
Integer Programming for Causal Structure Learning in the Presence of Latent Variables

Rui Chen¹ Sanjeeb Dash² Tian Gao²

Abstract

The problem of finding an ancestral acyclic directed mixed graph (ADMG) that represents the causal relationships between a set of variables is an important area of research on causal inference. Most existing score-based structure learning methods focus on learning directed acyclic graph (DAG) models without latent variables. A number of score-based methods have recently been proposed for the ADMG learning, yet they are heuristic in nature and do not guarantee an optimal solution. We propose a novel exact score-based method that solves an integer programming (IP) formulation and returns a score-maximizing ancestral ADMG for a set of continuous variables that follow a multivariate Gaussian distribution. We generalize the state-of-the-art IP model for DAG learning problems and derive new classes of valid inequalities to formulate an IP model for ADMG learning. Empirically, our model can be solved efficiently for medium-sized problems and achieves better accuracy than state-of-the-art score-based methods as well as benchmark constraint-based methods.

1. Introduction

Causal graphs are graphical models representing dependencies and causal relationships between a set of variables. One class of the most common causal graphs, often known as Bayesian networks (BNs), is modeled by directed acyclic graphs (DAGs) in which a direct causal relationship between two nodes is indicated by a directed edge. However, its structure learning problem is NP-hard (Chickering et al., 2004). Many exact and approximate algorithms for learning BN structures from data have been developed, including

¹Department of Industrial and Systems Engineering, University of Wisconsin-Madison, Madison, Wisconsin, USA ²IBM Research, Yorktown Heights, New York, USA. Correspondence to: Rui Chen <rchen234@wisc.edu>.

score-based and constraint-based approaches. Score-based methods use a score – such as Bayesian scores or Bayesian information criterion (BIC) – to measure the goodness of fit of different graphs over the data, and then use a search procedure – such as hill-climbing (Heckerman et al., 1995; Tsamardinos et al., 2006a; Gámez et al., 2011), forward-backward search (Chickering, 2002), dynamic programming (Singh & Moore, 2005; Silander & Myllymaki, 2006; Gao & Wei, 2018), A* (Yuan & Malone, 2013) or integer programming (Jaakkola et al., 2010; Cussens, 2011; Cussens et al., 2016) – in order to find the best graph. On the other hand, constraint-based structure learning algorithms use (conditional) independence tests to decide on the existence of edges between all pairs of variables. Popular algorithms include the SGS algorithm (Spirtes et al., 2000), PC algorithm (Spirtes et al., 2000), and IC algorithm (Pearl, 2000).

Despite their wide application (Pearl, 2000), it is known that DAG models are not closed under marginalization (Tsirlis et al., 2018). This implies that DAGs cannot be used to model structures involving latent variables which are the most common cases in practice. For example, in healthcare domains, there may be numerous unobserved factors such as gene expression. Self-reported family history and diets may also leave out some important information. Ancestral graphs (AGs) were proposed as a generalization of DAGs (Richardson & Spirtes, 2002). AG models include all DAG models and are closed under marginalization and conditioning. AGs capture the independence relations among observed variables without explicitly including latent variables in the model. In this work, we assume no selection biases (no undirected edges in AGs). Hence there are two types of edges in the AGs: *directed* edges represent direct or ancestral causal relationships between variables and *bidirected* edges represent latent confounded relationships between variables. A bidirected edge between two nodes means that there exists at least one latent confounder that causes both nodes, and neither node causes the other node. AGs considered in this work are also called ancestral acyclic directed mixed graphs (ADMGs) (Bhattacharya et al., 2020) in the literature.

Causal structure learning algorithms for ancestral ADMGs can also be divided into two main classes: constraint-based and score-based methods. Constraint-based methods apply

conditional independence tests on the data to infer graph structures. Score-based methods search through possible graph structures to optimize a criterion for model selection. There are also hybrid methods that use both conditional independent tests and some scoring criterion that measures the likelihood of the data. In the setting of learning DAGs from observational data, score-based methods often achieve better performance than constraint-based methods as score-based methods are less sensitive to error propagation (Spirtes, 2010). For ADMG learning, existing methods are mostly constraint-based, including the FCI algorithm (Spirtes et al., 2000; Zhang, 2008) and the conservative FCI (cFCI) (Ramsey et al., 2012) to name a couple. Several score-based or hybrid approaches have been proposed for ancestral ADMG structure learning over continuous Gaussian variables in recent years (Triantafillou & Tsamardinos, 2016; Tsirlis et al., 2018; Bernstein et al., 2020; Chobtham & Constantinou, 2020) but are all greedy or local search algorithms. In this paper, we close the gap and propose an exact score-based solution method based on an integer programming (IP) formulation for ancestral ADMG learning. Our method is inspired by existing DAG learning IP formulations, and we derive new classes of valid inequalities to restrict the learned graph to be a valid ADMG. Empirical evaluation shows that the proposed method outperforms existing score-based DAG and ADMG structure learning algorithms.

The paper is organized as follows. In Section 2, we define the ancestral ADMG learning problem. In Section 3, we propose an integer programming formulation for obtaining the score-maximizing ancestral ADMG. In Section 4, we present experiments to compare our method with existing baselines.

2. Ancestral ADMG Learning

2.1. Preliminaries

We briefly review some related concepts. DAGs are directed graphs without directed cycles. A directed mixed graph $\mathcal{G} = (V, E_d, E_b)$ consists of a set of nodes V , a set of directed edges (\rightarrow) $E_d \subseteq \{(i, j) : i, j \in V, i \neq j\}$, and bidirected edges (\leftrightarrow) $E_b \subseteq \{\{i, j\} : i, j \in V, i \neq j\}$ between certain pairs of nodes. Given a directed edge $a \rightarrow b$, b is the head node, and a the tail node. We call node a an ancestor of node b in \mathcal{G} if there is a directed path from a to b in \mathcal{G} or $a = b$. We call node a a spouse (resp., parent) of node b in \mathcal{G} if there is a bidirected edge between a and b (resp., a directed edge from a to b) in \mathcal{G} . We denote the set of ancestors, the set of parents, and the set of spouses of node a in \mathcal{G} by $\text{an}_{\mathcal{G}}(a)$, $\text{pa}_{\mathcal{G}}(a)$, and $\text{sp}_{\mathcal{G}}(a)$, respectively. For $i \in V$ and $W \subseteq V$, $W \rightarrow i$ denotes that W is the parent set of i . A directed mixed graph \mathcal{G} is called an *ancestral ADMG* if the following condition holds for all pairs of nodes a and b in \mathcal{G} :

- If $a \neq b$ and $b \in \text{an}_{\mathcal{G}}(a) \cup \text{sp}_{\mathcal{G}}(a)$, then $a \notin \text{an}_{\mathcal{G}}(b)$.

In other words, \mathcal{G} is an ancestral ADMG if it contains no directed cycles ($a \rightarrow c \rightarrow \dots \rightarrow b \rightarrow a$) or almost directed cycles. An almost directed cycle is of the form $a \rightarrow c \rightarrow \dots \rightarrow b \leftrightarrow a$; in other words, $\{a, b\} \in E_b$ is a bidirected edge, and $a \in \text{an}_{\mathcal{G}}(b)$. Given a directed mixed graph \mathcal{G} , the *districts* define a set of equivalence classes of nodes in \mathcal{G} . The district for node a is defined as the connected component of a in the subgraph of \mathcal{G} induced by all bidirected edges, i.e.,

$$\{b : b \leftrightarrow \dots \leftrightarrow a \text{ in } \mathcal{G} \text{ or } a = b\}.$$

Given a district D of \mathcal{G} , the directed mixed graph \mathcal{G}_D is a subgraph of \mathcal{G} defined as follows. The node set of \mathcal{G}_D consists of all nodes in D and their parents. The bidirected edges in \mathcal{G}_D are the bidirected edges in \mathcal{G} that connect nodes in D . The directed edges in \mathcal{G}_D are the directed edges in \mathcal{G} where the head nodes are contained in D . We say that \mathcal{G}_D is the subgraph *implied* by D . In this paper, we call subgraph C of \mathcal{G} a *c-component* if C is a subgraph implied by some district of \mathcal{G} . Note that if a c-component has no bidirected edges, then any district must consist of a single node, and the c-component consists of a number of directed edges all with the same head node. Some authors use c-component as a synonym for district.

2.2. Calculating Graph Scores

We assume the ground truth data is modeled by the following linear equations:

$$X = MX + \epsilon.$$

Here M is a $d \times d$ matrix, $X = (X_1, \dots, X_d)^T$ are variables of the model, and $\epsilon = (\epsilon_1, \dots, \epsilon_d)^T$ is the error vector which follows a multivariate Gaussian distribution $N(0, \Sigma)$ ¹. Neither M nor Σ can be observed. Our goal is to find the best-fitting ancestral ADMG \mathcal{G} and an associated parameterization (M, Σ) satisfying $M_{ij} = 0$ if $i = j$ or $j \rightarrow i$ is not in \mathcal{G} and $\Sigma_{ij} = 0$ if $i \leftrightarrow j$ is not in \mathcal{G} . A score of the graph measures how well the graph structure represents the data. We use BIC (Schwarz, 1978) scores for all graphs in this paper. The BIC score for graph \mathcal{G} is given by

$$\text{BIC}_{\mathcal{G}} = 2 \ln(L_{\mathcal{G}}(\hat{\Sigma})) - \ln(N)(2|V| + |E|).$$

Here $\hat{\Sigma}$ is the maximum likelihood estimate of Σ given the graph representation \mathcal{G} and $\ln(L_{\mathcal{G}}(\hat{\Sigma}))$ is the associated log-likelihood while N denotes the number of samples, $|V|$ and $|E|$ denote the number of nodes and the number of edges in \mathcal{G} , respectively. Given a fixed graph and the empirical

¹We focus on the multivariate Gaussian case in this work as the scoring function can be factorized, although our work may also apply to discrete cases per different factorization rules (Richardson, 2009; Evans & Richardson, 2014).




Figure 1: Decomposition of the BIC score for an ancestral ADMG.

covariance matrix Q , the maximum likelihood estimate of the parameters can be found by applying the residual iterative conditional fitting algorithm in (Drton et al., 2009). According to (Nowzohour et al., 2017), $\ln(L_{\mathcal{G}}(\hat{\Sigma}))$ can be decomposed by c-components in \mathcal{G} . Specifically, let \mathcal{D} denote all districts of \mathcal{G} . Then

$$\begin{aligned} \ln(L_{\mathcal{G}}(\hat{\Sigma})) = & \\ & -\frac{N}{2} \sum_{D \in \mathcal{D}} \left[|D| \ln(2\pi) + \log\left(\frac{|\hat{\Sigma}_{\mathcal{G}_D}|}{\prod_{j \in \text{pa}_{\mathcal{G}}(D)} \hat{\sigma}_{D_j}^2}\right) + \right. \\ & \left. \frac{N-1}{N} \text{tr}(\hat{\Sigma}_{\mathcal{G}_D}^{-1} Q_D - |\text{pa}_{\mathcal{G}}(D) \setminus D|) \right]. \quad (1) \end{aligned}$$

Here $\text{pa}_{\mathcal{G}}(D)$ denotes the union of parent sets of nodes in D , $\hat{\Sigma}_{\mathcal{G}_D}$ is the maximum log-likelihood for subgraph \mathcal{G}_D and $\hat{\sigma}_{D_j}^2$ denotes diagonal entry of $\hat{\Sigma}_{\mathcal{G}_D}$ corresponding to node j . Note that the districts partition the nodes of \mathcal{G} , and the edges of \mathcal{G} are partitioned by the subgraphs \mathcal{G}_D . The BIC score for graph \mathcal{G} can be expressed as a sum of local scores of its c-components. For example, in Figure 1, the BIC score of the ancestral ADMG is equal to the sum of local scores of four c-components represented by different colors. Nodes with the same color belong to the same district and directed edges indicate their parents. When the graph is a DAG, other decomposable scoring functions can also be used (Silander et al., 2008).

Given the above decomposition property of BIC scores, one can immediately formulate the ancestral ADMG learning problem as an optimization problem, where one takes as input the set \mathcal{C} of all possible candidate c-components defined on the input random variables, and considers each subset of \mathcal{C} where the union of edges in the c-components in the subset form an ancestral ADMG. This is however impractical as the size of \mathcal{C} grows exponentially with increasing number of nodes/random variables. In a similar fashion, the DAG learning problem can be solved by considering candidate c-components where each c-components is implied by a single node district. Even with this restriction, the number of c-components that need to be considered in DAG learning is exponential, and a common approach to deal with this issue is to only consider c-components with a bounded number of nodes. The model we propose for ADMG learning is defined over general c-components, although in our computational experiments we restrict the set of candidate

c-components to allow our model to be solved within a reasonable amount of time. In Section 2.3, we further consider eliminating sub-optimal c-components, which significantly reduces the size of the search space in some cases.

2.3. Pruning the list of candidate c-components

The set of ancestral ADMGs defined over a set of nodes is closed under deletion of directed and bidirected edges. Similarly, the operation of removing all parents of a node and replacing all bidirected edges incident to the node with outgoing directed edges transforms an ancestral ADMG into another ancestral ADMG. Both of these operations also result in the transformation of the c-components associated with an ADMG. If the c-components resulting from applying the above transformation to another c-component have a higher combined score than the original c-component, then we can assert that the original c-component need never be considered in an optimal solution of the ADMG learning problem. We use the above ideas and their hybrids to define two ways of pruning \mathcal{C} in our implementation. Firstly, if removing one or more edges from a c-component results in a set of c-components with a higher combined score, we prune the original c-component. For example, the c-component $\{A\} \rightarrow B \leftrightarrow C \leftarrow \{D\}$ can be pruned if it has a lower score than the sum of the scores of $B \leftarrow \{A\}$ and $C \leftarrow \{D\}$. Similarly, we prune a c-component if the second operation above leads to c-components with a higher combined score. For example, we can prune $\{A\} \rightarrow B \leftrightarrow C \leftarrow \{D\}$ if it has a score lower than the sum of the scores of $B \leftarrow \emptyset$ and $C \leftarrow \{B, D\}$. A hybrid of the previous two pruning ideas is also implemented. For example, we can combine both operations to transform $\{A\} \rightarrow B \leftrightarrow C \leftarrow \{D\}$ into the c-components $B \leftarrow \emptyset$ and $C \leftarrow \{B\}$.

3. An Integer Programming Formulation

Integer programming (IP) is a mathematical optimization tool for modeling and solving optimization problems involving variables that are restricted to discrete values. For any c-component C implied by a district, let D_C denote the implying district, let E_C denote the pairs of nodes in D_C that are connected by a bidirected edge, and for each node $i \in D_C$, let $W_{C,i}$ denote the parent set of node i in c-component C . Given a set of c-components \mathcal{C} , finding a score-maximizing ancestral ADMG with all its district-implied subgraphs in \mathcal{C} can be straightforwardly formulated as an integer program as follows:

$$\max_{z \in \{0,1\}^{\mathcal{C}}} \sum_{C \in \mathcal{C}} s_C z_C \quad (2)$$

$$\text{s.t.} \quad \sum_{C: i \in D_C} z_C = 1, i \in \{1, \dots, d\} \quad (3)$$

$$\mathcal{G}(z) \text{ is acyclic and ancestral.} \quad (4)$$




Figure 2: Depiction of the cluster inequality for $S = \{A, B, C, D\}$.

Here z_C is a binary variable indicating if the c-component C is chosen, s_C is the local score of c-component C and $\mathcal{G}(z)$ is the directed mixed graph whose district-implied subgraphs are exactly the c-components C with $z_C = 1$. Constraints (3) enforce the condition each node i is contained in a single district (that implies a chosen c-component). Constraint (4) implies that the resulting directed mixed graph $\mathcal{G}(z)$ contains no directed or almost directed cycles. Standard IP solvers can only deal with linear inequality constraints, and we next discuss how to represent constraint (4) by a set of linear inequality constraints.

3.1. Avoiding Directed Cycles

Several (mixed-)integer programming (MIP) formulations exist for constraining a graph to be acyclic. A description and a comparison of some of these different anti-cycle formulations for the DAG learning problem with continuous data, including the linear ordering formulation (Grötschel et al., 1985), the topological ordering formulation (Park & Klabjan, 2017), and the cycle elimination formulation, are given in (Manzour et al., 2020). A new layered network formulation was also proposed in (Manzour et al., 2020). One particular class of anti-cycle constraints called cluster inequalities/constraints was introduced in (Jaakkola et al., 2010) and often performs better than the other methods above and results in a very tight integer programming formulation for DAG learning problems. Let $I(\cdot)$ be a variable indicating whether or not the substructure described between the parentheses is present in the graph. Specifically, $I(W \rightarrow i)$ indicates whether or not $\text{pa}_{\mathcal{G}(z)}(i) = W$. Cluster constraints have the following form:

$$\sum_{i \in S} \sum_{W: W \cap S = \emptyset} I(W \rightarrow i) \geq 1, \quad \forall S \subseteq \{1, 2, \dots, d\}. \quad (5)$$

Inequalities (5) encode the constraint that for a set S of nodes in an acyclic graph, there must exist at least one node in S whose parent set has no intersection with S . In Figure 2, there is a directed cycle connecting nodes $\{A, B, C, D\}$, and it violates the cluster inequality for $S = \{A, B, C, D\}$,

which would be satisfied if node D has all its parents outside S . The generation of cluster constraints is essential for the state-of-the-art IP-based DAG learning solver GOBNILP (Bartlett & Cussens, 2017). The same inequalities can be applied to ancestral ADMG learning by introducing auxiliary variables $I(W \rightarrow i)$ defined as

$$I(W \rightarrow i) := \sum_{C \in \mathcal{C}: i \in D_C, W_{C,i} = W} z_C. \quad (6)$$

In our case, inequalities (5) can be further strengthened as some of the original variables z_C may be double counted when they are present in multiple auxiliary variables of the form $I(W \rightarrow i)$. We propose the following *strengthened cluster constraints* for ancestral ADMG learning:

$$\sum_{i \in S} \sum_{C \in \mathcal{C}: i \in D_C, W_{C,i} \cap S = \emptyset} z_C \geq 1, \quad \forall S \subseteq \{1, 2, \dots, d\}. \quad (7)$$

Proposition 1. *Inequalities (7) are satisfied by all solutions of the integer program (2)-(4). If $z \in \{0, 1\}^d$ satisfies (3) and (7), then z corresponds to an ADMG, i.e., $\mathcal{G}(z)$ contains no directed cycles.*

Proof. We first show the validity of (7). Let \bar{z} be a feasible solution of the integer program (2)-(4). Assume that inequality (7) with S set to \bar{S} is violated by \bar{z} , i.e.,

$$\sum_{i \in \bar{S}} \sum_{C \in \mathcal{C}: i \in D_C, W_{C,i} \cap \bar{S} = \emptyset} \bar{z}_C = 0. \quad (8)$$

By constraint (3), for each $i \in \{1, \dots, d\}$, there exists C_i satisfying $i \in D_{C_i}$ such that $\bar{z}_{C_i} = 1$. By (8), we have $W_{C_i, i} \cap \bar{S} \neq \emptyset$ for each $i \in \bar{S}$. It follows that each node i in \bar{S} has a parent in \bar{S} in $\mathcal{G}(\bar{z})$, which contradicts the feasibility of \bar{z} since $\mathcal{G}(\bar{z})$ contains a directed cycle with nodes in \bar{S} .

Next we show that $\mathcal{G}(z)$ is an ADMG if z satisfies (3) and (7). Constraints (3) imply that $\mathcal{G}(z)$ is a valid directed mixed graph as exactly one c-component having i in its implying district is active. We only need to show that $\mathcal{G}(z)$ contains no directed cycles. Assume for contradiction that $\mathcal{G}(z)$ contains a directed cycle $i_0 \rightarrow i_1 \rightarrow \dots \rightarrow i_K \rightarrow i_0$. Let $S' = \{i_0, i_1, \dots, i_K\}$. By constraint (3), for each $i \in \{1, \dots, d\}$, there exists exactly one C_i satisfying $i \in D_{C_i}$ such that $z_{C_i} = 1$. The directed cycle $i_0 \rightarrow i_1 \rightarrow \dots \rightarrow i_K \rightarrow i_0$ implies that $W_{C_i, i} \cap S' \neq \emptyset$ for each $i \in S'$. Therefore, by constraint (3), $\sum_{C \in \mathcal{C}: i \in D_C, W_{C,i} \cap S' = \emptyset} z_C = 0$ for each $i \in S'$, which contradicts inequality (7) with $S = S'$. \square

3.2. Avoiding Almost Directed Cycles

Ancestral ADMGs are ADMGs without almost directed cycles. To encode the constraint that the graph contains no almost directed cycles, we define auxiliary variables




Figure 3: Depiction of the bicluster inequality for $S = \{A, B, C\}$ and $(i, j) = (B, C)$.

$I(W^1 \rightarrow i, W^2 \rightarrow j, i \leftrightarrow j)$ and $I(i \leftrightarrow j)$ in a manner similar to (6), and give the following inequalities for all $S \subseteq \{1, 2, \dots, d\}$ with $i, j \in S$ and $i < j$:

$$\sum_{W^1: W^1 \cap S = \emptyset} \sum_{W^2: W^2 \cap S = \emptyset} I(W^1 \rightarrow i, W^2 \rightarrow j, i \leftrightarrow j) + \sum_{v \in S \setminus \{i, j\}} \sum_{W: W \cap S = \emptyset} I(W \rightarrow v) \geq I(i \leftrightarrow j). \quad (9)$$

We refer to inequalities (9) as the bicluster inequalities/constraints. Bicluster inequalities encode the constraint that if a bidirected edge $i \leftrightarrow j$ is present in some ancestral ADMG, then, for any set of nodes S containing i and j , either some node $v \in S \setminus \{i, j\}$ has all its parents outside of S or both parent sets of i and j have no intersection with S . As shown in Figure 3, where B and C are connected by a bidirected edge, either the parents of A must lie outside $S = \{A, B, C\}$, or the parents of both B and C must lie outside S . Similar to cluster inequalities, bicluster inequalities can also be strengthened when written in the original z variable space as some c-component variables are double-counted on the left-hand side of (9). Also some c-component variables might be contradicting the presence of the bidirected edge $i \leftrightarrow j$, and therefore cannot be active when $I(i \leftrightarrow j)$ is active and should be removed from the left-hand side.

Proposition 2. *The following inequalities are valid for the integer program (2)-(4):*

$$\sum_{C \in \mathcal{C}(S; \{i, j\})} z_C \geq I(i \leftrightarrow j), \quad \forall S \subseteq \{1, 2, \dots, d\}$$

with $i \in S, j \in S, i < j$ (10)

where

$$\begin{aligned} \mathcal{C}(S; \{i, j\}) = & \{C \in \mathcal{C} : \{i, j\} \in E_C, (W_{C,i} \cup W_{C,j}) \cap S = \emptyset\} \cup \\ & \{C \in \mathcal{C} : |\{i, j\} \cap D_C| \neq 1, \exists v \in D_C \cap S \setminus \{i, j\} \\ & \text{s.t. } W_{C,v} \cap S = \emptyset\}. \end{aligned}$$

Any $z \in \{0, 1\}^d$ satisfying (3), (7) and (10) corresponds to an ancestral ADMG.

Proof. We first show validity of (10). Let \bar{z} be a feasible solution of (2)-(4). Assume for contradiction that inequality (10) with $S = \bar{S}$ and $(i, j) = (\bar{i}, \bar{j})$ is violated by \bar{z} . Note that the right-hand side value of (10) can only be binary and (10) cannot be violated if the right hand side is 0 since z is binary. Then there exists a bidirected edge between \bar{i} and \bar{j} in $\mathcal{G}(\bar{z})$. Since $z = \bar{z}$ violates (10) with $S = \bar{S}$ and $(i, j) = (\bar{i}, \bar{j})$, $\bar{z}_C = 0$ for each $C \in \mathcal{C}(\bar{S}; \{\bar{i}, \bar{j}\})$. Next we construct an almost directed cycle in $\mathcal{G}(\bar{z})$. Because variable $I(\bar{i} \leftrightarrow \bar{j})$ is active, there exists $C_{\bar{i}\bar{j}} \in \mathcal{C}$ such that $\{\bar{i}, \bar{j}\} \in E_{C_{\bar{i}\bar{j}}}$ and $\bar{z}_{C_{\bar{i}\bar{j}}} = 1$. Since $\bar{z}_C = 0$ for each $C \in \mathcal{C}(\bar{S}; \{\bar{i}, \bar{j}\})$, we have $(W_{C_{\bar{i}\bar{j}, \bar{i}} \cup W_{C_{\bar{i}\bar{j}, \bar{j}}}) \cap \bar{S} \neq \emptyset$. Without loss of generality, assume $W_{C_{\bar{i}\bar{j}, \bar{i}}} \cap \bar{S} \neq \emptyset$. Let $i_0 = \bar{i}$. We recursively define i_k as follows:

- Let i_k be a parent of i_{k-1} such that $i_k \in \bar{S}$.

Then this sequence either finds a directed cycle because of the finiteness of the number of nodes, which contradicts the feasibility of \bar{z} , or ends up with some node i_K that has no parent in \bar{S} . If it is the second case, we next show that i_K can only be j which implies an almost directed cycle $i_K \rightarrow i_{K-1} \rightarrow \dots \rightarrow i_0 \leftrightarrow i_K$ in $\mathcal{G}(\bar{z})$. First of all, $i_0 = \bar{i}$ has a parent in \bar{S} since $W_{C_{\bar{i}\bar{j}, \bar{i}}} \cap \bar{S} \neq \emptyset$. Let $\bar{C} \in \mathcal{C}$ be such that $i_K \in D_{\bar{C}}$ and $z_{\bar{C}} = 1$. Assume for contradiction that $i_K \neq j$. Then $i_K \in D_{\bar{C}} \cap S \setminus \{i, j\}$ and $W_{\bar{C}, i_K} \cap \bar{S} = \emptyset$. Also note that either $\{i, j\} \cap D_{\bar{C}} = \emptyset$ or $\{i, j\} \subseteq D_{\bar{C}}$ since otherwise it contradicts the activeness of $I(i \leftrightarrow j)$. This implies that $\bar{C} \in \mathcal{C}(\bar{S}; \{\bar{i}, \bar{j}\})$ which contradicts the fact that $\bar{z}_C = 0$ for each $C \in \mathcal{C}(\bar{S}; \{\bar{i}, \bar{j}\})$.

We next show that $\mathcal{G}(z)$ is an ancestral ADMG if z satisfies (3), (7) and (10). By Proposition 1, $\mathcal{G}(z)$ is an ADMG. We only need to show that $\mathcal{G}(z)$ contains no almost directed cycles. Assume for contradiction that $\mathcal{G}(z)$ contains an almost directed cycle $i_0 \rightarrow i_1 \rightarrow \dots \rightarrow i_K \leftrightarrow i_0$. Let $S' = \{i_0, i_1, \dots, i_K\}$. We can then easily verify that $z_C = 0$ for any $C \in \mathcal{C}(S'; \{i_0, i_K\})$ which contradicts (10) with $S = S', i = \min\{i_0, i_K\}$ and $j = \max\{i_0, i_K\}$. \square

We provide a simple example of a strengthened bicluster inequality when the problem contains only 3 variables.

Example 1. *Let $d = 3$ and \mathcal{C} be the collection of all possible c-components on three nodes. Consider $S = \{1, 2, 3\}$, $i = 2$ and $j = 3$. Then the associated strengthened bicluster inequality (after simplification) is*

$$z_{0 \rightarrow 1} \geq z_{\{1\} \rightarrow 2 \leftrightarrow 3 \leftarrow \{1\}} + z_{0 \rightarrow 2 \leftrightarrow 3 \leftarrow \{1\}} + z_{\{1\} \rightarrow 2 \leftrightarrow 3 \leftarrow \emptyset}. \quad (11)$$

Note that at most one of the c-components on the right-hand side of (11) ($\{1\} \rightarrow 2 \leftrightarrow 3 \leftarrow \{1\}, \emptyset \rightarrow 2 \leftrightarrow 3 \leftarrow \{1\}$)

and $\{1\} \rightarrow 2 \leftrightarrow 3 \leftarrow \emptyset$ can be active due to (3). Inequality (11) enforces the constraint that c -component $\emptyset \rightarrow 1$ must be active when one of the three c -components on the right-hand side of (11) is active.

So far we have given an integer programming formulation of the ancestral ADMG learning problem. However, solving this integer program efficiently is nontrivial. The first step to solve an integer program is often solving some polyhedral relaxation of the problem. Directly solving the relaxed problem over the polyhedron defined by (3), (7) and (10) is computationally infeasible because of the exponential number of constraints. A common practice for solving such problems with exponentially many constraints is to further relax the polyhedral relaxation to be the polyhedron defined by only some of these constraints, find an optimal solution, solve a corresponding *separation problem*, and add violated constraints to the relaxation if they exist. The separation problem finds constraints that are violated by the current solution and this process is repeated until convergence. If the optimal solution of the relaxed problem is integral, such a separation problem translates to just finding a directed or almost directed cycle in a directed mixed graph which can easily be accomplished by depth-first search. However, the separation problem can be much harder to solve when the solution is fractional. We provide separation heuristics based on variants of Karger’s random contraction algorithm (Karger, 1993) in the supplement.

4. Empirical Evaluation

We conduct a set of experiments to compare our IP model, AGIP, with existing state-of-the-art baselines. We use DAGIP to represent existing DAG IP models such as the one in GOBNILP (Bartlett & Cussens, 2017), which is the same IP model as AGIP when all candidate c -components are implied by single-node districts. The solution obtained from DAGIP is equivalent to any exact score-based method for generating DAG solutions. We also compare with non-IP-based approaches, namely M^3HC (Tsirlis et al., 2018), FCI (Spirtes et al., 2000; Zhang, 2008), and cFCI (Ramsey et al., 2012).

To measure solution quality, we use a few different metrics. When comparing against score-based methods, our optimization model objective is to maximize a score, such as the BIC score (Schwarz, 1978) for model selection, and the metric is the solution score. To compare with constraint-based methods which do not have objective scores, the solution graph is converted to a partial ancestral graph (PAG), which characterizes a class of Markov equivalent AG solutions, and then compared with the ground truth PAG. We use structural Hamming distance (SHD) (Tsamardinos et al., 2006b), which is the number of edge operations (addition or deletion of an undirected edge, addition, removal or reversion

of the orientation of an edge) between the predicted graph and the ground truth graph. Finally, precision and recall (Tillman & Spirtes, 2011) are also compared. They are the number of correct edges with correct orientations in the predicted graph divided by the number of edges in the predicted graph and by the number of edges in the ground truth graph, respectively.

All experiments are run on a Windows laptop with 16GB RAM and an Intel Core i7-7660U processor running at 2.5GHz. Integer programs are solved using the optimization solver Gurobi 9.0.3.

4.1. Experiment 1: Exact Graph Recovery




Figure 4: Ground truth DAG (left), AG (middle), and PAG (right) of the four-node graph.

We first test on a four-node example, where the data (10000 data points) is simulated from a five-node DAG model (see Figure 4) with node 0 being unobserved. The purpose of this experiment is to show that for small graphs, where we can practically enumerate and use all possible c -component variables, and with enough samples, exact graph recovery is possible with AGIP. We test score-based methods AGIP, DAGIP and M^3HC on this example. In AGIP, we consider all possible c -components with arbitrary sizes. In DAGIP, we consider all possible c -components implied by single-node districts.




Figure 5: Solutions obtained from AGIP (left), DAGIP (middle) and M^3HC (right).

These three methods generate three different solutions. Comparing the scores of the solutions, we observe that $\text{score}(\text{AGIP}) > \text{score}(\text{DAGIP}) > \text{score}(\text{M}^3\text{HC})$ in this example. Both AGIP and M^3HC correctly capture the skeleton (presence of edges between each pair of nodes) of the ground truth AG. Only the AGIP solution is Markov equivalent to the ground truth AG, i.e., the AGIP solution encodes

the same set of conditional independence relationships as the ground truth AG. This result shows that our method is consistent for large samples and finds the exact solution.

4.2. Experiment 2: Random Graphs

We further experiment on a set of randomly generated ADMGs following the procedure described in (Triantafillou & Tsamardinos, 2016). For each instance, a random permutation is applied to $d + l$ variables in a DAG as the ordering of the variables. For each variable i , a set of up to 3 variables that have higher ordering than i is randomly chosen as the parent set of i . The resulting DAG is then assigned a randomly generated conditional linear Gaussian parameterization. Within the $d + l$ variables in the DAG, l of them are randomly chosen as latent variables and marginalized which result in an AG over the observed d variables. For each such graphical model, a sample of $N \in \{1000, 10000\}$ realizations of the observed variables is simulated to create the instance. For fixed (d, l, N) , 10 instances with parameters d , l , and N are generated.

4.2.1. COMPARISON BETWEEN AGIP, DAGIP AND M³HC SOLUTIONS

To guarantee efficiency, we restrict the sizes of c-components considered in the AGIP and DAGIP methods. AGIP considers c-components implied by a single-node district with up to 3 parents or a two-node district with up to 1 parent per node, while DAGIP considers c-components implied by a single-node district with up to 3 parents.

We want to emphasize that ADMG learning can be much harder than DAG learning. In the state-of-the-art DAG learning IP model (Bartlett & Cussens, 2017), assuming n variables and parent sets with maximum size k , there are at least $n \binom{n-1}{k} = \Theta(n^{k+1})$ IP variables in total for fixed k but increasing n (before pruning). For the ADMG learning problem, assuming a maximum district size of p nodes, the number of IP variables in our IP model (AGIP) is at least $\binom{n}{p} \binom{n-p}{k}^p = \Theta(n^{p(k+1)})$ for fixed k, p but increasing n (before pruning). When $p = 2$ (the minimum required to model bidirected edges), AGIP has the square of the number of IP variables in DAGIP. With our setting of the experiments, AGIP has roughly double the IP variables of DAGIP. One possible way to deal with this explosive growth is to add a collection of potentially good c-components with large-districts to AGIP rather than all possible ones.

In these experiments, we fix $d + l = 20$ with varying $l \in \{2, 4, 6\}$. We first show the tightness of our IP formulation. For each combination of (d, l, N) , we report in Table 1 the average number of binary variables before and after pruning, average pruning time, average integrality gap at the root node, and average IP solution time for the AGIP model over the 10 instances. We observe that the solution time

decreases as l increases. On the other hand, the solution time increases when the number of samples N increases, since fewer c-component variables can be pruned during the pruning phase. For all cases, the IP formulation has a pretty small gap at the root node which is always below 1% on average.

We next compare the qualities of solutions obtained from AGIP, DAGIP, and M³HC in terms of the scores. Since the feasible region of the AGIP model is strictly larger than DAGIP, the optimal score obtained from AGIP is guaranteed to be at least as good as the optimal score obtained from DAGIP. In Table 2, we report the average difference (improvement) in score compared with M³HC and the number of instances where AGIP improves the score over DAGIP for each combination of (d, l, N) . Both AGIP and DAGIP produce solutions better than the greedy algorithm M³HC, although M³HC in principle searches over a larger space of graph structures. In fact, M³HC finds slightly better solutions for only 4 of the 60 instances we tested and performed strictly worse than AGIP on the other 53 instances. Therefore, we can conclude that the exact methods DAGIP and AGIP are better at obtaining solutions with higher scores than the greedy method M³HC on randomly generated instances. We also observe that AGIP improves over DAGIP on 13 of 60 instances. There are two particular explanations for this. Firstly, the implied ground truth AG might be Markov equivalent to a DAG in which case the DAG solution can be optimal. Secondly, the “non-DAG” candidate c-components are limited as AGIP only considers additionally c-components implied by a two-node district with up to 1 parent each node compared with DAGIP.

4.2.2. COMPARISON BETWEEN SCORE-BASED METHODS AND CONSTRAINT-BASED METHODS

We compare score-based methods AGIP and M³HC with constraint-based methods FCI and cFCI in Table 3. All methods have better precision and recall values when the sample size N increases. FCI seems to perform worse than the other 3 methods, and on average, AGIP has the best SHD, precision, and recall.

We observe that AGIP has a better recall in 5 out of 6 settings. AGIP has significantly better performance than the other methods when $(d, l, N) = (18, 2, 10000)$.

To see the impact of latent variables on the performance of these methods, we also regenerate instances with fixed $d = 18$ and $N = 10000$ but varying $l \in \{2, 4, 6\}$. For each l , we plot on Figure 6 the mean (\pm standard deviation) of SHD, precision, and recall of estimated graphs from each method over 10 instances (see the supplement for detailed results). The performance of each method tends to drop as the number of latent variables increases. The drop is most significant for AGIP, presumably due to the restriction on the

Table 1: Tightness of the AGIP formulation.

| (d, l, N) | Avg # bin vars before pruning | Avg # bin vars after pruning | Avg pruning time (s) | Avg root gap (%) | Avg solution time (s) |
|----------------|-------------------------------|------------------------------|----------------------|------------------|-----------------------|
| (18, 2, 1000) | 59229 | 4116 | 19.1 | 0.65 | 60.4 |
| (16, 4, 1000) | 39816 | 3590 | 13.6 | 0.43 | 41.0 |
| (14, 6, 1000) | 20671 | 1788 | 3.9 | 0.54 | 8.9 |
| (18, 2, 10000) | 59229 | 9038 | 33.0 | 0.67 | 323.2 |
| (16, 4, 10000) | 39816 | 7378 | 21.4 | 0.53 | 215.4 |
| (14, 6, 10000) | 20671 | 3786 | 6.4 | 0.56 | 47.2 |

Table 2: Comparing scores of AGIP, DAGIP, and M³HC.

| (d, l, N) | Avg improvement in score compared with M ³ HC | | # instances where AGIP improves over DAGIP in score |
|----------------|--|--------|---|
| | AGIP | DAGIP | |
| (18, 2, 1000) | 82.75 | 82.32 | 3/10 |
| (16, 4, 1000) | 90.03 | 89.33 | 5/10 |
| (14, 6, 1000) | 34.84 | 34.68 | 3/10 |
| (18, 2, 10000) | 373.44 | 373.44 | 0/10 |
| (16, 4, 10000) | 147.96 | 147.54 | 1/10 |
| (14, 6, 10000) | 150.52 | 150.44 | 1/10 |

Table 3: Comparison between AGIP, M³HC, FCI, and cFCI for ancestral ADMG Learning.

| (d, l, N) | SHD | | | | Precision (%) | | | | Recall (%) | | | |
|----------------|-------------|-------------------|------|-------------|---------------|-------------------|------|-------------|-------------|-------------------|------|------|
| | AGIP | M ³ HC | FCI | cFCI | AGIP | M ³ HC | FCI | cFCI | AGIP | M ³ HC | FCI | cFCI |
| (18, 2, 1000) | 36.0 | 32.0 | 32.9 | 25.9 | 41.7 | 37.6 | 30.9 | 49.5 | 41.9 | 30.8 | 24.7 | 38.5 |
| (16, 4, 1000) | 23.4 | 32.1 | 34.9 | 29.7 | 58.0 | 35.9 | 27.6 | 41.0 | 56.6 | 29.3 | 22.8 | 32.5 |
| (14, 6, 1000) | 30.1 | 34.7 | 34.1 | 30.6 | 43.2 | 36.0 | 26.3 | 42.0 | 37.0 | 24.8 | 18.7 | 30.0 |
| (18, 2, 10000) | 12.6 | 26.4 | 30.8 | 24.8 | 80.3 | 54.1 | 29.6 | 50.0 | 78.7 | 49.3 | 26.2 | 44.1 |
| (16, 4, 10000) | 24.6 | 23.9 | 26.7 | 23.1 | 60.6 | 56.3 | 43.7 | 51.1 | 58.4 | 51.6 | 38.5 | 45.6 |
| (14, 6, 10000) | 26.5 | 23.5 | 27.7 | 23.3 | 47.7 | 54.2 | 38.6 | 53.1 | 44.5 | 46.5 | 33.2 | 44.4 |
| overall | 25.5 | 28.8 | 31.2 | 26.2 | 55.3 | 45.7 | 32.8 | 47.8 | 52.8 | 38.7 | 27.4 | 39.2 |

c-component size. The ground truth graphs associated with our randomly generated instances can have c-components with large parental sizes and district sizes, especially in cases with more latent variables.

4.3. Experiment 3: Non-DAG ADMGs

Although AGIP generates a solution with the same or better score than DAGIP, it returns a strictly better score than DAGIP on a small fraction (13/60) of instances in Experiment 2. One possibility is that all ground-truth graphs are DAGs. We next compare AGIP and DAGIP on graphs that

are not "DAG-representable", i.e., the ground truth AG is not Markov-equivalent to any DAG. We randomly generate AGs with $(d, l) = (10, 10)$ following the scheme described in Section 4.2. We pick the first 5 AGs that have at least one bidirected edge in the corresponding PAGs (which implies they are not DAG-representable). All of these 5 AGs are provided in the supplement. For each of the 5 AGs, we generate 10 instances of 10000 realizations of the model, each with a different randomly generated conditional linear Gaussian parametrization. In addition to the original candidate c-components, we consider in AGIP c-components that are implied by three-node districts with up to 1 parent

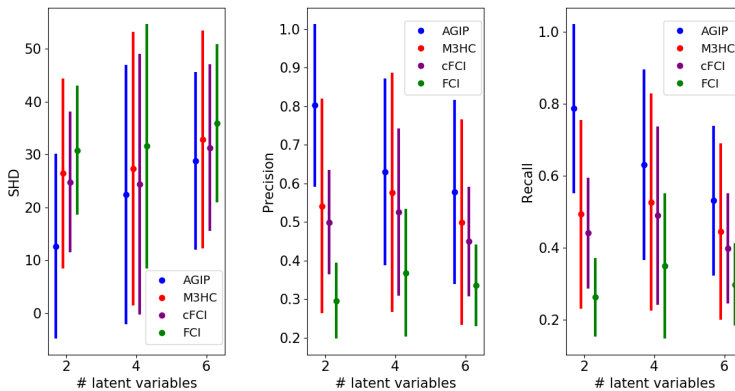


Figure 6: Performance of different methods when the number of latent variables increases.

Table 4: Comparison between AGIP and DAGIP on graphs that are not DAG-representable.

| Graph index | Avg SHD | | Avg precision (%) | | Avg recall (%) | | # instances where AGIP improves over DAGIP in score |
|-------------|-------------|-------------|-------------------|-------------|----------------|-------------|---|
| | AGIP | DAGIP | AGIP | DAGIP | AGIP | DAGIP | |
| 1 | 6.7 | 6.6 | 63.7 | 59.5 | 64.4 | 60.0 | 10/10 |
| 2 | 9.2 | 10.5 | 59.4 | 50.5 | 63.0 | 52.0 | 7/10 |
| 3 | 8.0 | 8.8 | 67.3 | 64.8 | 63.8 | 60.0 | 5/10 |
| 4 | 29.8 | 29.8 | 27.4 | 29.2 | 17.6 | 19.0 | 4/10 |
| 5 | 21.7 | 23.0 | 30.0 | 27.6 | 27.3 | 24.7 | 2/10 |
| overall | 15.1 | 15.7 | 49.6 | 46.3 | 47.2 | 43.1 | 28/50 |

for each node. Table 4 contains SHD, precision, and recall values of DAGIP and AGIP solutions.

AGIP has a strictly better score than DAGIP in 28 out of the 50 instances we considered. Note that if the scores for optimal AGIP and DAGIP solutions are identical (as in 22 instances), then it is not possible to predict which solution will perform better on metrics such as SHD, precision, or recall. AGIP performs better than DAGIP in precision and recall on the 1st, 2nd, 3rd, and 5th AGs, and performs better in SHD on the 2nd, 3rd, and 5th AGs. AGIP performs slightly worse than DAGIP on the 4th AG, which has a very large c-component containing 7 nodes in the district. All of the five AGs contain at least one c-component that is not covered by AGIP (nor by DAGIP). But considering more c-components does help improve the solution quality on the average, which illustrates the advantage of AGIP.

5. Conclusions and Future Work

We presented an integer-programming based approach for learning ancestral ADMGs from observational data. Our

main contributions are: 1) an IP formulation for the ancestral ADMG learning problem; 2) new classes of valid inequalities for efficient solution of the IP model; 3) numerical experiments to compare our model with existing methods. Empirical evaluation shows that our method has promising performance and can generate solutions with better accuracy than existing state-of-the-art learning methods. To our knowledge, this is the first exact score-based method for solving such problems in the presence of latent variables. For future work, extending the current approach to allow efficient solution with more c-components could further improve the solution quality of the proposed method. Adding other classes of valid inequalities to strengthen the current IP formulation is another direction worth exploring.

Acknowledgements

We thank Dr. James Luedtke for discussion on the problem formulation. We thank anonymous reviewers for their insightful comments.

References

- Bartlett, M. and Cussens, J. Integer linear programming for the bayesian network structure learning problem. *Artificial Intelligence*, 244:258–271, 2017.
- Bernstein, D., Saeed, B., Squires, C., and Uhler, C. Ordering-based causal structure learning in the presence of latent variables. In *International Conference on Artificial Intelligence and Statistics*, pp. 4098–4108, 2020.
- Bhattacharya, R., Nagarajan, T., Malinsky, D., and Shpitser, I. Differentiable causal discovery under unmeasured confounding. *arXiv preprint arXiv:2010.06978*, 2020.
- Chickering, D. M. Optimal structure identification with greedy search. *Journal of machine learning research*, 3: 507–554, 2002.
- Chickering, D. M., Heckerman, D., and Meek, C. Large-sample learning of Bayesian networks is NP-hard. *Journal of Machine Learning Research*, 5:1287–1330, 2004.
- Chobtham, K. and Constantinou, A. C. Bayesian network structure learning with causal effects in the presence of latent variables. *arXiv preprint arXiv:2005.14381*, 2020.
- Cussens, J. Bayesian network learning with cutting planes. In *Conference on Uncertainty in Artificial Intelligence*, pp. 153–160, 2011.
- Cussens, J., Haws, D., and Studený, M. Polyhedral aspects of score equivalence in Bayesian network structure learning. *Mathematical Programming*, pp. 1–40, 2016.
- Drton, M., Eichler, M., and Richardson, T. Computing maximum likelihood estimates in recursive linear models with correlated errors. *Journal of Machine Learning Research*, 10(10):2329–2348, 2009.
- Evans, R. J. and Richardson, T. Markovian acyclic directed mixed graphs for discrete data. *The Annals of Statistics*, pp. 1452–1482, 2014.
- Gao, T. and Wei, D. Parallel bayesian network structure learning. In *International Conference on Machine Learning*, pp. 1671–1680, 2018.
- Grötschel, M., Jünger, M., and Reinelt, G. On the acyclic subgraph polytope. *Mathematical Programming*, 33(1): 28–42, 1985.
- Gámez, J., Mateo, J., and Puerta, J. Learning Bayesian networks by hill climbing: efficient methods based on progressive restriction of the neighborhood. *Data Mining and Knowledge Discovery*, 22(1-2):106–148, 2011.
- Heckerman, D., Geiger, D., and Chickering, D. M. Learning Bayesian networks: The combination of knowledge and statistical data. *Machine learning*, 20(3):197–243, 1995.
- Jaakkola, T., Sontag, D., Globerson, A., and Meila, M. Learning bayesian network structure using lp relaxations. In *International Conference on Artificial Intelligence and Statistics*, pp. 358–365, 2010.
- Karger, D. R. Global min-cuts in rnc, and other ramifications of a simple min-cut algorithm. In *Annual ACM-SIAM Symposium on Discrete Algorithms*, volume 93, pp. 21–30, 1993.
- Manzour, H., Küçükyavuz, S., Wu, H.-H., and Shojaie, A. Integer programming for learning directed acyclic graphs from continuous data. *INFORMS Journal on Optimization*, pp. ijoo–2019, 2020.
- Nowzohour, C., Maathuis, M. H., Evans, R. J., and Bühlmann, P. Distributional equivalence and structure learning for bow-free acyclic path diagrams. *Electronic Journal of Statistics*, 11(2):5342–5374, 2017.
- Park, Y. W. and Klabjan, D. Bayesian network learning via topological order. *The Journal of Machine Learning Research*, 18(1):3451–3482, 2017.
- Pearl, J. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2000.
- Ramsey, J., Zhang, J., and Spirtes, P. L. Adjacency-faithfulness and conservative causal inference. *arXiv preprint arXiv:1206.6843*, 2012.
- Richardson, T. A factorization criterion for acyclic directed mixed graphs. In *Conference on Uncertainty in Artificial Intelligence*, pp. 462–470, 2009.
- Richardson, T. and Spirtes, P. Ancestral graph markov models. *The Annals of Statistics*, 30(4):962–1030, 2002.
- Schwarz, G. Estimating the dimension of a model. *The annals of statistics*, 6(2):461–464, 1978.
- Silander, T. and Myllymaki, P. A simple approach for finding the globally optimal Bayesian network structure. In *Conference on Uncertainty in Artificial Intelligence*, pp. 445–452, 2006.
- Silander, T., Roos, T., Kontkanen, P., and Myllymäki, P. Factorized normalized maximum likelihood criterion for learning bayesian network structures. In *Proceedings of the 4th European Workshop on Probabilistic Graphical Models*, pp. 257–264, 2008.
- Singh, A. P. and Moore, A. W. Finding optimal Bayesian networks by dynamic programming. Technical report, Carnegie Mellon University, 2005.
- Spirtes, P. Introduction to causal inference. *Journal of Machine Learning Research*, 11(5), 2010.

- Spirtes, P., Glymour, C. N., Scheines, R., and Heckerman, D. *Causation, prediction, and search*. MIT press, 2000.
- Tillman, R. and Spirtes, P. Learning equivalence classes of acyclic models with latent and selection variables from multiple datasets with overlapping variables. In *International Conference on Artificial Intelligence and Statistics*, pp. 3–15, 2011.
- Triantafillou, S. and Tsamardinos, I. Score-based vs constraint-based causal learning in the presence of confounders. In *UAI Workshop on Causation: Foundation to Application*, pp. 59–67, 2016.
- Tsamardinos, I., Brown, L., and Aliferis, C. The max-min hill-climbing Bayesian network structure learning algorithm. *Machine Learning*, 65(1):31–78, 2006a.
- Tsamardinos, I., Brown, L. E., and Aliferis, C. F. The max-min hill-climbing bayesian network structure learning algorithm. *Machine learning*, 65(1):31–78, 2006b.
- Tsirlis, K., Lagani, V., Triantafillou, S., and Tsamardinos, I. On scoring maximal ancestral graphs with the max-min hill climbing algorithm. *International Journal of Approximate Reasoning*, 102:74–85, 2018.
- Yuan, C. and Malone, B. Learning optimal bayesian networks: A shortest path perspective. *Journal of Artificial Intelligence Research*, 48:23–65, 2013.
- Zhang, J. On the completeness of orientation rules for causal discovery in the presence of latent confounders and selection bias. *Artificial Intelligence*, 172(16-17): 1873–1896, 2008.