^*|$ to compare systems with different sparsity), where $\text{SHD}(\hat{\mathcal{G}}, \mathcal{G}^*)/|\mathcal{E}^*| = 0$ corresponds to $\hat{\mathcal{G}} = \mathcal{G}^*$. More details are given in Supp. A.3.

Causal Discovery Algorithms The proposed test can also be applied to DAGs inferred by causal discovery algorithms. However, as many algorithms use (some of the) CIs either explicitly (constraint-based) or implicitly (score-based) for constructing the DAG, evidence in favor of an inferred DAG can only come from those CIs that were neither used by the algorithm, nor implied, via semi-graphoid axioms (Pearl and Paz 1986; Geiger, Verma, and Pearl 1990), by those CIs used by the algorithm. Nonetheless, we evaluate our test on graphs inferred by LiNGAM (Shimizu et al. 2006), CAM (Bühlmann, Peters, and Ernest 2014), and NOTEARS (Zheng et al. 2018). We chose those algorithms as they are not solely based on CIs (note, however, that e.g. in LinGAM independence of noise entails CI). More details are given in Supp. C.3.

5.2 Synthetic Data and Graphs

For evaluating our method on synthetic data we sample random \mathcal{G}^* under the Erdős-Rényi model (Erdős and Rényi 1959) with $n \in \{10, 20, 30\}$ nodes and an expected degree $d \in \{1, 2, 3\}$, denoted as ER- n - d . To generate synthetic data from \mathcal{G}^* , conditionals are modeled as additive noise models $X_i = f_i(\text{Pa}_i^{\mathcal{G}^*}) + N_i$ with N_i sampled from a normal distribution and f_i either being random (nonlinear) MLPs, or

a linear combination of the node’s parents. The exogenous variables are sampled from a standard normal, uniform, or Gaussian mixture distribution.

For all experiments on synthetic data we sample $T = 10^3$ node permutations and use datasets with $N = 10^3$ observations. To investigate the effect of N and T on p_{LMC} we run ablation studies on nonlinear data with $N, T \in \{10^1, 10^2, 10^3, 10^4\}$. More information on implementation and parameter choices is provided in Supp. A.4.

5.3 Real Data

To evaluate our proposed metric on real-world data, we consider three datasets with established consensus graphs serving as ground truth. We provide further details on the data in Supp. A.5.

Protein Signaling Network (Sachs et al. 2005) This open dataset contains quantitative measurements of the expression levels of $n = 11$ phosphorylated proteins and phospholipids in the human primary T cell signaling network. The $N = 853$ observational measurements, corresponding to individual cells, were acquired via intracellular multicolor flow cytometry (Sachs et al. 2005). The consensus DAG contains 19 edges ($d \approx 3.45$).

Auto MPG (Quinlan 1993) The Auto MPG dataset contains eight attributes (three multivalued discrete and five continuous) with the fuel consumption in miles per gallon (mpg) for $N = 398$ unique car models. While the original use of the data was to predict mpg of a car, in line with previous works on causal inference (e.g. Wang and Mueller 2017; Teshima and Sugiyama 2021), we use a consensus network ($n = 6, 9$ edges, $d = 3$).

Application Performance Monitoring (APM) We collect trace data of microservices in a distributed system hosted on Amazon Web Services (AWS). The traces contain latency information on incoming and outbound requests of each service, averaged over 20 min, making observations approximately i.i.d. In total we ran the application for six days, leading to $N = 432$ observations. We test the working hypothesis that the transpose of the dependency graph ($n = 39, 40$ edges, $d \approx 2.05$) of the application is the true causal DAG (for a discussion of this hypothesis, see Sec. 4).

6 Results

6.1 Simulated Graphs

Nonlinear Mechanisms Figure 3a depicts mean p_{LMC} for 50 synthetic graphs of various size and sparsity with nonlinear mechanisms modeled using random MLPs. As expected, we find that the average p_{LMC} monotonically decreases with increasing amount of domain knowledge for both models of domain experts. When the domain expert has complete knowledge ($\hat{\mathcal{G}} = \mathcal{G}^*$, corresponding to $|K|/|V| = 1$ for DE- V and $\text{SHD}/|\mathcal{E}| = 0$ for DE- \mathcal{E}), we reject the hypothesis that the DAG is as bad as a random node permutation with significance level $\alpha = 1\%$ for all configurations ($p_{\text{LMC}} < 0.005$).

	Sachs	Auto MPG	APM
$p_{\text{LMC}}^{\mathcal{G}^*, \mathcal{D}}$	0.031 ± 0.006	0.03 ± 0.0052	0.00
95% conf. int.	[0.020, 0.042]	[0.019, 0.04]	–
$\phi_{\text{LMC}}^{\mathcal{G}^*, \mathcal{D}}$	0.08	0.50	0.22

Table 1: Mean p_{LMC} and standard deviation (over 50 trials) for the consensus graphs of the real-world data sets with 95% confidence interval. For $\alpha = 5\%$ we reject the hypotheses that the graphs are as bad as random ones, despite high fractions of violations $\phi_{\text{LMC}}^{\mathcal{G}^*, \mathcal{D}}$ (3).

	NOTEARS	CAM	LiNGAM
$p_{\text{LMC}}^{\hat{\mathcal{G}}, \mathcal{D}}$	0.548 ± 0.237	0.0724 ± 0.0866	0.0362 ± 0.1290
$\text{SHD}/ \mathcal{E} $	2.780 ± 0.241	1.880 ± 0.198	1.0600 ± 0.0806

Table 2: p_{LMC} and SHD for graphs inferred by causal discovery on the Sachs et al. (2005) data.

Gaussian-Linear Mechanisms In the supplemental, Fig. A3 we report p_{LMC} for synthetic graphs of various size and sparsity and with linear-gaussian mechanisms. Similar to DAGs with nonlinear mechanisms, we observe that our metric strictly decreases with increasing amount of domain knowledge for both models of domain experts. If $\hat{\mathcal{G}} = \mathcal{G}^*$, we would not falsify the true graph using our metric and significance level $\alpha = 1\%$.

Effect of Number of Sampled Permutations and Number of Observations Further, we investigate the effect of the number of permutations T and sample size N on p_{LMC} for synthetic DAGs with nonlinear mechanisms (Tab. A1). To limit the running time of the experiment, we only evaluate p_{LMC} for the true graph, i.e. $\hat{\mathcal{G}} = \mathcal{G}^*$. Here, we notice that on average our metric is consistently below 0.05, and therefore we would not reject the true graph with significance level $\alpha = 5\%$. The only exception to this are graphs with few nodes (ER-10- d), evaluated on very few ($N = 10$) samples.

6.2 Real World Applications

Figure 3b shows mean p_{LMC} over 50 sampled given DAGs for the three real-world datasets. Similar to the experiments with synthetic data, we find that with increasing amount of domain knowledge (higher $|K|/|V|$, lower $\text{SHD}/|\mathcal{E}|$) p_{LMC} is strictly decreasing. When $\hat{\mathcal{G}} = \mathcal{G}^*$ we reject the hypotheses that the given graphs are as bad as a random node permutation at $\alpha = 5\%$ for all datasets.

Furthermore, we find that for all real-world datasets the fraction of LMC violations $\phi_{\text{LMC}}^{\mathcal{G}^*, \mathcal{D}}$ (3) is higher than the expected type I error rate of 5% for the significance level $\alpha = 5\%$ we used for all conditional independence tests throughout this work (c.f. Tab. 1). Thus, using $\phi_{\text{LMC}}^{\mathcal{G}^*, \mathcal{D}}$ as a metric, we would falsely reject the true causal graph, naively assuming our CI tests would have valid level.

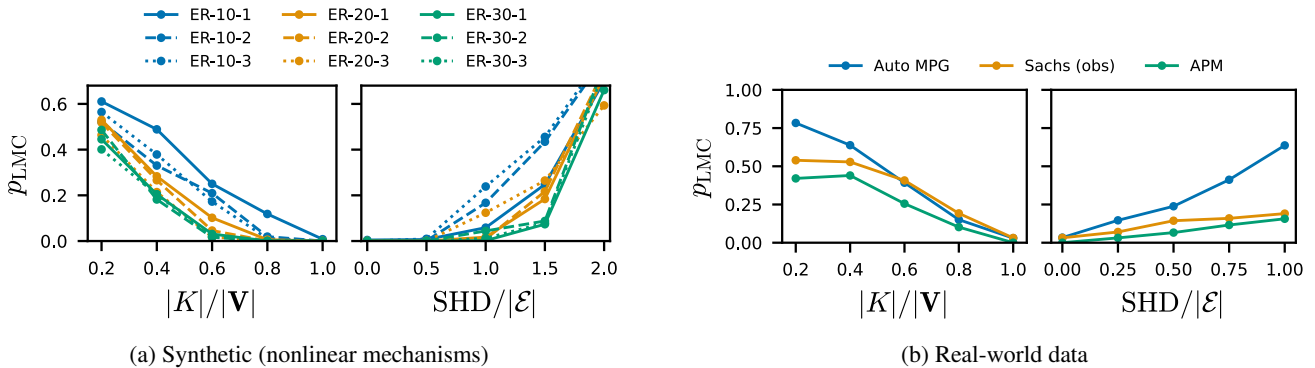


Figure 3: Mean p_{LMC} for two types of domain experts, simulated via $DE-V$ (left; smaller numbers correspond to less domain knowledge) and $DE-E$ (right; smaller numbers correspond to more domain knowledge). On synthetic data (a) for the true DAG ($|K|/|V| = 1$; $SHD/|E| = 0$), we reject the null that the DAG is as bad as random with $\alpha = 1\%$ for all configurations. \hat{G} is falsified with the same α if $|K|/|V| \leq 0.6$ or $SHD/|E| \geq 1.5$. On real-world data (b), for the true DAG, we reject the null that it is as bad as random with $\alpha = 5\%$ and \hat{G} is falsified with the same α for $|K|/|V| \leq 0.8$ or $SHD/|E| \geq 0.5$ for all datasets.

Nodes n	10	50	100	200
Runtime [s]	5 ± 3	97 ± 7	507 ± 18	$3,386 \pm 87$

Table 3: Runtime of p_{LMC} for large graphs with up to 200 nodes. All graphs were modeled as $ER-n-1$, $n \in \{10, 50, 100, 200\}$. Data were generated using nonlinear conditionals and $N = 1000$ samples. For each test we sample 100 permutations, sufficient to reject the null at $\alpha = 1\%$. As CI test we employed the GCM with boosted decision trees as regressor. See Supp. A.6 for more details.

6.3 Causal Discovery Algorithms

While the main scope of this work is to evaluate user-given graphs, we conduct additional experiments with \hat{G} inferred via causal discovery. On the Sachs et al. (2005) data we find that graphs inferred by NOTEARS and CAM are not significantly better than random, whereas graphs inferred by LiNGAM are not falsified using our metric at $\alpha = 5\%$ (Tab. 2). Furthermore, a ranking based on our metric is in accordance with an SHD ranking (NOTEARS $>$ CAM $>$ LiNGAM). Both inequalities are significant with $p < 0.001$ for SHD and p_{LMC} when tested using a Wilcoxon signed-rank test. For further experimental results see Supp. C.3.

6.4 Runtime

Large graphs may entail thousands of CIs that need to be tested in order to compute our metric. Therefore, we evaluated the feasibility of applying p_{LMC} to graphs with up to 200 nodes (Tab. 3). We find that runtimes are reasonably fast ($< 1h$) even for very large graphs. Note, that domain experts and causal discovery methods would likely take much longer to come up with a DAG \hat{G} in the first place. E.g., on DAGs with 50 and 100 nodes CAM and NOTEARS are about an order of magnitude slower, respectively (Rolland et al. 2022; Lachapelle et al. 2020). More information on the algorithmic complexity of our metric is provided in Supp. B.

7 Discussion

In this work we addressed the lack of a suitable metric to evaluate an estimated DAG on observed data. To this end we discussed an existing absolute metric that, without a baseline comparison, is difficult to interpret.

We defined a set of properties that a suitable baseline should satisfy and found that sampling random node permutations of the given DAG is a natural way to satisfy these requirements. Using this baseline, we derived a novel metric which comprises two tests that measure first, how characteristic a given graph is in the sense that it is falsifiable by testing CIs and second, whether the given graph is significantly better than a random one in terms of CIs. Using graphs originating from emulated domain experts and causal discovery algorithms we evaluated our method on two types of data. Synthetic data with known true DAG and real-world data with a consensus graph established in the literature. Furthermore, we contributed a novel data set from cloud monitoring where an estimate of the true causal graph is given by the inverted dependency graph.

We acknowledge that our baseline may seem like a low bar to clear and other baselines could be considered in future work. However, we argue that besides being a natural choice that satisfies a number of desirable properties (P1 – P3 in Sec. 4) our experiments show that ‘better than random’ is a surprisingly high threshold that simulated domain experts and causal discovery algorithms fail to meet in many settings. We also note that applying our algorithm to DAGs inferred by causal discovery algorithms should be done with caution as elaborated in Sec. 5.1 and other approaches may be better suited for this purpose (Faller et al. 2024). Furthermore, in certain applications, specific suggestions for local edge improvements that go beyond the simple report of the triplets that result in LMC violations may be desired, and we leave this interesting direction for future work. Another promising direction could be to extend our node-permutation baseline to the likelihood $p(\mathcal{D}|\hat{G})$ (c.f. Sec. 2).

References

- Acid, S.; and de Campos, L. M. 2003. Searching for Bayesian Network Structures in the Space of Restricted Acyclic Partially Directed Graphs. *Journal of Artificial Intelligence Research*, 18(1): 445–490.
- Amrhein, V.; Greenland, S.; and McShane, B. 2019. Scientists rise up against statistical significance. *Nature*, 567(7748): 305–307.
- Blöbaum, P.; Götz, P.; Budhathoki, K.; Mastakouri, A. A.; and Janzing, D. 2024. DoWhy-GCM: An Extension of DoWhy for Causal Inference in Graphical Causal Models. *Journal of Machine Learning Research*, 25(147): 1–7.
- Budhathoki, K.; Minorics, L.; Bloebaum, P.; and Janzing, D. 2022. Causal Structure-Based Root Cause Analysis of Outliers. In *Proceedings of the 39th International Conference on Machine Learning*, 2357–2369. PMLR.
- Bühlmann, P.; Peters, J.; and Ernest, J. 2014. CAM: Causal Additive Models, High-Dimensional Order Search and Penalized Regression. *The Annals of Statistics*, 42(6): 2526–2556.
- Chalupka, K.; Perona, P.; and Eberhardt, F. 2015. Visual Causal Feature Learning. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*, UAI’15, 181–190.
- Chickering, D. M. 1996. Learning Bayesian Networks Is NP-Complete. In Fisher, D.; and Lenz, H.-J., eds., *Learning from Data: Artificial Intelligence and Statistics V*, Lecture Notes in Statistics, 121–130. New York, NY: Springer.
- Chickering, D. M.; Heckerman, D.; and Meek, C. 2004. Large-Sample Learning of Bayesian Networks Is NP-Hard. *Journal of Machine Learning Research*, 5: 1287–1330.
- Claassen, T.; Mooij, J. M.; and Heskes, T. 2013. Learning Sparse Causal Models Is Not NP-hard. In *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence*, UAI’13, 172–181.
- Erdős, P.; and Rényi, A. 1959. On Random Graphs I. *Publicationes Mathematicae Debrecen*, 6: 290–297.
- Faller, P. M.; Vankadara, L. C.; Mastakouri, A. A.; Locatello, F.; and Janzing, D. 2024. Self-Compatibility: Evaluating Causal Discovery without Ground Truth. In *Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS)*.
- Friedman, N.; and Koller, D. 2003. Being Bayesian About Network Structure. A Bayesian Approach to Structure Discovery in Bayesian Networks. *Machine Learning*, 50(1): 95–125.
- Geiger, D.; Verma, T.; and Pearl, J. 1990. Identifying Independence in Bayesian Networks. *Networks*, 20(5): 507–534.
- Giudici, P.; and Castelo, R. 2003. Improving Markov Chain Monte Carlo Model Search for Data Mining. *Machine Learning*, 50(1): 127–158.
- Glymour, C.; Zhang, K.; and Spirtes, P. 2019. Review of Causal Discovery Methods Based on Graphical Models. *Frontiers in Genetics*, 10.
- Huber, W.; Carey, V. J.; Long, L.; Falcon, S.; and Gentleman, R. 2007. Graphs in Molecular Biology. *BMC Bioinformatics*, 8(Suppl 6): S8.
- Lachapelle, S.; Brouillard, P.; Deleu, T.; and Lacoste-Julien, S. 2020. Gradient-Based Neural DAG Learning. In *International Conference on Learning Representations*.
- Lauritzen, S. 2020. *Lecture Notes: Lectures on Graphical Models*, volume 3. University of Copenhagen.
- Lauritzen, S. L.; Dawid, A. P.; Larsen, B. N.; and Leimer, H.-G. 1990. Independence Properties of Directed Markov Fields. *Networks*, 20(5): 491–505.
- Lemeire, J.; Meganck, S.; Cartella, F.; and Liu, T. 2012. Conservative Independence-Based Causal Structure Learning in Absence of Adjacency Faithfulness. *International Journal of Approximate Reasoning*, 53(9): 1305–1325.
- Madigan, D.; York, J.; and Allard, D. 1995. Bayesian Graphical Models for Discrete Data. *International Statistical Review / Revue Internationale de Statistique*, 63(2): 215–232.
- Mastakouri, A.; Schölkopf, B.; and Janzing, D. 2019. Selecting causal brain features with a single conditional independence test per feature. In *Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS’19)*.
- Meek, C. 1995. Strong Completeness and Faithfulness in Bayesian Networks. In *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence (UAI’95)*, 411–418.
- Mooij, J. M.; Peters, J.; Janzing, D.; Zscheischler, J.; and Schölkopf, B. 2016. Distinguishing Cause from Effect Using Observational Data: Methods and Benchmarks. *Journal of Machine Learning Research*, 17(1): 1532–4435.
- Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. The Morgan Kaufmann Series in Representation and Reasoning. San Francisco, California: Morgan Kaufmann.
- Pearl, J. 2009. *Causality: Models, Reasoning and Inference*. Cambridge: Cambridge University Press, second edition.
- Pearl, J.; and Paz, A. 1986. Graphoids: Graph-Based Logic for Reasoning about Relevance Relations or When Would x Tell You More about y If You Already Know z? In *Proceedings of the 7th European Conference on Artificial Intelligence (ECAI’86)*, volume 2, 357–363.
- Peters, J.; and Bühlmann, P. 2014. Identifiability of Gaussian Structural Equation Models with Equal Error Variances. *Biometrika*, 101(1): 219–228.
- Peters, J.; and Bühlmann, P. 2015. Structural Intervention Distance for Evaluating Causal Graphs. *Neural Computation*, 27(3): 771–799.
- Peters, J.; Janzing, D.; and Schölkopf, B. 2017. *Elements of Causal Inference: Foundations and Learning Algorithms*. Adaptive Computation and Machine Learning Series. Cambridge, MA, USA: The MIT Press.
- Pingault, J.-B.; O’Reilly, P. F.; Schoeler, T.; Ploubidis, G. B.; Rijdsdijk, F.; and Dudbridge, F. 2018. Using Genetic Data to Strengthen Causal Inference in Observational Research. *Nature Reviews Genetics*, 19(9): 566–580.

- Pitchforth, J.; and Mengersen, K. 2013. A Proposed Validation Framework for Expert Elicited Bayesian Networks. *Expert Systems with Applications*, 40(1): 162–167.
- Popper, K. R. 2005. *The logic of scientific discovery*. Routledge. German original: “Logik der Forschung. Zur Erkenntnistheorie der modernen Naturwissenschaft” (1934).
- Quinlan, J. R. 1993. Combining Instance-Based and Model-Based Learning. In *Proceedings of the Tenth International Conference on Machine Learning (ICML’93)*, 236–243.
- Ramsey, J.; and Andrews, B. 2018. FASK with Interventional Knowledge Recovers Edges from the Sachs Model. arxiv:1805.03108.
- Reynolds, R. J.; Scott, R. T.; Turner, R. T.; Iwaniec, U. T.; Bouxsein, M. L.; Sanders, L. M.; and Antonsen, E. L. 2022. Validating Causal Diagrams of Human Health Risks for Spaceflight: An Example Using Bone Data from Rodents. *Biomedicines*, 10(9): 2187.
- Rolland, P.; Cevher, V.; Kleindessner, M.; Russell, C.; Janzing, D.; Schölkopf, B.; and Locatello, F. 2022. Score Matching Enables Causal Discovery of Nonlinear Additive Noise Models. In *Proceedings of the 39th International Conference on Machine Learning (ICML’22)*, 18741–18753.
- Rose, H. E. 2009. Action and the Orbit–Stabiliser Theorem. In Rose, H., ed., *A Course on Finite Groups*, Universitext, 91–111. London: Springer.
- Sachs, K.; Perez, O.; Pe’er, D.; Lauffenburger, D. A.; and Nolan, G. P. 2005. Causal Protein-Signaling Networks Derived from Multiparameter Single-Cell Data. *Science*, 308(5721): 523–529.
- Shah, R. D.; and Peters, J. 2020. The Hardness of Conditional Independence Testing and the Generalised Covariance Measure. *The Annals of Statistics*, 48(3): 1514–1538.
- Shimizu, S.; Hoyer, P. O.; Hyvärinen, A.; and Kerminen, A. 2006. A Linear Non-Gaussian Acyclic Model for Causal Discovery. *Journal of Machine Learning Research*, 7(72): 2003–2030.
- Shrier, I.; and Platt, R. W. 2008. Reducing Bias through Directed Acyclic Graphs. *BMC Medical Research Methodology*, 8(1): 70.
- Spirtes, P.; Glymour, C.; and Scheines, R. 2001. *Causation, Prediction, and Search*. The MIT Press.
- Teshima, T.; and Sugiyama, M. 2021. Incorporating Causal Graphical Prior Knowledge into Predictive Modeling via Simple Data Augmentation. In *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence (UAI’21)*, 86–96.
- Textor, J.; van der Zander, B.; Gilthorpe, M. S.; Liskiewicz, M.; and Ellison, G. T. 2016. Robust Causal Inference Using Directed Acyclic Graphs: The R Package ‘Dagitty’. *International Journal of Epidemiology*, 45(6): 1887–1894.
- Tsamardinos, I.; Brown, L. E.; and Aliferis, C. F. 2006. The Max-Min Hill-Climbing Bayesian Network Structure Learning Algorithm. *Machine Learning*, 65(1): 31–78.
- Uhler, C.; Raskutti, G.; Bühlmann, P.; and Yu, B. 2013. Geometry of the Faithfulness Assumption in Causal Inference. *The Annals of Statistics*, 41(2): 436–463.
- Vidgen, B.; and Yasseri, T. 2016. P-Values: Misunderstood and Misused. *Frontiers in Physics*, 4.
- Wang, J.; and Mueller, K. 2017. Visual Causality Analysis Made Practical. In *2017 IEEE Conference on Visual Analytics Science and Technology (VAST)*, 151–161.
- Wang, L.; and Sahbi, H. 2013. Directed Acyclic Graph Kernels for Action Recognition. In *2013 IEEE International Conference on Computer Vision*, 3168–3175.
- Wang, T.; Huang, J.; Zhang, H.; and Sun, Q. 2020. Visual Commonsense Representation Learning via Causal Inference. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 1547–1550.
- Wasserstein, R. L.; and Lazar, N. A. 2016. The ASA Statement on p-Values: Context, Process, and Purpose. *The American Statistician*, 70(2): 129–133.
- Zhang, K.; Peters, J.; Janzing, D.; and Schölkopf, B. 2011. Kernel-Based Conditional Independence Test and Application in Causal Discovery. In *Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence (UAI’11)*, 804–813.
- Zheng, X.; Aragam, B.; Ravikumar, P. K.; and Xing, E. P. 2018. DAGs with NO TEARS: Continuous Optimization for Structure Learning. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems (NeurIPS’18)*.