

Causal Discovery for Observational Sciences Using Supervised Machine Learning

ANNE HELBY PETERSEN^{1,*}, JOSEPH RAMSEY², CLAUD THORN EKSTRØM¹, AND
PETER SPIRITES²

¹*Section of Biostatistics, Department of Public Health, University of Copenhagen, Denmark*

²*Department of Philosophy, Carnegie Mellon University, USA*

Abstract

Causal inference can estimate causal effects, but unless data are collected experimentally, statistical analyses must rely on pre-specified causal models. Causal discovery algorithms are empirical methods for constructing such causal models from data. Several asymptotically correct discovery methods already exist, but they generally struggle on smaller samples. Moreover, most methods focus on very sparse causal models, which may not always be a realistic representation of real-life data generating mechanisms. Finally, while causal relationships suggested by the methods often hold true, their claims about causal non-relatedness have high error rates. This non-conservative error trade off is not ideal for observational sciences, where the resulting model is directly used to inform causal inference: A causal model with many missing causal relations entails too strong assumptions and may lead to biased effect estimates. We propose a new causal discovery method that addresses these three shortcomings: Supervised learning discovery (SLdisco). SLdisco uses supervised machine learning to obtain a mapping from observational data to equivalence classes of causal models. We evaluate SLdisco in a large simulation study based on Gaussian data and we consider several choices of model size and sample size. We find that SLdisco is more conservative, only moderately less informative and less sensitive towards sample size than existing procedures. We furthermore provide a real epidemiological data application. We use random subsampling to investigate real data performance on small samples and again find that SLdisco is less sensitive towards sample size and hence seems to better utilize the information available in small datasets.

Keywords *causal learning; causality; deep learning; neural network; observational data; structure learning*

1 Introduction

Questions of cause and effect are prevalent in many scientific applications: Why is the global temperature rising? What are the economic consequences of raising a tax? How can we prevent the development of depression? Classic scientific methodology suggests that such questions ought to be answered by relying on experimentation, controlling the supposed cause in order to quantify its effect. However, many interesting and important questions do not allow for such experimentation – it may be impossible, too costly or not ethically permissible. In such cases, empirical sciences generally rely on theory for constructing causal models and use *observational*

*Corresponding author. Email: ahpe@sund.ku.dk.

data for causal inference, i.e. causal effect estimation. While this methodology has provided many useful results, relying on theoretical specifications of causal models induces risk of confirmation bias and greatly limits the scope of topics that can be studied.

Causal discovery provides an empirical alternative to theory-driven causal model specification. Here, causal models are inferred from observational data by use of conditional independence tests or model scoring. Exhaustive searches among candidate causal models are generally not computationally feasible, and hence, existing causal discovery algorithms rely on either sequential testing, greedy search procedures or a combination of both. While several existing algorithms have been proven to be asymptotically correct (including the Peter-Clark (PC) algorithm (Kalisch and Bühlman, 2007) and greedy equivalence search (GES) (Chickering, 2002)), their finite sample properties are compromised by the sequential nature of the procedures. Each algorithmic step involves one or more statistical decisions, which will then inform subsequent steps. But statistical decisions are of course subject to statistical error, and these statistical errors propagate in an unfortunate manner, which makes the algorithms questionable for small or moderate sample sizes.

Moreover, existing algorithms have generally been developed and evaluated with focus on sparse causal models, that is, causal models where there are only few cause-effect links between the involved variables. This makes computations simpler, but sparsity of causal mechanisms may not be a meaningful or natural assumption in many sciences relying on observational data for causal inference. For example in epidemiology, causal etiological models often include numerous complexly intertwined cause-effect relationships that can hardly be described as sparse (De Stavola et al., 2006).

Finally, causal discovery methods have traditionally been evaluated with focus on their performance for *positive* findings: It has been prioritized to ensure that causal mechanisms claimed to be present are truly present. This focus has been useful for informing experimental sciences where such claimed causal relationships can then afterwards be tested and quantified using controlled experiments. However, in truly observational sciences, such as epidemiology, where no interventions are feasible, this evaluation metric may not be as useful. A major concern in epidemiology is the presence of *confounding* and *selection* variables for a given causal inference task and therefore, a causal discovery procedure that provides good guarantees on *negative* findings may be more useful: Knowing that a certain causal effect is for example *not* confounded will help inform what statistical procedures can be used to estimate it, and avoids conditioning on variables that will bias the estimation of this target (i.e. inadvertently conditioning on mediators or colliders).

To address these problems, we propose a new approach for causal discovery using supervised machine learning, supervised learning discovery (SLdisco). We train a machine learning model on simulated data, where the true causal model is known, thereby obtaining a classification function that takes in new observational data and outputs a causal model suggestion. SLdisco is non-sequential in the sense that it learns the full causal model jointly, and hence it is less sensitive to statistical error occurring from small or moderate sample sizes. It seeks to estimate the true model density, and hence it has no built-in preference towards sparse (or dense) causal model. This agnosticism gives it relatively better performance on dense data generating mechanisms. Finally, the construction of the classification function includes a classification threshold, which will provide a direct means to control trade off between different error types, and hence SLdisco can be optimized to focus on correctly identifying negative rather than positive causal relationships.

Our proposed method assumes no latent confounding or selection variables, and seeks to learn the Markov equivalence class of the causal model, as represented via a completed partially

directed acyclic graph (CPDAG). This article is the first to propose to use supervised machine learning for this task.

We provide a proof-of-concept showcase of the potential of SLdisco. We use simple convolutional neural networks as the machine learning model and we train these models on simulated linear Gaussian data. The choice of linear Gaussian data is theoretically attractive, since several existing classical causal discovery algorithms are asymptotically correct for this type of data. This implies that we can directly measure and compare differences in performance due to finite samples, because we know that the involved methods will produce perfect results in the large sample limit.

The article is structured as follows: We first introduce key concepts in Section 2. In Section 3 we provide a thorough exposition of the three issues mentioned above for two popular existing causal discovery methods, the PC and greedy equivalence search (GES) algorithms. In Section 4 we present our proposed methodology, SLdisco. Section 5 describes which metrics we use to evaluate the performance of the different discovery algorithms. We evaluate SLdisco in a large simulation study in Section 6, and in a real epidemiological data application in Section 7. Finally, we discuss strengths and limitations of our method in Section 8 and provide suggestions for future research.

1.1 Related Work

This article is the first to propose using supervised machine learning for CPDAG discovery. However, other authors have proposed to use supervised machine learning for other causal discovery tasks.

Li et al. (2020) look at the special case where the full DAG is identifiable from observational data (so-called linear non-Gaussian additive noise (LiNGAM) setting and the linear Gaussian equivariance setting), and applies a neural network to this task. However, they do not consider influence of sample size or graph density, and their evaluation setup does not inform non-conservative performance. They specifically consider only very sparse graphs and all graphs in each of their training datasets have the same density. Hence their results are conditional on knowing the correct graph density. This is not a realistic assumption in observational sciences.

Yu et al. (2019) and Xu and Xu (2021) propose methods using neural networks in the same special case where the full DAGs is identifiable, and they also consider only a fixed average graph density, and do not study the influence of sample size.

Ke et al. (2022) also address this special case using neural networks, and they furthermore train their model on a mixture of observational and experimental data. Hence, in observational sciences, where experimental data is not available, this method cannot be used.

Zheng et al. (2018) and Zheng et al. (2020) tackle the same special case of identified DAGs, but provide a general framework for formulating the graphical acyclicity criterion as a constraint for a continuous optimization problem.

Finally, Goudet et al. (2018) propose a neural network based method for learning the directions of causal relationships. Their method assumes that the graph skeleton is known and then seeks to orient each edge by looking only at pairs of adjacent variables. While this avoids error propagation, it does not make use of the full information in the data, and the graph skeleton would still have to be learned with other approaches.

In conclusion, while the idea of using supervised machine learning in causal discovery is not new, no other work use this methodology to infer CPDAGs; only DAG discovery and orientation

classification has previously been addressed. This also means that the models described above cannot be used for CPDAG discovery and hence are not suitable benchmarks for our task.

2 Representing Causal Models

DAGs and CPDAGs We focus on causal data generating mechanisms that can be described by directed acyclic graphs (DAGs). We will here only introduce the concepts informally and refer to Appendix A for formal definitions and an overview of notation. We will use $\mathbf{X} = X_1, \dots, X_p$ to refer to a set of p random variables. In a causal DAG, nodes represent random variables, and edges represent causal links. If $X_i \rightarrow X_j$ in a DAG, X_i is a cause of X_j , or, equivalently, intervening on X_i will result in changes in X_j (but not vice versa). Figure 1a provides an example of a DAG. Causal models that can be represented by DAGs have two important assumptions: First, we assume acyclicity, which means that there cannot be any feedback loops in the data generating mechanism. Secondly, we assume that there are no unmeasured variables that are parents of more than one variable (no latent confounders), nor any unobserved selection (no conditioning on latent colliders).

The causal statements of a DAG can be related to conditional independence via the Markov property, which allows us to use the graphical relation d -separation to read off conditional independencies: If a set $S \in \mathbf{X} \setminus \{X_i, X_j\}$ d -separates X_i and X_j , then it holds that $X_i \perp\!\!\!\perp X_j \mid S$. In order to draw causal inference from observational data, we generally need the converse implication as well, namely that conditional independence implies causal unrelatedness. We refer to this as the *faithfulness assumption*, and equipped with both the Markov property and the faithfulness assumption, we obtain a crucial link from testable statistical properties (conditional independence) to causal statements. However, the full DAG is generally not uniquely identified from conditional independencies, as more than one DAG can lead to the same conditional independence statements. All DAGs that entail the same conditional independence statements comprise an equivalence class, usually denoted the *Markov equivalence class*, and can be represented by a completed partially directed acyclic graph (CPDAG) (Peters et al., 2017). A CPDAG has directed edges wherever all members of the equivalence class agree on the orientation of an edge, and undirected edges whenever there is ambivalence about the orientation of the edge. All edge orientations in a CPDAG are related to its *v-structures* (Pearl, 2009), that is, triplets of nodes $X_i \rightarrow X_j \leftarrow X_k$ where X_i and X_k are not adjacent. Note that all DAGs in a Markov equivalence class share the same adjacencies, also known as the graph skeleton (Pearl, 2009). Hence a CPDAG will have the same skeleton as any of its member DAGs. Figure 1 provides an example of a DAG and the CPDAG that represents its Markov equivalence class.




Figure 1: An example of a DAG (a) and the corresponding CPDAG (b) which describes its Markov equivalence class.

Adjacency Matrices DAGs and CPDAGs can also be represented by *adjacency matrices* instead of their graphical objects. A DAG with p nodes is represented by a $p \times p$ matrix M where the (i, j) th element is 1 if and only if there is an edge $X_j \rightarrow X_i$ in the DAG. For CPDAGs, which also have undirected edges, the adjacency matrix is defined as follows:

$$\begin{aligned} M[i, j] = 0 \text{ and } M[j, i] = 1 &\Leftrightarrow X_i \rightarrow X_j \\ M[i, j] = 1 \text{ and } M[j, i] = 0 &\Leftrightarrow X_i \leftarrow X_j \\ M[i, j] = 1 \text{ and } M[j, i] = 1 &\Leftrightarrow X_i - X_j \\ M[i, j] = 0 \text{ and } M[j, i] = 0 &\Leftrightarrow X_i \perp\!\!\!\perp X_j \end{aligned}$$

As examples, the adjacency matrices for DAG and CPDAG in Figure 1 are given by M_1 and M_2 below, respectively:

$$M_1 = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \end{bmatrix} \quad \text{and} \quad M_2 = \begin{bmatrix} 0 & 0 & 0 & 1 & 1 \\ 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 1 \\ 1 & 0 & 1 & 1 & 0 \end{bmatrix}$$

3 Observational Science Challenges for Selected Existing Causal Discovery Procedures

A *causal discovery procedure* takes in a dataset and outputs a representation of its causal data generating mechanism. Unless additional assumptions about the data generating mechanism are made (e.g., distributional assumptions or assumptions about its functional form), only the CPDAG is identifiable. We will now present the three observational science causal discovery challenges in the context of two specific algorithms, PC and GES. These algorithms are chosen because they represent the two main paradigms within causal discovery: Constraint-based methods (PC) and score-based methods (GES). Constraint-based methods use conditional independence tests to recover the true CPDAG, while score-based methods search through candidate CPDAGs and score them. Both methods rely on numerous sequential statistical decisions in a manner that is problematic for small sample performance and induces a bias towards sparse models, as we will explain further below.

3.1 Small Sample and Dense Graph Challenges for the PC Algorithm

The PC algorithm (Spirtes and Glymour, 1991) reconstructs a causal CPDAG from observational data as follows: First, it starts with a completely connected graph over the nodes. It then seeks to prune away edges by looking for *separating sets*. For each adjacent pair of variables, X_i and X_j , the algorithm searches for a set S among nodes that are connected to X_i or X_j such that $X_i \perp\!\!\!\perp X_j \mid S$. If such a (possibly empty) set exists, the edge between X_i and X_j is removed. The algorithm will first consider small separating sets (starting with $S = \emptyset$, i.e., marginal independence), and then, if necessary, increase the size of S . Starting with small candidates for S is attractive in the interest of efficient and precise computations, as conditional independence tests for larger S are both more computationally expensive and require more statistical power. When all possible edge pruning steps have been carried out, the algorithm orients edges by

applying a complete list of orientation rules, which make use of 1) v -structure properties, and 2) the assumption of acyclicity.

In empirical applications, the decisions about conditional independence rely on statistical tests. No generally valid, uniformly consistent test for conditional independence exists (Shah and Peters, 2020), but under certain distributional assumptions, valid tests are available. For example, for jointly Gaussian variables, conditional independence can be asserted by testing for vanishing partial correlations. Given a valid conditional independence test, it has been proven that PC is asymptotically correct (Spirtes and Glymour, 1991).

We wish to highlight two important characteristics of this algorithm that explain why PC struggles with two of the three issues mentioned above, namely small sample performance and dense graph performance:

1. The algorithm conducts a large number of sequential tests, and the results of former tests inform which tests will later be conducted. If for example a test reveals that two variables, X_i and X_j are separable, the edge between them is removed, and they are then no longer adjacent and will not be considered again for future separation attempts. Moreover, because only variables connected with X_i or X_j are considered for candidate separating sets, this also has consequences for future tests involving the two variables: If X_i and X_j are considered separable, and hence the edge between them is removed, this may imply that X_j will never again be considered to be included in separating sets for X_i (if there are no other paths connecting them). But if the original test conclusion was faulty, for example due to statistical error, such an error may then propagate through its influence on subsequent tests. A similar problem occurs in the orientation step.
2. The algorithm is biased towards sparse graphs. In dense graphs, separating sets will often be large, as there are many possible paths between two given variables that must all be closed. As the size of the separating set increases, the relative statistical power decreases, and hence the type II error increases, which means that the probability of falsely concluding conditional independence between any two variables is larger. This will lead to removing too many edges and hence PC will be biased towards sparse graphs.

These two characteristics explain why PC may not perform ideally on small or moderate sample sizes (due to error propagation), nor on dense graphs (due to preference towards sparse solutions).

3.2 Small Sample and Dense Graph Challenges for the GES Algorithm

The greedy equivalence search (GES) (Chickering, 2002) algorithm reconstructs a causal CPDAG from observational data as follows: First, it starts with an empty graph, where no nodes are connected. Then, a forward selection type step is carried out: Edges are greedily added until no further edge additions result in increasing a chosen score. Afterwards, a backwards elimination type step is conducted: Edges are greedily removed until no further edge removals result in increasing the score. Finally, the procedure seeks to further increase the score by considering single-edge orientation reversals.

In terms of the statistical error propagation and sparse graph preference issues, we make the following remarks regarding GES:

1. The greedy search strategy implies that former statistical decisions (score comparisons) affect subsequent statistical decisions. Hence statistical errors propagate.
2. Using an empty graph as starting point for a greedy search results in a preference towards sparsity on small samples. On small samples, statistical error will be larger, and this may result in a scoring function with several local optima, even if the scoring function is convex as

the sample size tends to infinity. If such a local optimum exists, the greedy search procedure of GES may get stuck in this optimum. Because sparse graphs are considered first, there is an increased risk of such local optimum being systematically too sparse.

In conclusion, as for PC, on small or moderate samples, GES will tend to propagate errors and have a built-in preference towards sparse graphs.

3.3 Error Trade Off: A Need for Conservative Discovery

Any causal discovery procedure based on empirical data will of course be subject to some amount of statistical error. We will now discuss what type of error is most critical for observational sciences, using epidemiology as a guiding example, and contrast that to the error trade offs inherent in PC and GES.

In epidemiology, the typical causal inference pipeline involves two steps before an actual data analysis is initiated (Greenland et al., 1999):

1. Construct causal model (often by drawing a DAG)
2. Analyze DAG to assert whether the causal estimand of interest (e.g., an average treatment effect) is identifiable from observational data and with which analytical approach and which epidemiological study design.

Importantly, what analytical approach is appropriate is dictated by the structure of the DAG. Which effects are readily identifiable from a given DAG or CPDAG can be determined using for example Pearl's *do-calculus* (Pearl, 2009) (for DAGs) or Perković's *causal identification formula* (Perkovic, 2020) (for CPDAGs). And if the causal effect of interest is identifiable, an appropriate causal inference method can then be applied to obtain an effect estimate, for example regression adjustment, inverse probability weighting, the G-formula or targeted minimum loss estimation. But even if a causal effect of interest first appears unidentified, the epidemiological toolbox includes another remedy: Study designs. By making use of auxiliary variables, natural experiments or stratification among the observations, it may be possible to make otherwise unidentifiable effects estimable. Examples of such designs are instrumental variable designs, sibling comparison designs, difference-in-difference designs and case-control designs.

When asserting whether a certain analytical approach or design is appropriate given a DAG (or CPDAG), *absent* adjacencies are crucial. For example, an absent adjacency may make it appear as if a certain effect of interest is unconfounded, which will make it appear identifiable. On the other hand, a spurious adjacency in the CPDAG may make actually identifiable effects appear unidentified. But as long as the spurious adjacency is undirected, it will not induce biased effect estimates.

If a theory-based causal model is replaced by a causal discovery estimate as the first step, we will therefore prefer an estimated model which rather has too many than too few adjacencies, and rather too many undirected edges than too many falsely oriented edges. We will refer to such a model as a *conservative* estimate of the causal model, because it (ideally) only makes causal claims that we are quite certain about at the cost of making fewer causal claims altogether. Hence, it errs on the side of inconclusiveness, much like the traditional use of the term *conservativeness* in statistics (where for example a confidence interval with larger than nominal coverage is considered conservative).

Using a conservative CPDAG estimate will imply a conservative causal identification analysis: We may not successfully identify effects that *are* present, but unidentified effects will *not* appear identifiable. For example, if an effect is unconfounded in the estimated causal model, it will generally also be unconfounded in the true causal model. But if it is confounded in the

estimated causal model, it may or may not be confounded in the true model.

Since PC and GES are asymptotically correct and have a preference towards sparse graphs on finite samples, they will generally *not* estimate a supergraph of the true skeleton and hence they will not be conservative. On the contrary, they were specifically designed to aid pre-experimental screening where the causal discovery algorithm is used to identify the strongest causal relationships in a given observational dataset, so the experimenter can subsequently design a controlled experiment where the relevant variables are studied further. For this purpose, it is most important that claimed causal relationships are very likely to also be present in the true causal model, while it is less crucial that estimated absent causal relationships reflect the true model. We return to simulation-based evidence of these claims in Section 6.1.

In conclusion, for this purpose of observational science causal discovery, the error-tradeoff of PC and GES may not be very useful; these algorithms will provide models where effects that are in fact confounded appear unconfounded. This can result in analysts choosing inappropriate analytical tools or designs and, as an effect, obtain faulty causal conclusions along with biased causal effect estimates. Instead, a conservative causal discovery procedure is needed.

4 Supervised Machine Learning for Causal Discovery

We here propose a supervised machine learning approach for conducting causal discovery, and we will refer to this method as *supervised learning discovery* (SLdisco). SLdisco is motivated by the three limitations of existing causal discovery algorithms highlighted above. The main ideas of SLdisco is to 1) avoid poor small sample performance due to statistical error propagation by learning the full causal structure jointly, 2) avoid a preference towards sparse (or dense) graphs by building a machine learning model that is agnostic towards outputted sparsity, and 3) obtain a direct means to affect error trade offs by outputting probabilities of edge/orientation presence rather than binary decisions, whose error trade offs can then be controlled using thresholding and post-processing. SLdisco consists of three steps (further details below):

1. Simulate training data with known data generating mechanisms.
2. Train a machine learning model on the simulated training data using true CPDAGs as labels.
3. Use the resulting classification function as a one-step causal discovery procedure on real data.

We use the following notation: The terms variable and node are used interchangeably. We use p to denote the number of nodes in the DAG/CPDAG. We denote the p variables by X_1, \dots, X_p , and we use n to denote sample sizes for each simulated dataset, while b_{train} and b_{test} are used to denote the size of the training data and testing data, respectively.

4.1 Data Simulation

We limit the scope of the method to linear Gaussian data generating mechanisms. In this case, the correlation matrix is a sufficient statistic and hence it is meaningful to only consider these as input for the discovery method. Note, however, that the choice of Gaussianity in this simulation setup is not only in the interest of computational simplicity. It is also useful from a theoretical point of view, because it makes it possible to separate algorithmic correctness from finite sample performance when evaluating the PC and GES algorithms. As mentioned above, these algorithms are complete and correct on Gaussian data, which means that all errors must be due to finite samples. We elaborate further on this point in the Discussion.

The data are simulated in a three-step procedure:

1. Construct DAG with randomly drawn density (0–80% missing edges compared to fully connected).
2. Simulate from a linear Gaussian structural equation model according to the DAG, and compute the correlation matrix.
3. Construct CPDAG adjacency matrix corresponding to the DAG.

Technical details about the data simulation are provided in Appendix B, and here we only provide a brief overview.

The DAGs in step 1) are constructed such that the distribution of the number of edges is uniform on the interval $(m(p)_{\min}, m(p)_{\max})$ where $m(p)_{\max} = \sum_{i=1}^{p-1} i$ (fully connected DAG) and $m(p)_{\min} = m(p)_{\max} \cdot 0.2$. We thus consider a broad distribution of graph densities and include very dense graphs in the simulation setup.

We simulate each variable X_i according to the DAG (which holds information about parent sets) using a structural equation:

$$X_i := \sum_{X_j \in \text{pa}(X_i)} X_j \cdot \beta_{j,i} + \epsilon_i$$

where $\epsilon_i \sim N(0, \sigma_i^2)$ independently. All parameter values ($\beta_{j,i}$ and σ_i) are drawn randomly.

As data features, we use correlation matrices computed on the full dataset of X_1, \dots, X_p variables. As data labels, we use the CPDAGs encompassing the true DAGs. We represent the CPDAGs by their adjacency matrices.

4.2 Machine Learning Model

We use neural networks to construct a mapping from correlation matrices to adjacency matrices. In the special case where the DAG is fully identifiable from observational data (e.g., linear Gaussian data generating mechanisms with equal residual variances), this learning task is achievable since 1) DAG discovery can be formulated as a continuous optimization problem (Zheng et al., 2018, 2020) and 2) a neural network is a universal approximator (Hornik et al., 1989). Hence, with oracle correlation matrix information (corresponding to $n \rightarrow \infty$), a sufficiently large training data set and a good neural network architecture, it is possible to estimate adjacency matrices at any level of accuracy in this special case. Clearly, oracle correlation matrix information is not available in real-data settings, so in practice, with a sufficiently sophisticated architecture and sufficiently large training data size (which may both be feasible to obtain), we should be able to achieve a level of accuracy that is *only* limited by how well we estimate the correlation matrix (or, more generally, a sufficient statistic). We find it plausible that a similar result holds for our more general case of CPDAG estimation, and we believe that the results provided below provide evidence for this claim.

In this article, we do not attempt to provide such a sufficiently sophisticated architecture or training data size. Instead, we present results for a simple, albeit useful convolutional neural network with a set choice of training data size ($b_{\text{train}} = 1,000,000$), in order to provide a proof of concept. Note that appropriate size of the training data pairs will of course depend strongly on the number of nodes in the data generating DAG.

4.2.1 Neural Network and Post-Processing

We use a simple convolutional neural network to learn the relationship between correlation matrices and CPDAG adjacency matrices. Technical details about its architecture and training are provided in Appendix C.

Most importantly, for the final output layer, we use a sigmoid activation function (also known as the *expit* function, i.e. the inverse of the logit function). This means that our neural network model outputs a matrix of probabilities rather than an adjacency matrix. Each element in this probability matrix then corresponds to the estimated probability of that element in the adjacency matrix being one.

This matrix of probabilities then needs to be post-processed to obtain an adjacency matrix. Let O denote the probability matrix. We propose and evaluate two alternative approaches for this, both using a threshold $\tau \in (0, 1)$:

Cut-off: Set adjacency matrix elements to 1 if they are larger than τ , and set all others to zero.

Backwards PC-orientation (BPCO): First use the cut-off post processing method and check whether the resulting matrix, M , is a proper CPDAG adjacency matrix. If so, return M . Otherwise, repeat the following steps until M is a proper CPDAG adjacency matrix:

1. Among the remaining non-zero elements in M , choose the element that had the lowest probability in O and set this element to zero. If M is now a proper CPDAG, return it.
2. Otherwise, create a new adjacency matrix \tilde{M} by modifying M such that only the skeleton and v -structures are preserved. Then apply the orientation rules from the PC algorithm (Meek, 1995) to \tilde{M} and check whether the resulting matrix is a proper CPDAG.
 - (a) If so, return it.
 - (b) If not, discard \tilde{M} and go to step 1.

Note that BPCO will always produce a proper CPDAG:¹ If M ends up with only one edge, \tilde{M} will only have a single undirected edge, which is a proper CPDAG. Moreover, BPCO will always either have the same skeleton as the cutoff method or a sparser one, since the procedure in step 2 only changes edge orientations, not adjacencies.

5 Metrics

As discussed in Section 3.3, we aim for a causal discovery procedure that is conservative, and our choice of relevant metrics will of course reflect this. We will divide the question into two parts: 1) Adjacency performance and 2) orientation performance. In both cases, we report mean metrics both over the full test data, as well as stratified by the true graph density. Moreover, we also consider the estimated number of edges, and compare this to the true number of edges, in order to assess which threshold level (τ) obtains the closest approximation to the true number of edges.

5.1 Adjacency Metrics

Adjacency performance only involves the skeleton of the CPDAG, and as mentioned above, we wish to estimate a supergraph of that skeleton to achieve a conservative causal model estimate. This motivates our first metric of interest, namely the negative predictive value (NPV). Let TP, TN, FP, FN denote the number of true positive, true negative, false positive and false negative

¹This is also theoretically true for PC, however, in practical implementations, the algorithm may return outputs that e.g., have cycles. This occurs when the algorithm receives conflicting information about conditional independencies, which can happen due to assumptions that are not fulfilled or because of finite-sample statistical uncertainty. In contrast, for BPCO we explicitly check whether the output is a proper CPDAG and hence it always returns a CPDAG.

classifications (in this case: adjacencies), respectively. Then the NPV is defined as

$$\text{NPV} = \frac{\text{TN}}{\text{TN} + \text{FN}}$$

i.e., the probability of an estimated missing adjacency to also be missing in the true skeleton. A high value of NPV will thus reflect a conservative skeleton estimate: Missing adjacencies are generally trustworthy. However, one could of course obtain a NPV of 1 simply by trivially returning a fully connected skeleton (using the convention that $\text{NPV} = 1$ for this degenerate case where there are no classified negatives). Thus we also need to consider an secondary metric that describes the error for positive findings, i.e. adjacencies estimated to be present.

For this purpose, we will use the F1 score:

$$\text{F1} = 2 \cdot \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}}.$$

F1 describes a tradeoff between two metrics, precision and recall, where

$$\text{precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \quad \text{and} \quad \text{recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}$$

F1 thus summarizes errors related to positive findings: Precision (also known as the positive predictive value) measures the probability of an estimated adjacency to be a true adjacency, while recall (also known as sensitivity) measures the probability of the estimated graph including a given true adjacency.

In order to obtain a conservative *and* informative model we thus aim for 1) a large NPV (to obtain a supergraph of the skeleton) and 2) an acceptable F1 (so that the supergraph is close to the true skeleton).

5.2 Orientation Metrics

We measure orientation performance by use of conditional endpoint metrics, that is, for each correctly identified adjacency, we consider whether each edge endpoint is correctly estimated. A correctly placed arrowhead is considered a true positive, a correctly placed tail is a true negative, a falsely placed arrowhead is a false positive, while a falsely placed tail is a false negative.

As discussed in Section 3.3, a conservative causal model should rather have too many undirected edges than too many falsely directed edges. Hence, for orientation, a conservative procedure should prioritize obtaining a high value of precision (i.e. positive predictive value). With a high precision, estimated arrowheads will generally be trustworthy, while estimated tails may or may not reflect the true graph. As for the adjacency metrics, a trivially large value of precision can of course be obtained by simply not orienting any edges. Hence, we again need a secondary metric that summarizes errors related to the opposite (now: false) classifications.

To this end, we propose a negative classification equivalent to F1:

$$G1 = 2 \cdot \frac{\text{NPV} \cdot \text{specificity}}{\text{NPV} + \text{specificity}}$$

where

$$\text{specificity} = \frac{\text{TN}}{\text{TN} + \text{FP}}$$

This metric measures the tradeoff between two types of negative classification errors, just as F1 does for positive classification errors.

A conservative causal model should hence have 1) a large orientation precision (thus only claiming orientations that have a high probability of being true), and 2) an acceptable G1 (so that some edges are forced to be oriented).

6 Simulation Study

We evaluate SLdisco by comparing estimated (pseudo) adjacency matrices with true adjacency matrices on a test data set. We vary the following aspect of the procedure and compare the resulting performances:

- The number of observations in each simulated training dataset, n_{train} . We consider $n_{\text{train}} \in \{50, 100, 500, 1000, 5000, 10000, 50000\}$.
- The number of variables in the data generating mechanism, p . We consider $p \in \{5, 10, 20\}$.
- The post processing procedure for converting probability matrices to (pseudo) adjacency matrices. We consider the cutoff and BPCO methods.
- The threshold value, τ . We consider $\tau \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6\}$.

We consider all combinations of the above.

We compare SLdisco with PC and GES. For PC, we consider a sequence of test significance levels, namely $\alpha \in \{10^{-8}, 10^{-4}, 10^{-3}, 0.01, 0.05, 0.1, 0.2, 0.5, 0.8\}$. For GES, we use the Bayesian information criterion (BIC) to score models. In Appendix D, we furthermore compare this choice to an alternative score, namely a modified version of BIC with a larger penalty for the number of variables. We first compare PC and GES in order to choose a single benchmark causal discovery setting for each of the two, which simplifies subsequent comparisons with SLdisco.

For each evaluation scenario, we train the neural network on $b_{\text{train}} = 1,000,000$ correlation matrix/adjacency matrix pairs, and we test the procedure on $b_{\text{test}} = 5000$ correlation matrix/adjacency matrix pairs. For each scenario, we report performance metrics averaged over these 5000 repetitions.

Data simulations and computations for SLdisco and PC are conducted in R. We use the `keras` package (Allaire and Chollet, 2021) for neural network training, the `pcalg` package (Kalisch et al., 2012) for PC computations, and the `causalDisco` (Petersen, 2022) package for evaluations. GES computations are performed in TETRAD (Ramsey et al., 2018) using the fast GES (FGES) version of GES (Ramsey et al., 2017). All code, as well as the simulated data, is available online at <https://github.com/annemme/SLdisco>.

6.1 PC and GES Comparison

Figure 2 compares PC and GES with BIC scoring. For the small $p = 5$ graphs, we see that it is possible to choose a significance level for PC (around 0.05) so that the two methods perform equally well in terms of adjacency F1, adjacency NPV and orientation precision when $n \geq 500$. However, for $p \in \{10, 20\}$, GES consistently produces better performance for the adjacency metrics, no matter the choice of significance level for PC. For the orientation metrics, GES performs better than PC in most cases, and in the cases where PC is superior, both methods perform very poorly (for example, orientation G1 for $p = 20$). All in all, we thus find GES to be the strongest competitor.

In Appendix D, we furthermore compare two different versions of GES using different choices of score and find that standard BIC performs best overall (see Figure 7). We therefore use GES

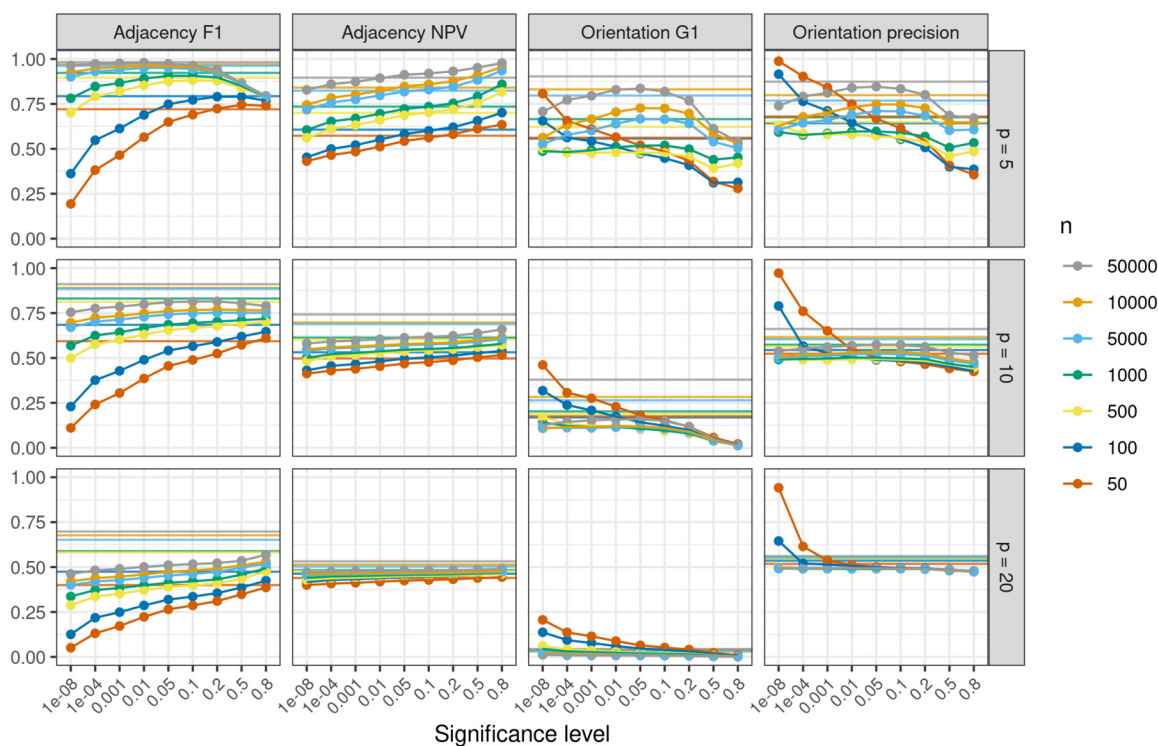


Figure 2: All metrics for PC (lines with dots) for varying values of significance level, compared to GES with BIC scoring (lighter straight lines).

with BIC score as our main GES benchmark in the following.

For simplicity, we also only report one setting for PC below, and we use PC with significance level $\alpha = 0.1$ because this results in good overall performance across metrics (see Figure 2).

6.2 Estimated Number of Edges

Figure 8 (Appendix D) presents the number of estimated edges for SLdisco. SLdisco is able to obtain an average number of edges that is rather close to the truth by setting the threshold appropriately. We find the overall closest approximation of the true number of edges at $\tau = 0.4$ for $p = 5$, at $\tau = 0.4$ for $p = 10$, and at $\tau = 0.3$ for $p = 20$. In the following, for simplicity, we primarily report results from SLdisco with these τ settings, but we note that even better approximations to the true number of edges can be achieved by allowing the value of τ to also depend on n . We propose that in practice, τ may be selected by referring to results like these.

We also find a rather good approximation of the correct number of edges for SLdisco when stratifying by quartiles of true numbers of edges (see Figure 9 in Appendix D), though SLdisco underestimates the number of edges for sparse graphs when $p = 20$.

Figure 3 shows the estimated number of edges for SLdisco with the settings outlined above, compared to PC and GES. For $p = 5$, we see that SLdisco estimates the true number of edges well for small n , but on average adds a little more than one too many edges for large n . GES and PC, on the other hand, underestimates the number of edges quite severely when n is small, but estimates the true number of edges very well for large n . For $p = 10$, SLdisco provides a good approximation to the true number of edges across all values of n . GES performance varies




Figure 3: Estimated and true numbers of edges for SLdisco with $\tau = 0.4$ for $p \in \{5, 10\}$ and $\tau = 0.3$ for $p = 20$, compared with GES and PC. The results are stratified by number of nodes (p), sample size (n) and method. The black, straight line marks the average true number of edges for each value of p .

strongly with n , with underestimation for small n and (slight) overestimation for large n . PC underestimates the number of edges for all n , but less severely as n increases. Finally, for $p = 20$, SLdisco consistently overestimates the number of edges for all n , while GES and PC consistently underestimate, although with less underestimation as n increases.

All in all, we thus find that SLdisco is more prone to overestimate than underestimate the number of edges, while GES and PC generally are more at risk of underestimation. Moreover, the results from SLdisco vary much less with sample size than the results from GES and PC, which are generally very poor for small n .

6.3 Adjacency Results

Figure 4 shows adjacency results. We find that the two SLdisco post-processing methods perform very similarly, so we will refer to them collectively. For small values of n , we find that SLdisco strongly outperforms PC and GES across all values of p . For $p \in \{5, 10\}$, GES and PC converge towards the performance of SLdisco as n increases, and GES even outperforms SLdisco slightly for $p = 10$ and the largest values of n . For $p = 20$, however, GES and PC perform very poorly, while SLdisco remains stronger.

We now turn to F1 results. For $p \in \{5, 10\}$, SLdisco only outperforms GES and PC for the smallest values of n , while PC – and especially GES – achieves higher scores for large n . For $p = 20$, SLdisco is again stronger for smaller values of n , but performs similarly to GES for large n . PC, on the other hand, performs rather poorly in this scenario across all n .

For SLdisco, we find that in some cases, performance does not increase with n . This is likely due to overfitting; all neural networks were trained for the same number of epochs (see Appendix C), and since the classification task should become easier as n increases, this means that we may be overfitting for larger values of n .

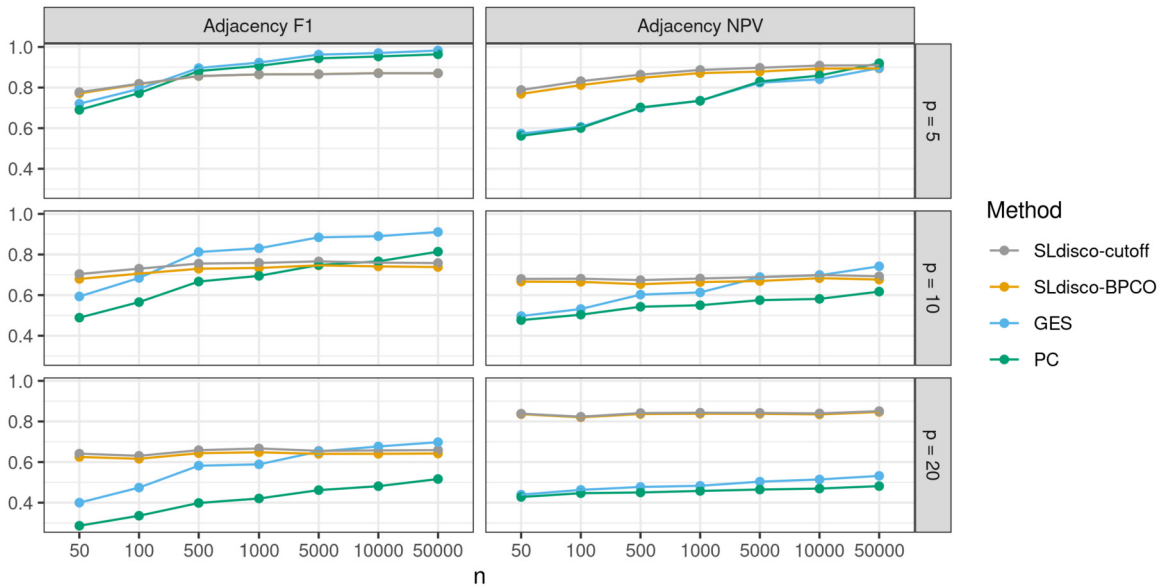


Figure 4: Adjacency F1 and adjacency negative predictive values (NPV). The results are stratified by number of nodes (p), sample size (n) and method.

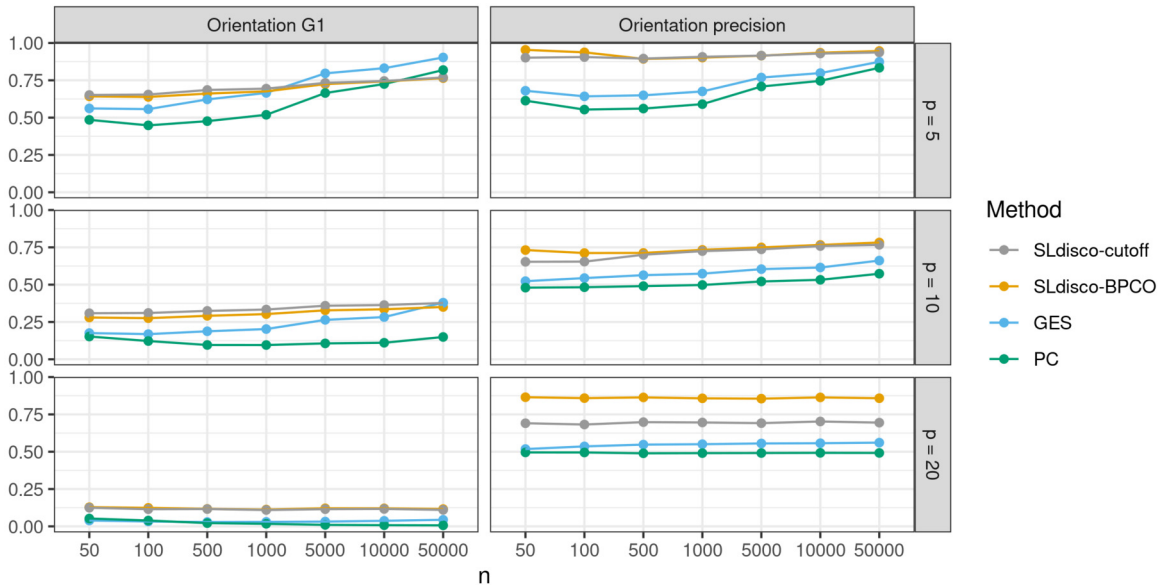


Figure 5: Conditional orientation G1 and conditional orientation precision scores. The results are stratified by number of nodes (p), sample size (n) and method.

6.4 Orientation Results

Figure 5 presents conditional orientation results. For orientation precision, we see that both SLdisco methods outperform GES and PC for all combinations of n and p . We find that SLdisco-BPCO outperforms SLdisco-cutoff quite markedly when $p = 20$.

For orientation G1, we again find that SLdisco outperforms the competitors, except for the scenarios where $p = 5$ and n is large. In these scenarios, GES, and to a lesser extent also PC, outperforms SLdisco. We note, however, that all four methods perform very poorly when $p \in \{5, 10\}$.

We again see a tendency to overfitting for SLdisco, as with the adjacency metrics.

6.5 Results Stratified by True Graph Density

Figure 10 in Appendix D focuses on the BPCO post-processing method, and compare only with GES, which has proven to be the strongest competitor above. The four metrics from above are stratified by graph density, measured as quartiles of the true number of edges. We only present results for $p = 10$, but find similar results for $p \in \{5, 20\}$.

While GES struggles increasingly with adjacency NPV, orientation G1 and orientation precision as the graph density increases, SLdisco is strongest on relatively dense graphs for these metrics. For the dense graphs, SLdisco is hence able to outperform GES for most sample sizes.

But on the sparsest graphs, SLdisco struggles especially with orientation G1 and adjacency F1. In contrast, GES adjacency F1 performance is less affected by graph density. As in the above, we again see that GES depends more strongly on sample size, while SLdisco is able to produce more similar results for both small and large samples.

7 Application

For the real data application, we apply PC, GES and SLdisco to data from the Metropolit Cohort (Osler et al., 2004, 2006). This is a longitudinal dataset following a cohort of Danish men from their birth in 1953 with data collections in childhood and adulthood, as well as register-based information and followup during other periods of their lives. This is a high-quality dataset with good validity and low degree of measurement error (Osler et al., 2006).

We use a subset of 10 variables also used by Petersen et al. (2021), which includes follow-up until 2018. While that study focused on the life course etiology of depression in early old age, the simplified application provided here is intended to explore what social factors across the life course may be relevant for understanding alcohol use in adulthood. The 10 variables are all assigned to one of four periods: Birth, childhood, young adulthood, or adulthood. Table 1 provides an overview of the 10 variables, as well as the data sources used to construct them.

The original dataset also contains a number of binary variables, but we use only numeric variables for a more direct comparison with the simulation study. One numeric variable (*total years of smoking*) is dropped so that we obtain exactly 10 variables, which means that we can easily apply the pretrained SLdisco models from the simulation study. We do not believe that smoking history is a likely cause of alcohol use (or any of the other variables included here), and hence we do not expect omission of this variable to be crucial.

We use only complete cases and condition of cohort members being alive and residing in Denmark in 2018. This results in a dataset of $n = 2928$ observations. Appendix E provides information about marginal distributions and pairwise associations in this data.

7.1 Methods

We of course do not know the true causal model for this data, so we cannot readily apply the same evaluation strategy as in Section 6. Instead, we will:

Table 1: Overview of variables used in the Metropolit application. The data were constructed from the following sources and linked using unique personal identification numbers: A) Midwife notes, collected retrospectively in 1965. B) School survey conducted in 1965. C) The Danish Conscription Database, where data are typically collected when the subjects are approximately 18 years old (1971). D) National Danish socio-economic registers. E) Survey conducted in 2004. Further details about the variables and the data sources are provided in Petersen et al. (2021).

Period	Variable	Data source
Birth	Length	A
Birth	Weight	A
Childhood	Creativity score	B
Childhood	Intelligence score	B
Youth	BMI	C
Youth	Height	C
Youth	Intelligence score	C
Adulthood	Disposable income	D
Adulthood	Alcohol binging frequency	E
Adulthood	BMI	E

1. Discuss plausibility of the causal models proposed by SLdisco, GES and PC, respectively.
2. Evaluate how the performance of these methods varies with sample size.

For each causal discovery method, we use the settings that provided the closest estimation of the true number of edges in the simulation study for $p = 10$ and $n = 5000$ (the closest value we have considered to $n_{\text{metropolit}} = 2928$). We use PC with a significance level of $\alpha = 0.1$, GES with BIC scoring and SLdisco with BPCO post-processing and $\tau = 0.4$.

To address plausibility (task 1), we apply PC, GES and SLdisco to correlation matrices computed on the full Metropolit dataset. We use an SLdisco model trained on datasets with $n = 5000$. We make use of the fact that the Metropolit dataset includes temporal information; each variable is assigned to a specific period and this provides us with a ground truth about *some* of the potential causal orientations: Any edge oriented against the direction of time is necessarily false.

To compare the performance across varying sample sizes (task 2), we use random subsampling. Even though we do not know the ground truth causal model, the more data we have, the better we should be able to estimate it. Hence, we can use CPDAGs estimated from the full Metropolit dataset as a “best estimate”, and compare these graphs to graphs estimated on from correlations matrices computed on random subsets of the Metropolit data. This will allow us to evaluate the impact of small sample sizes on real data. We draw a random subsample for each $n \in \{50, 100, 500, 1000\}$, compute the corresponding correlation matrices, and use these as input for PC, GES and SLdisco. The resulting causal models are then compared with the same method’s “best estimate”, as obtained from task 1. We report the comparisons using the same four metrics as in the simulation study.

We use the `pcalg` (Kalisch et al., 2012) R package for the PC and GES computations, and the `causalDisco` R package (Petersen, 2022) for evaluation and plotting. All code, as well as correlation matrices, is available online at

<https://github.com/annennenne/SLdisco>

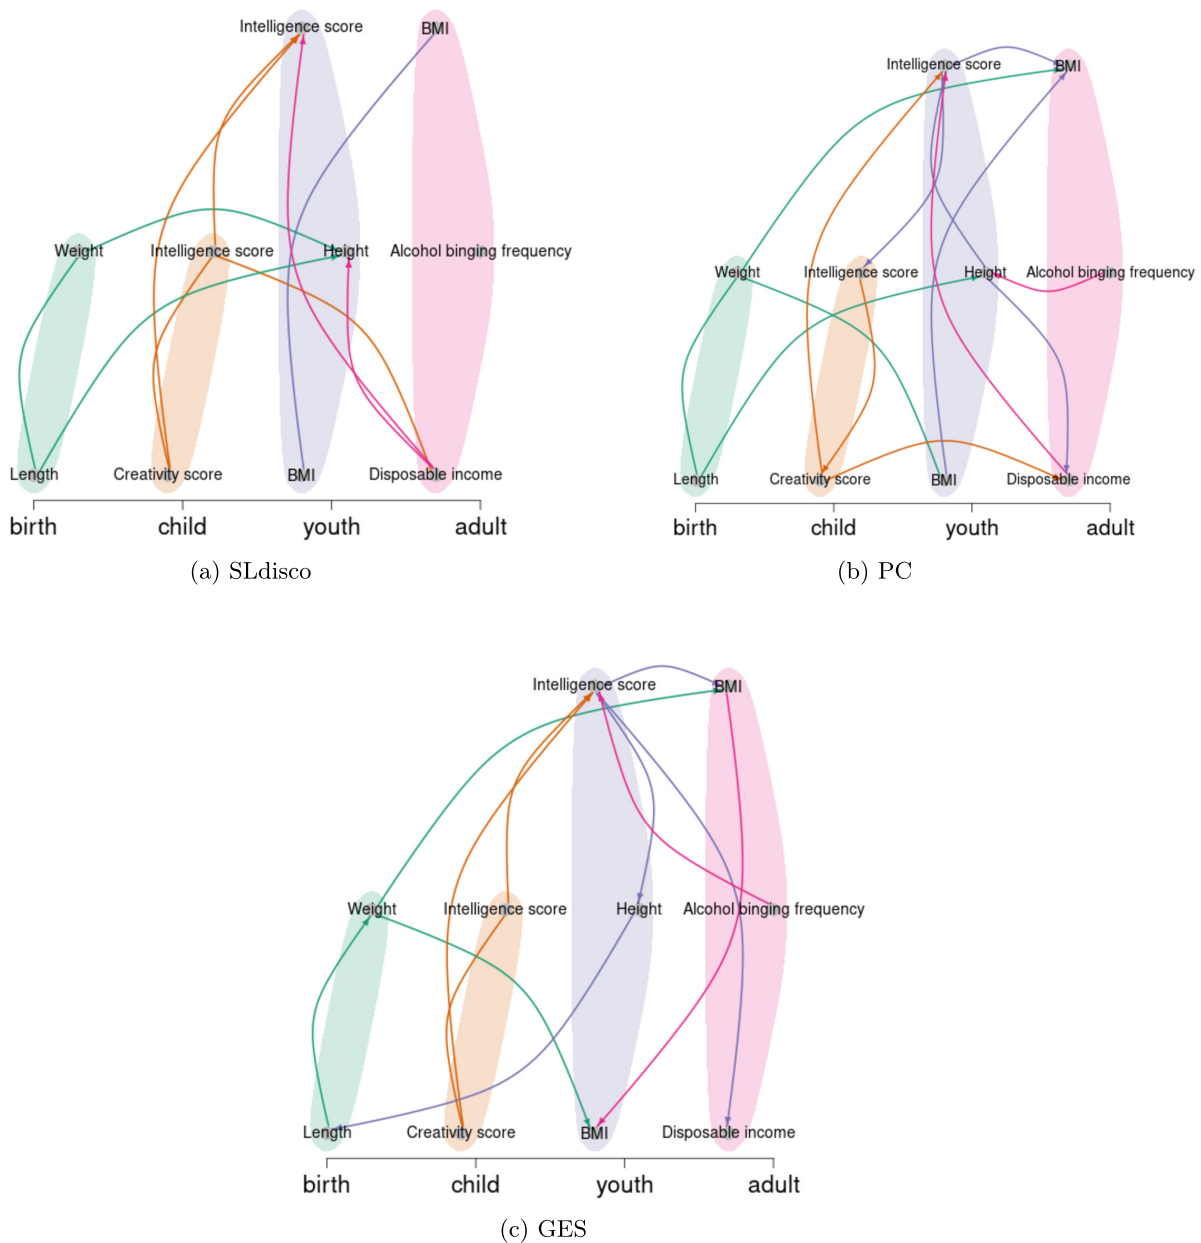


Figure 6: CPDAG estimates for the Metropolit example based on all $n = 2928$ observations using SLdisco (a), PC (b) and GES (c).

7.2 Results

Figure 6 shows the estimated CPDAGs for the Metropolit data for SLdisco, PC and GES, respectively. The variables are arranged in temporal periods (marked in colors) and edges are colored according to the period the originate from, if they are oriented, or simply their first period, if they are unoriented. None of the three methods are provided with external information about the temporal order of the variables.

7.2.1 Plausibility of the Estimated Models

For SLdisco we see that a number of plausible causal links are identified: edges between birth weight, birth length and height in youth, edges between creativity score in childhood and intelligence score in childhood and youth, and an edge between body mass index (BMI) in youth and BMI in adulthood. But the CPDAG also includes two edges that were erroneously oriented against the direction of time: An edge from disposable income in adulthood to height in youth, as well as an edge from disposable income in adulthood to intelligence score in youth. The CPDAG also includes an undirected edge between intelligence score in childhood and disposable income in adulthood, which may or may not be plausible. On the one hand, it would be more obvious that a specific effect of intelligence should go via intelligence score in youth. But on the other hand, socio-economic status may affect intelligence test performance in children (Von Stumm and Plomin, 2015), and hence the edge between intelligence score in childhood and adulthood income could represent an indirect effect due to lack of social mobility. SLdisco does not identify any direct causes of alcohol bingeing among the other included variables.

PC also suggest the wrongly oriented edge between disposable income in adulthood and intelligence score in youth, as well as two other erroneously oriented edges: One from alcohol bingeing frequency in adulthood to height in youth and one from intelligence score in youth to intelligence score in childhood. PC furthermore recovers a number of plausible causal links, e.g., an effect of BMI in youth on BMI in adulthood, and an effect of birth length on height in youth. However, the CPDAG also shows a number of less plausible causal links, for example a causal link between height (youth) and alcohol bingeing (adulthood), height (youth) and income (adulthood), and a direct effect of birth weight on BMI in adulthood (on top of the more plausible effect via BMI in youth). Furthermore, we note that in the current application, PC does not return a proper CPDAG; the graph in Figure 6 (b) includes a cycle: *Childhood creativity score* \rightarrow *youth intelligence score* \rightarrow *childhood intelligence score* \rightarrow *childhood creativity score*. PC may produce graphs that are not proper CPDAGs if there is conflicting information with respect to separating sets, which can occur if some of the algorithm's assumptions (no unobserved confounding, faithfulness and a correct statistical test of independence) are not fulfilled (Kalisch et al., 2012). We return to this point below. Note that this issue cannot occur for BPCO, because BPCO explicitly checks whether the output is a proper CPDAG, and iterates further if not.

GES also finds three edges against the direction of time (BMI in adulthood to BMI in youth, alcohol bingeing in adulthood to intelligence score in youth, and height in youth to birth length), and plausible causal links between the variables related to intelligence and creativity scores. It shares a number of plausible causal links with SLdisco and PC, e.g., links between intelligence score and creativity score variables, as well as links between birth weight and birth length. Among less plausible causal links is the suggested effect of birth weight on adult BMI (not mediated through BMI in youth), and the edge between intelligence score and height, both in youth.

It should be noted that all three methods rest on a strong assumption that may not be reasonable for this example, namely that there are no unobserved confounders or selection variables. If the example does suffer from such unobserved variables it may induce spurious edges in the outputs of PC (Spirtes et al., 2000) and GES (Ogarrio et al., 2016), which may explain why these methods – somewhat surprisingly – produce denser CPDAGs than SLdisco in this application. The fact that PC produces a graph with a cycle also lends evidence to this explanation, as this can be expected behavior when applying PC to data with unobserved confounding (Kalisch et al., 2012). We do not know how density estimation for SLdisco is affected by presence of unobserved confounding, and leave this question to future research. Unobserved confounding

Table 2: Results for subsampled Metropolit data applications. For each method, the subsample based results are compared to results from the same method based on the full Metropolit dataset ($n = 2928$). n_{sub} denotes the number of observations in the subsampled dataset.

Method	n_{sub}	Adjacency metrics		Orientation metrics	
		F1	NPV	G1	Precision
SLdisco	50	0.67	0.88	0.75	1.00
	100	0.67	0.88	0.75	1.00
	500	0.89	0.95	0.55	1.00
	1000	0.95	0.97	0.80	1.00
PC	50	0.53	0.78	0.33	1.00
	100	0.53	0.78	0.33	1.00
	500	0.72	0.85	0.20	0.50
	1000	0.75	0.86	0.33	0.71
GES	50	0.56	0.82	0.33	1.00
	100	0.67	0.85	0.29	1.00
	500	0.64	0.86	0.00	1.00
	1000	0.76	0.89	0.00	0.25

can also lead to wrong edge orientations for both PC and GES (Spirtes et al., 2000; Ogarrio et al., 2016), and hence it may explain some of the faulty oriented edges that we found above.

Moreover, all three methods rely on the data being jointly Gaussian, and this assumptions may not be fully satisfied for some variables, especially the three adulthood variables which have unimodal but skewed distributions (see Figure 11 in Appendix E). All in all, the results presented here should thus be interpreted with some caution, and more work should be dedicated to further investigating the impact of the strong assumptions that SLdisco relies on, and whether they may be relaxed.

7.2.2 Influence of Sample Size

Table 2 presents results from the subsampling experiment. We see that SLdisco produces very similar CPDAGs across the different values of n , as reflected by generally large metric values. The large values of the adjacency metrics show that the estimated graph skeletons are quite similar, and increasingly so when n increases. Unoriented edges are preserved across subsamples (as evident from the orientation precisions), whereas estimation of unoriented edges differs a bit more. In comparison, GES produces lower values and hence is more sensitive towards sample size, and PC is even more sensitive than GES.

8 Discussion

We have proposed a new method for causal discovery based on supervised machine learning, SLdisco, and we have evaluated it focusing on the three highlighted challenges for using existing causal discovery methods in observational sciences, namely 1) poor small sample performance, 2) bias towards sparse graphs, and 3) error-tradeoffs that result in non-conservative causal model estimates.

In a simulation study, we found that SLdisco produces more conservative causal models than both PC and GES: SLdisco generally achieves larger adjacency NPVs and orientation precisions than PC and GES, which means that adjacency absence and oriented edges proposed by SLdisco are more trustworthy. These increases in conservativeness come at the price of modestly reduced informativeness; SLdisco is overall not as good at detecting adjacency presence (lower adjacency F1) and orients fewer edges than PC and GES (lower orientation G1).

Moreover, we found that SLdisco shows promising potential for estimating the correct graph density. We also find that SLdisco is more robust towards graph density and that it shows little sensitivity towards sample size.

In the application, we found SLdisco to produce a mostly plausible CPDAG, whereas PC and GES identified more implausible causal links or orientations. However, their results were also denser, and this suggests that the strong assumptions underlying all three methods – Gaussianity and absence of unobserved confounding – may not be fulfilled for this application, and hence further real data assessments are needed. We were however able to systematically assess the influence of sample size on the estimated causal models in the application as well, and we found that SLdisco was again much less sensitive towards sample size than the other two methods. SLdisco thus makes better use of information available in small sample data, both on simulated and real life data. All in all, we conclude that SLdisco holds the potential to address the three challenges for causal discovery in observational sciences.

We have compared SLdisco with two well-known and well-studied procedures for causal discovery, namely PC and GES, and we have considered a simplified setting with Gaussian data and no unobserved confounding. This is obviously a limitation of the study in terms of generalizability, as we will return to below. However, it is also a strength in terms of validity of the conclusions. PC and GES are correct and complete (with respect to CPDAGs) in the large sample limit for the type of data studied here, and hence we are able to directly study the finite-sample statistical properties of the methods without blurring the picture by other factors. The only reason for any deviance from fully correctly estimated CPDAGs are necessarily due to finite sample properties of the methods. By comparing SLdisco to these two methods, and by focusing on Gaussian data, we thus get a good picture of how sequential testing (for PC) or greedy searching (for GES) impacts finite sample causal discovery performance.

We now turn to the limitations imposed by the use of Gaussian training data for SLdisco. Gaussianity may not always be a reasonable assumption, and it is not clear to what extent the performance of SLdisco depends on this. It is difficult to conduct a general simulation-based investigation into the practical influence of deviations from Gaussianity: We would have to simulate data under a specific alternative distribution, and would hence run into the very same generalization issues, and still not be much more knowledgeable about the performance on real data, where the data generating mechanism is unknown. The real data application did include some variables whose distributions were quite skewed and hence deviates from Gaussianity. Nonetheless, SLdisco produced a reasonable CPDAG with only two definitely erroneous edges. Both of these edges were problematic due to their orientation, but plausible in terms of their adjacencies. This could suggest that Gaussianity is not a crucial assumption for SLdisco, at least when the variable distributions are unimodal as in our application. Applying SLdisco to more data with varying distributions will be helpful in gaining a better understanding of the influence of the Gaussianity assumption.

On a related note, our choice to use correlation matrices as input is of course only equivalent with using the full data under the assumption of Gaussianity, where the correlation matrix is a sufficient statistics (if we disregard the data scale). Our suggested procedure could easily be

modified to use full datasets, rather than correlation matrices, as its input. Using full datasets as inputs for SLdisco should be more robust towards deviations from Gaussianity and would allow SLdisco to utilize information in the data beyond partial correlations. Using full data would, however, impose a larger burden on computer memory, especially when n is large.

The assumption of no unobserved confounding, on the other hand, is potentially more crucial. This assumption is very strong and might not be fulfilled in many observational science applications. The CPDAGs proposed by GES and PC in the application may suggest that there are issues due to unobserved confounding in the Metropolit dataset. If so, it is interesting that SLdisco does not seem to be influenced by deviation from this assumption in the same way as PC and GES: It appears that SLdisco does not add extra spurious adjacencies – like the other two methods are known to do. A possible explanation for this could be that SLdisco is more robust towards unobserved confounding, but further work is needed to assess if this holds. In either case, extending SLdisco to also handle unobserved confounding could be achieved rather easily simply by simulating data in the presence of unobserved confounding. However, causal model equivalence classes would then have to be represented by partial ancestral graphs (PAGs) (instead of CPDAGs) (Richardson and Spirtes, 2002), which are more complicated graphical objects that include more edge types (bidirected, directed, undetermined, and combinations of the latter two). Expanding SLdisco to also handle data with unobserved confounding would therefore also require some modifications to the output layer of the proposed neural networks, and the orientation metrics used for evaluations would have to be generalized to handle more edge orientation possibilities.

Another limitation of our suggested machine learning procedure is our proposed neural network architecture. The architecture is clearly simplistic and we expect that performance can be greatly improved by substituting our approach with a more sophisticated and tailored architecture. Li et al. (2020) remark that the causal discovery task is *permutation equivariant* in the sense that permuting variable order in the input (i.e. permuting both rows and columns in the correlation matrix) results in the same permutation in the output (permuting both rows and columns in the adjacency matrix). They propose to use a equivariant neural network to natively support this feature, although with a different causal discovery task in mind (DAG discovery). Using such a permutation equivariant architecture in our setting could also be interesting. This would also ensure that the method becomes completely order invariant, which means that different orderings of inputted variables result in the same outputted adjacency matrix. This is an attractive and natural property to request of a causal discovery method, however, for some commonly used approaches, including PC, it is not completely fulfilled (Colombo et al., 2014).

We considered only a few, rather small, values of p (number of nodes). This was primarily in the interest of simplicity of disseminating the results. But we would like to make two points concerning the value of p : First, when the aim is conservative discovery for observational sciences, and hence the expected CPDAG of interest is quite dense, we believe that it is not relevant to consider very large values of p . A dense graph with e.g. 100 nodes would be completely uninterpretable and not aid applied scientists in understanding the causal mechanisms they wish to study. Secondly, when using convolutional neural networks, large values of p are not generally a computational issue; on the contrary, larger p would have allowed us to reuse existing network architectures designed for e.g., image classification, where p measures the number of pixels. We have performed some exploratory experiments reusing for example the MobileNet architecture (Howard et al., 2017), and while they produced slightly better performance for the “large” $p = 20$ graphs, they did not scale well to smaller graphs (results not included). Hence, we believe that the SLdisco procedure could be well-suited for much larger p as well. However, this would then

require a larger training data set, as the search space of Gaussian CPDAGs would become much larger, because the number of DAGs grows super exponentially with p (Peters et al., 2017).

We found that the post-processing method used to convert probability matrices to (pseudo) adjacency matrices was not very influential on the overall performance. The BPCO, which ensures that the output is a proper CPDAG, generally obtained slightly worse performance than the cutoff method for the adjacency metrics, but comparable or better performance in the orientation metrics. It may be possible to construct a better post-processing method since there are no guarantees that our proposed greedy method will find an optimal solution from the estimated probability matrix. However, we believe that a more useful line of research would be to natively include CPDAG structure requirements in the neural network optimization, so that no or little post-processing becomes necessary. This could be done by changing the loss function used to optimize the neural network. The current choice of loss function, the binary cross-entropy loss, essentially considers each entry in the output matrix one by one and seeks to optimize it locally, without taking the structure of the full output matrix into account. Informally, this loss is not aware of the fact that it should be looking for a CPDAG. Naturally, if the neural network is sufficiently sophisticated and $n \rightarrow \infty$, $b_{\text{train}} \rightarrow \infty$, optimizing the adjacency matrix entry-by-entry will converge to the correct solution. But with finite n and b_{train} , we suspect much can be gained by choosing a loss function that better captures the estimation goal. However, constructing such a function is not straightforward. Neural network weights are optimized using a variation of gradient descent and hence the loss function needs to be differentiable, and there has to exist analytical expressions for its derivatives, as they are computed a very large number of times. We leave suggestions for candidates for such a loss function to future research.

A strength of SLdisco is the very low computational burden at inference time, that is, when the method is used on real data. A pretrained SLdisco model provides a mapping from input data to CPDAGs that does not need further optimization. SLdisco will therefore be easy to combine with e.g. stability selection or other bootstrapping procedures to provide more insights into the statistical uncertainty of causal discovery – a topic that is currently heavily under-described. At training time, SLdisco does require some computation time (approximately 100–200 seconds per epoch resulting in 4–8 hours for the models presented here). However, it is only necessary to train a neural network once for each n - p combination. Moreover, transfer learning across similar values of n may be permissible so that new models need to be trained less often. Note also that more efficient training schemes may be achievable by e.g. training first on a large n and then only fine-tuning weights when considering a different value of n . For p , such transfer learning is less natural, as we do not expect the mapping from correlation matrices to adjacency matrices to be the same for different values of p . This may be tested empirically in future studies, but requires that a given neural network can input correlation matrices of varying sizes. One way to obtain this feature is by simply zero-padding or up-sampling the input matrix in the first layers of the architecture. Alternatively, if for example graph neural networks are used, this feature is ensured natively (Scarselli et al., 2008). This neural network type may also be able to avoid two other issues raised here, lack of permutation equivariance and the need of post-processing. We believe that applying graph neural network in SLdisco will be an interesting and natural next step in further improving the method we have proposed here.

Supplementary Material

The following supplementary materials are available:

Appendices:

- A:** Terminology and notation
- B:** Details about data simulation
- C:** Details about the neural network
- D:** Results: Extra figures
- E:** Application: Extra figures

Application data: Correlation matrices from the Metropolit data application and estimated adjacency matrices.

GES simulation study results: Estimated adjacency matrices from the GES applications for the simulation study (estimated using TETRAD).

Neural network models: Trained neural network models from the simulation study (.h5 files).

Replication code: R code for replicating the simulation study and the application.

Acknowledgments

We thank Dan Saattrup Nielsen for guidance and help with the machine learning procedures. We thank Kun Zhang for his contributions to designing the project.

Funding

This work was funded by the Independent Research Fund Denmark (grant 8020-00031B) and the National Institutes of Health (grant AWD000044521364471).

References

- Allaire J, Chollet F (2021). *keras: R Interface to ‘Keras’*. R package version 2.4.0.
- Chickering DM (2002). Optimal structure identification with greedy search. *Journal of Machine Learning Research*, 3: 507–554.
- Colombo D, Maathuis MH, et al. (2014). Order-independent constraint-based causal structure learning. *Journal of Machine Learning Research*, 15(1): 3741–3782.
- De Stavola BL, Nitsch D, dos Santos Silva I, McCormack V, Hardy R, Mann V, et al. (2006). Statistical issues in life course epidemiology. *American Journal of Epidemiology*, 163(1): 84–96. <https://doi.org/10.1093/aje/kwj003>.
- Goudet O, Kalainathan D, Caillou P, Guyon I, Lopez-Paz D, Sebag M (2018). Learning functional causal models with generative neural networks. In: *Explainable and Interpretable Models in Computer Vision and Machine Learning* (HJ Escalante, S Escalera, I Guyon, X Baró, Y Güçlütürk, U Güçlü, M van Gerven, R van Lier, eds.), 39–80. Springer.
- Greenland S, Pearl J, Robins JM (1999). Causal diagrams for epidemiologic research. *Epidemiology*, 10(1): 37–48. <https://doi.org/10.1097/00001648-199901000-00008>.
- Hornik K, Stinchcombe M, White H (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5): 359–366. [https://doi.org/10.1016/0893-6080\(89\)90020-8](https://doi.org/10.1016/0893-6080(89)90020-8).
- Howard AG, Zhu M, Chen B, Kalenichenko D, Wang W, Weyand T, et al. (2017). Mobilenets: Efficient convolutional neural networks for mobile vision applications. arXiv preprint: <https://arxiv.org/abs/1704.04861>.
- Kalisch M, Bühlman P (2007). Estimating high-dimensional directed acyclic graphs with the pc-algorithm. *Journal of Machine Learning Research*, 8(3): 613–636.

- Kalisch M, Mächler M, Colombo D, Maathuis MH, Bühlmann P (2012). Causal inference using graphical models with the R package pcalg. *Journal of Statistical Software*, 47(11): 1–26. <https://doi.org/10.18637/jss.v047.i11>.
- Ke NR, Chiappa S, Wang J, Bornschein J, Weber T, Goyal A, et al. (2022). Learning to induce causal structure. arXiv preprint: <https://arxiv.org/abs/2204.04875>.
- Li H, Xiao Q, Tian J (2020). Supervised whole dag causal discovery. arXiv preprint: <https://arxiv.org/abs/2006.04697>.
- Meek C (1995). Causal inference and causal explanation with background knowledge. In: *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence, UAI'95*, 403–410. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- Ogarrio JM, Spirtes P, Ramsey J (2016). A hybrid causal search algorithm for latent variable models. In: *Conference on Probabilistic Graphical Models*, 368–379. PMLR.
- Osler M, Andersen AMN, Lund R, Batty GD, Hougaard CØ, Damsgaard MT, et al. (2004). Revitalising the metropolit 1953 danish male birth cohort: Background, aims and design. *Paediatric and Perinatal Epidemiology*, 18(5): 385–394. <https://doi.org/10.1111/j.1365-3016.2004.00584.x>.
- Osler M, Lund R, Kriegbaum M, Christensen U, Andersen AMN (2006). Cohort profile: The metropolit 1953 Danish male birth cohort. *International Journal of Epidemiology*, 35(3): 541–545. <https://doi.org/10.1093/ije/dyi300>.
- Pearl J (2009). *Causality*. Cambridge university press.
- Perkovic E (2020). Identifying causal effects in maximally oriented partially directed acyclic graphs. In: *Conference on Uncertainty in Artificial Intelligence*, 530–539. PMLR.
- Peters J, Janzing D, Schölkopf B (2017). *Elements of Causal Inference: Foundations and Learning Algorithms*. The MIT Press.
- Petersen AH (2022). *causalDisco: Tools for Causal Discovery on Observational Data*. R package version 0.9.1.
- Petersen AH, Osler M, Ekstrøm CT (2021). Data-driven model building for life-course epidemiology. *American Journal of Epidemiology*, 190(9): 1898–1907. <https://doi.org/10.1093/aje/kwab087>.
- Ramsey J, Glymour M, Sanchez-Romero R, Glymour C (2017). A million variables and more: The fast greedy equivalence search algorithm for learning high-dimensional graphical causal models, with an application to functional magnetic resonance images. *International Journal of Data Science and Analytics*, 3(2): 121–129. <https://doi.org/10.1007/s41060-016-0032-z>.
- Ramsey JD, Zhang K, Glymour M, Romero RS, Huang B, Ebert-Uphoff I, et al. (2018). Tetrad—a toolbox for causal discovery. In: *8th International Workshop on Climate Informatics*.
- Richardson T, Spirtes P (2002). Ancestral graph Markov models. *The Annals of Statistics*, 30(4): 962–1030. <https://doi.org/10.1214/aos/1031689015>.
- Scarselli F, Gori M, Tsoi AC, Hagenbuchner M, Monfardini G (2008). The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1): 61–80. <https://doi.org/10.1109/TNN.2008.2005605>.
- Shah RD, Peters J (2020). The hardness of conditional independence testing and the generalised covariance measure. *The Annals of Statistics*, 48(3): 1514–1538. <https://doi.org/10.1214/19-AOS1857>.
- Spirtes P, Glymour C (1991). An algorithm for fast recovery of sparse causal graphs. *Social Science Computer Review*, 9(1): 62–72. <https://doi.org/10.1177/089443939100900106>.
- Spirtes P, Glymour CN, Scheines R, Heckerman D (2000). *Causation, Prediction, and Search*.

MIT press.

- Von Stumm S, Plomin R (2015). Socioeconomic status and the growth of intelligence from infancy through adolescence. *Intelligence*, 48: 30–36. <https://doi.org/10.1016/j.intell.2014.10.002>.
- Xu C, Xu W (2021). Causal structure learning with one-dimensional convolutional neural networks. *IEEE Access*, 9: 162147–162155. <https://doi.org/10.1109/ACCESS.2021.3133496>.
- Yu Y, Chen J, Gao T, Yu M (2019). Dag-gnn: Dag structure learning with graph neural networks. In: *International Conference on Machine Learning*, 7154–7163. PMLR.
- Zheng X, Aragam B, Ravikumar P, Xing EP (2018). Dags with no tears: Continuous optimization for structure learning. arXiv preprint: <https://arxiv.org/abs/1803.01422>.
- Zheng X, Dan C, Aragam B, Ravikumar P, Xing E (2020). Learning sparse nonparametric dags. In: *International Conference on Artificial Intelligence and Statistics*, 3414–3425. PMLR.