

---

# Selective Greedy Equivalence Search: Finding Optimal Bayesian Networks Using a Polynomial Number of Score Evaluations

---

David Maxwell Chickering  
Microsoft Research  
Redmond, WA 98052  
dmax@microsoft.com

Christopher Meek  
Microsoft Research  
Redmond, WA 98052  
meek@microsoft.com

## Abstract

We introduce Selective Greedy Equivalence Search (SGES), a restricted version of Greedy Equivalence Search (GES). SGES retains the asymptotic correctness of GES but, unlike GES, has polynomial performance guarantees. In particular, we show that when data are sampled independently from a distribution that is perfect with respect to a DAG  $\mathcal{G}$  defined over the observable variables then, in the limit of large data, SGES will identify the equivalence class of  $\mathcal{G}$  after a number of score evaluations that is (1) polynomial in the number of nodes and (2) exponential in various complexity measures including maximum-number-of-parents, maximum-clique-size, and a new measure called *v-width* that is at least as small as—and potentially much smaller than—the other two. More generally, we show that for any hereditary and equivalence-invariant property  $\Pi$  known to hold in  $\mathcal{G}$ , we retain the large-sample optimality guarantees of GES even if we ignore any GES deletion operator during the backward phase that results in a state for which  $\Pi$  does not hold in the common-descendants subgraph.

## 1 INTRODUCTION

Greedy Equivalence Search (GES) is a score-based search algorithm that searches over equivalence classes of Bayesian-network structures. The algorithm is appealing because (1) for finite data, it explicitly (and greedily) tries to maximize the score of interest, and (2) as the data grows large, it is guaranteed—under suitable distributional assumptions—to return the generative structure. Although empirical results show that the algorithm is efficient in real-world domains, the number of search states that GES needs to evaluate in the worst case can be exponential in the number of domain variables.

In this paper, we show that if we assume the generative distribution is perfect with respect to some DAG  $\mathcal{G}$  defined over the observable variables, and if  $\mathcal{G}$  is known to be constrained by one of various graph-theoretic measures of complexity, then we can disregard all but a polynomial number of the backward search operators considered by GES while retaining the large-sample guarantees of the algorithm; we call this new variant of GES *selective greedy equivalence search* or *SGES*. Our complexity results are a consequence of a new understanding of the backward phase of GES, in which edges (either directed or undirected) are greedily deleted from the current state until a local minimum is reached. We show that for any *hereditary* and *equivalence-invariant* property known to hold in generative model  $\mathcal{G}$ , we can remove from consideration any edge-deletion operator between  $X$  and  $Y$  for which the property does not hold in the resulting induced subgraph over  $X$ ,  $Y$ , and their common descendants. As an example, if we know that each node has at most  $k$  parents, we can remove from consideration any deletion operator that results in a common child with more than  $k$  parents.

By casting limited *v-width* and two other complexity measures—maximum-clique size and maximum-parent-set size—as graph properties, we show how to enumerate directly over a polynomial number of edge-deletion operators at each step, and we show that we need only a polynomial number of calls to the scoring function to complete the algorithm.

The main contributions of this paper are theoretical. Our definition of the new SGES algorithm deliberately leaves unspecified the details of how to implement its forward phase; we prove our results for SGES given *any* implementation of this phase that completes with a polynomial number of calls to the scoring function. A naive implementation is to immediately return a complete (i.e., no independence) graph using *no* calls to the scoring function, but this choice is unlikely to be reasonable in practice, particularly in discrete domains where the sample complexity of this initial model will likely be a problem. Whereas we believe it an important direction, our paper does not explore practical al-

ternatives for the forward phase that have polynomial-time guarantees.

Our paper is organized as follows. In Section 2, we describe related work. In Section 3, we provide notation and background material. In Section 4, we present our new SGES algorithm, we show that it is optimal in the large-sample limit, and we provide complexity bounds when given an equivalence-invariant and hereditary property that holds on the generative structure. In Section 5, we present a simple synthetic experiment that demonstrates the value of restricting the backward operators in SGES. We conclude with a discussion of our results in Section 6.

## 2 RELATED WORK

It is useful to distinguish between approaches to learning the structure of graphical models as *constraint based*, *score based* or *hybrid*. Constraint-based approaches typically use (conditional) independence tests to eliminate potential models, whereas score-based approaches typically use a penalized likelihood or a marginal likelihood to evaluate alternative model structures; hybrid methods combine these two approaches. Because score-based approaches are driven by a global likelihood, they are less susceptible than constraint-based approaches to incorrect categorical decisions about independences.

There are polynomial-time algorithms for learning the best model in which each node has at most one parent. In particular, the Chow-Liu algorithm (Chow and Liu, 1968) used with any equivalence-invariant score will identify the highest-scoring tree-like model in polynomial time; for scores that are not equivalence invariant, we can use the polynomial-time maximum-branching algorithm of Edmonds (1967) instead. Gaspers et al. (2012) show how to learn *k-branchings* in polynomial time; these models are polytrees that differ from a branching by a constant  $k$  number of edge deletions.

Without additional assumptions, most results for learning non-tree-like models are negative. Meek (2001) shows that finding the maximum-likelihood path is NP-hard, despite this being a special case of a tree-like model. Dasgupta (1999) shows that finding the maximum-likelihood polytree (a graph in which each pair of nodes is connected by at most one path) is NP-hard, even with bounded in-degree for every node. For general directed acyclic graphs, Chickering (1996) shows that finding the highest marginal-likelihood structure is NP-hard, even when each node has at most two parents. Chickering et al. (2004) extend this same result to the large-sample case.

Researchers often assume that the “generative” distribution of the training data is *perfect* with respect to some model class in order to reduce the complexity of learning algorithms. Geiger et al. (1990) provide a polynomial-time

constraint-based algorithm for recovering a polytree under the assumption that the generative distribution is perfect with respect to a polytree; an analogous score-based result follows from this paper. The constraint-based PC algorithm of Sprites et al. (1993) can identify the equivalence class of Bayesian networks in polynomial time if the generative structure is a DAG model over the observable variables in which each node has a bounded degree; this paper provides a similar result for a score-based algorithm. Kalish and Buhlmann (2007) show that for Gaussian distributions, the PC algorithm can identify the right structure even when the number of nodes in the domain is larger than the sample size. Chickering (2002) uses the same DAG-perfectness-over-observables assumption to show that the greedy GES algorithm is optimal in the large-sample limit, although the branching factor of GES is worst-case exponential; the main result of this paper shows how to limit this branching factor without losing the large-sample guarantee. Chickering and Meek (2002) show that GES identifies a “minimal” model in the large-sample limit under a less restrictive set of assumptions.

Hybrid methods for learning DAG models use a constraint-based algorithm to prune out a large portion of the search space, and then use a score-based algorithm to select among the remaining (Friedman et al., 1999; Tsamardinos et al., 2006). Ordyniak and Szeider (2013) give positive complexity results for the case when the remaining DAGs are characterized by a structure with constant treewidth.

Many researchers have turned to exhaustive enumeration to identify the highest-scoring model (Gillispie and Perlman, 2001; Koivisto and Sood 2004; Silander and Myllymäki, 2006; Kojima et al, 2010). There are many complexity results for other model classes. Karger and Srebro (2001) show that finding the optimal Markov network is NP-complete for treewidth  $> 1$ . Narasimhan and Bilmes (2004) and Shahaf, Chechetka and Guestrin (2009) show how to learn approximate limited-treewidth models in polynomial time. Abeel, Koller and Ng (2005) show how to learn factor graphs in polynomial time.

## 3 NOTATION AND BACKGROUND

We denote a variable by an upper case letter (e.g.,  $A$ ) and a state or value of that variable by the same letter in lower case (e.g.,  $a$ ). We denote a set of variables by a bold-face capitalized letter or letters (e.g.,  $\mathbf{X}$ ). We use a corresponding bold-face lower-case letter or letters (e.g.,  $\mathbf{x}$ ) to denote an assignment of state or value to each variable in a given set. We use calligraphic letters (e.g.,  $\mathcal{G}, \mathcal{E}$ ) to denote statistical models and graphs.

A *Bayesian-network model* for a set of variables  $\mathbf{U}$  is a pair  $(\mathcal{G}, \theta)$ .  $\mathcal{G} = (\mathbf{V}, \mathbf{E})$  is a directed acyclic graph—or *DAG* for short—consisting of nodes in one-to-one correspondence with the variables and directed edges that connect

those nodes.  $\theta$  is a set of parameter values that specify all of the conditional probability distributions. The Bayesian network represents a joint distribution over  $\mathbf{U}$  that factors according to the structure  $\mathcal{G}$ .

The structure  $\mathcal{G}$  of a Bayesian-network model represents the independence constraints that must hold in the distribution. The set of all independence constraints implied by the structure  $\mathcal{G}$  can be characterized by the *Markov conditions*, which are the constraints that each variable is independent of its non-descendants given its parents. All other independence constraints follow from properties of independence. A distribution defined over the variables from  $\mathcal{G}$  is *perfect with respect to  $\mathcal{G}$*  if the set of independences in the distribution is equal to the set of independences implied by the structure  $\mathcal{G}$ .

Two DAGs  $\mathcal{G}$  and  $\mathcal{G}'$  are *equivalent*<sup>1</sup>—denoted  $\mathcal{G} \approx \mathcal{G}'$ —if the independence constraints in the two DAGs are identical. Because equivalence is reflexive, symmetric, and transitive, the relation defines a set of equivalence classes over network structures. We use  $[\mathcal{G}]_{\approx}$  to denote the equivalence class of DAGs to which  $\mathcal{G}$  belongs.

An equivalence class of DAGs  $\mathcal{F}$  is an *independence map (IMAP)* of another equivalence class of DAGs  $\mathcal{E}$  if all independence constraints implied by  $\mathcal{F}$  are also implied by  $\mathcal{E}$ . For two DAGs  $\mathcal{G}$  and  $\mathcal{H}$ , we use  $\mathcal{G} \leq \mathcal{H}$  to denote that  $[\mathcal{H}]_{\approx}$  is an IMAP of  $[\mathcal{G}]_{\approx}$ ; we use  $\mathcal{G} < \mathcal{H}$  when  $\mathcal{G} \leq \mathcal{H}$  and  $[\mathcal{H}]_{\approx} \neq [\mathcal{G}]_{\approx}$ .

Verma and Pearl (1991) show that two DAGs are equivalent if and only if they have the same *skeleton* (i.e., the graph resulting from ignoring the directionality of the edges) and the same *v-structures* (i.e., pairs of edges  $X \rightarrow Y$  and  $Y \leftarrow Z$  where  $X$  and  $Z$  are not adjacent). As a result, we can use a *partially directed acyclic graph*—or *PDAG* for short—to represent an equivalence class of DAGs: for a PDAG  $\mathcal{P}$ , the equivalence class of DAGs is the set that has the same skeleton and the same v-structures as  $\mathcal{P}$ <sup>2</sup>.

We extend our notation for DAG equivalence and the DAG IMAP relation to include the more general PDAG structure. In particular, for a PDAG  $\mathcal{P}$ , we use  $[\mathcal{P}]_{\approx}$  to denote the corresponding equivalence class of DAGs. For any pair of PDAGs  $\mathcal{P}$  and  $\mathcal{Q}$ —where one or both may be a DAG—we use  $\mathcal{P} \approx \mathcal{Q}$  to denote  $[\mathcal{Q}]_{\approx} = [\mathcal{P}]_{\approx}$  and we use  $\mathcal{P} \leq \mathcal{Q}$  to denote  $[\mathcal{Q}]_{\approx}$  is an IMAP of  $[\mathcal{P}]_{\approx}$ . To avoid confusion, for the remainder of the paper we will reserve the symbols  $\mathcal{G}$  and  $\mathcal{H}$  for DAGs.

For any PDAG  $\mathcal{P}$  and subset of nodes  $\mathbf{V}$ , we use  $\mathcal{P}[\mathbf{V}]$  to denote the subgraph of  $\mathcal{P}$  induced by  $\mathbf{V}$ ; that is,  $\mathcal{P}[\mathbf{V}]$  has

<sup>1</sup>We make the standard conditional-distribution assumptions of multinomials for discrete variables and Gaussians for continuous variables so that if two DAGs have the same independence constraints, then they can also model the same set of distributions.

<sup>2</sup>The definitions for the skeleton and set of v-structures for a PDAG are the obvious extensions to these definitions for DAGs.

as nodes the set  $\mathbf{V}$  and has as edges all those from  $\mathcal{P}$  that connect nodes in  $\mathbf{V}$ . We use  $\text{NA}_{X,Y}$  to denote, within a PDAG, the set of nodes that are *neighbors* of  $X$  (i.e., connected with an undirected edge) and also adjacent to  $Y$  (i.e., without regard to whether the connecting edge is directed or undirected).

An edge in  $\mathcal{G}$  is *compelled* if it exists in every DAG that is equivalent to  $\mathcal{G}$ . If an edge in  $\mathcal{G}$  is not compelled, we say that it is *reversible*. A *completed PDAG (CPDAG)*  $\mathcal{C}$  is a PDAG with two additional properties: (1) for every directed edge in  $\mathcal{C}$ , the corresponding edge in  $\mathcal{G}$  is compelled and (2) for every undirected edge in  $\mathcal{C}$  the corresponding edge in  $\mathcal{G}$  is reversible. Unlike non-completed PDAGs, the CPDAG representation of an equivalence class is unique. We use  $\text{Pa}_Y^{\mathcal{P}}$  to denote the *parents* of node  $Y$  in  $\mathcal{P}$ . An edge  $X \rightarrow Y$  is *covered* in a DAG if  $X$  and  $Y$  have the same parents, with the exception that  $X$  is not a parent of itself.

### 3.1 Greedy Equivalence Search

---

Algorithm *GES*( $\mathbf{D}$ )

---

**Input** : Data  $\mathbf{D}$

**Output**: CPDAG  $\mathcal{C}$

$\mathcal{C} \leftarrow \text{FES}(\mathbf{D})$

$\mathcal{C} \leftarrow \text{BES}(\mathbf{D}, \mathcal{C})$

**return**  $\mathcal{C}$

---

Figure 1: Pseudo-code for the GES algorithm.

The GES algorithm, shown in Figure 1, performs a two-phase greedy search through the space of DAG equivalence classes. GES represents each search state with a CPDAG, and performs transformation operators to this representation to traverse between states. Each operator corresponds to a DAG edge modification, and is scored using a DAG scoring function that we assume has three properties. First, we assume the scoring function is *score equivalent*, which means that it assigns the same score to equivalent DAGs. Second, we assume the scoring function is *locally consistent*, which means that, given enough data, (1) if the current state *is not* an IMAP of  $\mathcal{G}$ , the score prefers edge additions that remove incorrect independences, and (2) if the current state *is* an IMAP of  $\mathcal{G}$ , the score prefers edge deletions that remove incorrect dependences. Finally, we assume the scoring function is *decomposable*, which means we can express it as:

$$\text{Score}(\mathcal{G}, \mathbf{D}) = \sum_{i=1}^n \text{Score}(X_i, \text{Pa}_i^{\mathcal{G}}) \quad (1)$$

Note that the data  $\mathbf{D}$  is implicit in the right-hand side Equation 1. Most scores in the literature have these properties.

For the remainder of this paper, we assume they hold for our scoring function.

All of the CPDAG operators from GES are scored using differences in the DAG scoring function, and in the limit of large data, these scores are positive precisely for those operators that remove incorrect independences and incorrect dependences.

The first phase of the GES—called *forward equivalence search* or *FES*—starts with an empty (i.e., no-edge) CPDAG and greedily applies *GES insert* operators until no operator has a positive score; these operators correspond precisely to the union of all single-edge additions to all DAG members of the current (equivalence-class) state. After FES reaches a local maximum, GES switches to the second phase—called *backward equivalence search* or *BES*—and greedily applies *GES delete* operators until no operator has a positive score; these operators correspond precisely to the union of all single-edge deletions from all DAG members of the current state.

**Theorem 1. (Chickering, 2002)** *Let  $\mathcal{C}$  be the CPDAG that results from applying the GES algorithm to  $m$  records sampled from a distribution that is perfect with respect to DAG  $\mathcal{G}$ . Then in the limit of large  $m$ ,  $\mathcal{C} \approx \mathcal{G}$ .*

The role of FES in the large-sample limit is only to identify a state  $\mathcal{C}$  for which  $\mathcal{G} \leq \mathcal{C}$ ; Theorem 1 holds for GES under any implementation of FES that results in an IMAP of  $\mathcal{G}$ . The implementation details can be important in practice because what constitutes a “large” amount of data depends on the number of parameters in the model. In theory, however, we could simply replace FES with a (constant-time) algorithm that sets  $\mathcal{C}$  to be the no-independence equivalence class.

The focus of our analysis in the next section is on a modified version of BES, and the details of the delete operator used in this phase are important. In Figure 2, we show the preconditions, scoring function, and transformation algorithm for a delete operator. We note that we do not need to make any CPDAG transformations when *scoring* the operators; it is only once we have identified the highest-scoring (non-negative) delete that we need to make the transformation shown in the figure. After applying the edge modifications described in the **foreach** loop, the resulting PDAG  $\mathcal{P}$  is not necessarily completed and hence we may have to convert  $\mathcal{P}$  into the corresponding CPDAG representation. As shown by Chickering (2002), this conversion can be accomplished easily by using the structure of  $\mathcal{P}$  to extract a DAG that we then convert into a CPDAG by undirecting all reversible edges. The complexity of this procedure for a  $\mathcal{P}$  with  $n$  nodes and  $e$  edges is  $O(n \cdot e)$ , and requires no calls to the scoring function.

---

**Operator:**  $Delete(X, Y, \mathbf{H})$  applied to  $\mathcal{C}$

---

• **Preconditions**

$X$  and  $Y$  are adjacent  
 $\mathbf{H} \subseteq \mathbf{NA}_{Y,X}$   
 $\overline{\mathbf{H}} = \mathbf{NA}_{Y,X} \setminus \mathbf{H}$  is a clique

• **Scoring**

$Score(Y, \{\mathbf{Pa}_Y^{\mathcal{C}} \cup \overline{\mathbf{H}}\} \setminus X) - Score(Y, X \cup \mathbf{Pa}_Y^{\mathcal{C}} \cup \overline{\mathbf{H}})$

• **Transformation**

Remove edge between  $X$  and  $Y$   
**foreach**  $H \in \mathbf{H}$  **do**  
    Replace  $Y - H$  with  $Y \rightarrow H$   
    **if**  $X - H$  **then** Replace with  $X \rightarrow H$ ;  
**end**  
Convert to CPDAG

---

Figure 2: Preconditions, scoring, and transformation algorithm for a delete operator applied to a CPDAG.

## 4 SELECTIVE GREEDY EQUIVALENCE SEARCH

In this section, we define a variant of the GES algorithm called *selective GES*—or *SGES* for short—that uses a subset of the BES operators. The subset is chosen based on a given property  $\Pi$  that is known to hold for the generative structure  $\mathcal{G}$ . Just like GES, SGES—shown in Figure 3—has a forward phase and a backward phase.

For the forward phase of SGES, it suffices for our theoretical analysis that we use a method that returns an IMAP of  $\mathcal{G}$  (in the large-sample limit) using only a polynomial number of insert-operator score calls. For this reason, we call this phase *poly-FES*. A simple implementation of poly-FES is to return the no-independence CPDAG (with no score calls), but other implementations are likely more useful in practice.

The backward phase of SGES—which we call *selective backward equivalence search* (*SBES*)—uses only a subset of the BES delete operators. This subset must necessarily include all  $\Pi$ -consistent delete operators—defined below—in order to maintain the large-sample consistency of GES, but the subset can (and will) include additional operators for the sake of efficient enumeration.

The DAG properties used by SGES must be *equivalence invariant*, meaning that for any pair of equivalent DAGs, either the property holds for both of them or it holds for neither of them. Thus, for any equivalence-invariant DAG property  $\Pi$ , it makes sense to say that  $\Pi$  either holds or does not hold for a PDAG. As shown by Chickering (1995), a DAG property is equivalence invariant if and only if it is invariant to covered-edge reversals; it follows that the prop-

erty that each node has at most  $k$  parents is equivalence invariant, whereas the property that the length of the longest directed path is at least  $k$  is not. Furthermore, the properties for SGES must also be *hereditary*, which means that if  $\Pi$  holds for a PDAG  $\mathcal{P}$  it must also hold for all induced subgraphs of  $\mathcal{P}$ . For example, the max-parent property is hereditary, whereas the property that each node has *at least*  $k$  parents is not. We use *EIH property* to refer to a property that is equivalence invariant and hereditary.

**Definition 1.  $\Pi$ -Consistent GES Delete**

A GES delete operator  $Delete(X, Y, \mathbf{H})$  is  $\Pi$  consistent for CPDAG  $\mathcal{C}$  if, for the set of common descendants  $\mathbf{W}$  of  $X$  and  $Y$  in the resulting CPDAG  $\mathcal{C}'$ , the property holds for the induced subgraph  $\mathcal{C}'[X \cup Y \cup \mathbf{W}]$ .

In other words, after the delete, the property holds for the subgraph defined by  $X$ ,  $Y$ , and their common descendants.

---

Algorithm  $SGES(\mathbf{D}, \Pi)$

---

**Input** : Data  $\mathbf{D}$ , Property  $\Pi$

**Output**: CPDAG  $\mathcal{C}$

$\mathcal{C} \leftarrow$  poly-FES  
 $\mathcal{C} \leftarrow$  SBES( $\mathbf{D}, \mathcal{C}, \Pi$ )  
**return**  $\mathcal{C}$

---

Figure 3: Pseudo-code for the SGES algorithm.

---

Algorithm  $SBES(\mathbf{D}, \mathcal{C}, \Pi)$

---

**Input** : Data  $\mathbf{D}$ , CPDAG  $\mathcal{C}$ , Property  $\Pi$

**Output**: CPDAG

**Repeat**

**Ops**  $\leftarrow$  Generate  $\Pi$ -consistent delete operators for  $\mathcal{C}$   
 $Op \leftarrow$  highest-scoring operator in **Ops**  
**if** score of  $Op$  is negative **then return**  $\mathcal{C}$   
 $\mathcal{C} \leftarrow$  Apply  $Op$  to  $\mathcal{C}$

---

Figure 4: Pseudo-code for the SBES algorithm.

**4.1 LARGE-SAMPLE CORRECTNESS**

The following theorem establishes a graph-theoretic justification for considering only the  $\Pi$ -consistent deletions at each step of SBES.

**Theorem 2.** *If  $\mathcal{G} < \mathcal{C}$  for CPDAG  $\mathcal{C}$  and DAG  $\mathcal{G}$ , then for any EIH property  $\Pi$  that holds on  $\mathcal{G}$ , there exists a  $\Pi$ -consistent  $Delete(X, Y, \mathbf{H})$  that when applied to  $\mathcal{C}$  results in the CPDAG  $\mathcal{C}'$  for which  $\mathcal{G} \leq \mathcal{C}'$ .*

The proof of Theorem 2 can be found in Chickering and Meek (2015), an expanded version of this paper. The result is a consequence of an explicit characterization of, for a

given pair of DAGs  $\mathcal{G}$  and  $\mathcal{H}$  such that  $\mathcal{G} < \mathcal{H}$ , an edge in  $\mathcal{H}$  that we can either reverse or delete in  $\mathcal{H}$  such that for the resulting DAG  $\mathcal{H}'$ , we have  $\mathcal{G} \leq \mathcal{H}'^3$ .

**Theorem 3.** *Let  $\mathcal{C}$  be the CPDAG that results from applying the SGES algorithm to (1)  $m$  records sampled from a distribution that is perfect with respect to DAG  $\mathcal{G}$  and (2) EIH property  $\Pi$  that holds on  $\mathcal{G}$ . Then in the limit of large  $m$ ,  $\mathcal{C} \approx \mathcal{G}$ .*

**Proof:** Because the scoring function is locally consistent, we know poly-FES must return an IMAP of  $\mathcal{G}$ . Because SBES includes all the  $\Pi$ -consistent delete operators, Theorem 2 guarantees that, unless  $\mathcal{C} \approx \mathcal{G}$ , there will be a positive-scoring operator.  $\square$

**4.2 COMPLEXITY MEASURES**

In this section, we discuss a number of distributional assumptions that we can use with Theorem 3 to limit the number of operators that SGES needs to score. As discussed in Section 2, when we assume the generative distribution is perfect with respect to a DAG  $\mathcal{G}$ , then graph-theoretic assumptions about  $\mathcal{G}$  can lead to more efficient training algorithms. Common assumptions used include (1) a maximum parent-set size for any node, (2) a maximum-clique<sup>4</sup> size among any nodes and (3) a maximum treewidth. Treewidth is important because the complexity of exact inference is exponential in this measure.

We can associate a property with each of these assumptions that holds precisely when the DAG  $\mathcal{G}$  satisfies that assumption. Consider the constraint that the maximum number of parents for any node in  $\mathcal{G}$  is some constant  $k$ . Then, using “PS” to denote parent size, we can define the property  $\Pi_{PS}^k$  to be true precisely for those DAGs in which each node has at most  $k$  parents. Similarly we can define  $\Pi_{CL}^k$  and  $\Pi_{TW}^k$  to correspond to maximum-clique size and maximum treewidth, respectively.

For two properties  $\Pi$  and  $\Pi'$ , we write  $\Pi \subseteq \Pi'$  if for every DAG  $\mathcal{G}$  for which  $\Pi$  holds,  $\Pi'$  also holds. In other words,  $\Pi$  is a *more constraining* property than is  $\Pi'$ . Because the lowest node in any clique has all other nodes in the clique as parents, it is easy to see that  $\Pi_{PS}^k \subseteq \Pi_{CL}^{k-1}$ . Because the treewidth for DAG  $\mathcal{G}$  is defined to be the size of the largest clique minus one in a graph whose cliques are at least as large as those in  $\mathcal{G}$ , we also have  $\Pi_{TW}^k \subseteq \Pi_{CL}^{k-1}$ . Which property to use will typically be a trade-off between how reasonable the assumption is (i.e, less constraining properties are more reasonable) and the efficiency of the resulting algorithm (i.e., more constraining properties lead to faster algorithms).

<sup>3</sup>Chickering (2002) characterizes the *reverse* transformation of reversals/additions in  $\mathcal{G}$ , which provides an *implicit* characterization of reversals/deletions in  $\mathcal{H}$ .

<sup>4</sup>We use *clique* in a DAG to mean a set of nodes in which all pairs are adjacent.

We now consider a new complexity measure called  $v$ -width, whose corresponding property is less constraining than the previous three. For a DAG  $\mathcal{G}$ , the  $v$ -width is defined to be the maximum of, over all pairs of non-adjacent nodes  $X$  and  $Y$ , the size of the largest clique among common children of  $X$  and  $Y$ . In other words,  $v$ -width is similar to the maximum-clique-size bound, except that the bound only applies to cliques of nodes that are shared children of some pair of non-adjacent nodes. With this understanding it is easy to see that, for the property  $\Pi_{VW}^k$  corresponding to a bound on the  $v$ -width, we have  $\Pi_{CL}^k \subseteq \Pi_{VW}^k$ .

To illustrate the difference between  $v$ -width and the other complexity measures, consider the two DAGs in Figure 5. The DAG in Figure 5(a) has a clique of size  $K$ , and consequently a maximum-clique size of  $K$  and a maximum parent-set size of  $K - 1$ . Thus, if  $K$  is  $O(n)$  for a large graph of  $n$  nodes, any algorithm that is exponential in these measures will not be efficient. The  $v$ -width, however, is zero for this DAG. The DAG in Figure 5(b), on the other hand, has a  $v$ -width of  $K$ .




Figure 5: Two DAGs (a) and (b) having identical maximum clique sizes, similar maximum number of parents, and divergent  $v$ -widths.

In order to use a property with SGES, we need to establish that it is EIH. For  $\Pi_{PS}^k$ ,  $\Pi_{CL}^k$  and  $\Pi_{VW}^k$ , equivalence-invariance follows from the fact that all three properties are covered-edge invariant, and hereditary follows because the corresponding measures cannot increase when we remove nodes and edges from a DAG. Although we can establish EIH for the treewidth property  $\Pi_{TW}^k$  with more work, we omit further consideration of treewidth for the sake of space.

### 4.3 GENERATING DELETIONS

In this section, we show how to generate a set of deletion operators for SBES such that all  $\Pi$ -consistent deletion operators are included, for any  $\Pi \in \{\Pi_{PS}^k, \Pi_{CL}^k, \Pi_{VW}^k\}$ . Furthermore, the total number of deletion operators we generate is polynomial in the number of nodes in the domain and exponential in  $k$ .

Our approach is to restrict the  $Delete(X, Y, \mathbf{H})$  operators based on the  $\mathbf{H}$  sets and the resulting CPDAG  $\mathcal{C}'$ . In par-

ticular, we *rule out* candidate  $\mathbf{H}$  sets for which  $\Pi$  does not hold on the induced subgraph  $\mathcal{C}'[\mathbf{H} \cup X \cup Y]$ ; because all nodes in  $\mathbf{H}$  will be common children of  $X$  and  $Y$  in  $\mathcal{C}'$ —and thus a subset of the common descendants of  $X$  and  $Y$ —we know from Definition 1 (and the fact that  $\Pi$  is hereditary) that none of the dropped operators can be  $\Pi$ -consistent.

Before presenting our restricted-enumeration algorithm, we now discuss how to enumerate delete operators without restrictions. As shown by Andersson et al. (1997), a CPDAG is a chain graph whose undirected components are chordal. This means that the induced sub-graph defined over  $\mathbf{NA}_{Y,X}$ —which is a subset of the neighbors of  $Y$ —is an undirected chordal graph. A useful property of chordal graphs is that we can identify, in polynomial time, a set of *maximal cliques* over these nodes<sup>5</sup>; let  $\mathbf{C}_1, \dots, \mathbf{C}_m$  denote the nodes contained within these  $m$  maximal cliques, and let  $\bar{\mathbf{H}} = \mathbf{NA}_{Y,X} \setminus \mathbf{H}$  be the complement of the shared neighbors with respect to the candidate  $\mathbf{H}$ . Recall from Figure 2 that the preconditions for any  $Delete(X, Y, \mathbf{H})$  include the requirement that  $\bar{\mathbf{H}}$  is a clique. This means that for any valid  $\mathbf{H}$ , there must be some maximal clique  $\mathbf{C}_i$  that contains the entirety of  $\bar{\mathbf{H}}$ ; thus, we can generate all operators (without regard to any property) by stepping through each maximal clique  $\mathbf{C}_i$  in turn, initializing  $\mathbf{H}$  to be all nodes *not* in  $\mathbf{C}_i$ , and then generating a new operator corresponding to expanding  $\mathbf{H}$  by all subsets of nodes in  $\mathbf{C}_i$ . Note that if  $\mathbf{NA}_{Y,X}$  is itself a clique, we are enumerating over all  $2^{|\mathbf{NA}_{Y,X}|}$  operators.

As we show below, all three of the properties of interest impose a bound on the maximum clique size among nodes in  $\bar{\mathbf{H}}$ . If we are given such a bound  $s$ , we know that any “expansion” subset for a clique that has size greater than  $s$  will result in an operator that is not valid. Thus, we can implement the above operator-enumeration approach more efficiently by only generating subsets within each clique that have size at most  $s$ . This allows us to process each clique  $\mathbf{C}_i$  with only  $O(|\mathbf{C}_i + 1|^s)$  calls to the scoring function. In addition, we need not enumerate over *any* of the subsets of  $\mathbf{C}_i$  if, after removing this clique from the graph, there remains a clique of size greater than  $s$ ; we define the function  $FilterCliques(\{\mathbf{C}_1, \dots, \mathbf{C}_m\}, s)$  to be the subset of cliques that remain after imposing this constraint. With this function, we can define SELECTIVE-GENERATE-OPS as shown in Figure 6 to leverage the max-clique-size constraint when generating operators; this algorithm will in turn be used to generate all of the CPDAG operators during SBES.

**Example:** In Figure 7, we show an example CPDAG  $\mathcal{C}$  for which to run  $SELECTIVE-GENERATE-OPS(\mathcal{C}, X, Y, s)$  for various values of  $s$ . In the example, there is a single clique

<sup>5</sup>Blair and Peyton (1993) provide a good survey on chordal graphs and detail how to identify the maximal cliques while running maximum-cardinality search.

---

Algorithm SELECTIVE-GENERATE-OPS( $\mathcal{C}, X, Y, s$ )

---

**Input** : CPDAG  $\mathcal{C}$  with adjacent  $X, Y$  and limit  $s$

**Output**:  $\text{Ops} = \{\mathbf{H}_1, \dots, \mathbf{H}_m\}$

```

Ops  $\leftarrow \emptyset$ 
Generate maximal cliques  $\mathbf{C}_1, \dots, \mathbf{C}_m$  from  $\text{NA}_{Y,X}$ 
 $\mathbf{S} \leftarrow \text{FilterCliques}(\{\mathbf{C}_1, \dots, \mathbf{C}_m\}, s)$ 
foreach  $\mathbf{C}_i \in \mathbf{S}$  do
   $\mathbf{H}_0 \leftarrow \text{NA}_{Y,X} \setminus \mathbf{C}_i$ 
  foreach  $\mathbf{C} \subseteq \mathbf{C}_i$  with  $|\mathbf{C}| \leq s$  do
    Add  $\mathbf{H}_0 \cup \mathbf{C}$  to  $\text{Ops}$ 
  end
end
return  $\text{Ops}$ 

```

---

Figure 6: Algorithm to generate clique-size limited delete operators.

$\mathbf{C} = \{A, B\}$  in the set  $\text{NA}_{Y,X}$ , and thus at the top of the outer **foreach** loop, the set  $\mathbf{H}_0$  is initialized to the empty set. If  $s = 0$ , the only subset of  $\mathbf{C}$  with size zero is the empty set, and so that is added to  $\text{Ops}$  and the algorithm returns. If  $s = 1$  we add, in addition to the empty set, all singleton subsets of  $\mathbf{C}$ . For  $s \geq 2$ , we add all subsets of  $\mathbf{C}$ .  $\square$

Now we discuss how each of the three properties impose a constraint  $s$  on the maximum clique among nodes in  $\mathbf{H}$ , and consequently the selective-generation algorithm in Figure 6 can be used with each one, given an appropriate bound  $s$ . For both  $\Pi_{VW}^k$  and  $\Pi_{CL}^k$ , the  $k$  given imposes an explicit bound on  $s$  (i.e.,  $s = k$  for both). Because any clique in  $\mathbf{H}$  of size  $r$  will result in a DAG member of the resulting equivalence class having a node in that clique with at least  $r + 1$  parents (i.e.,  $r - 1$  from the other nodes in the clique, plus both  $X$  and  $Y$ ), we have for  $\Pi_{PS}^k$ ,  $s = k - 1$ .

We summarize the discussion above in the following proposition.

**Proposition 1.** *Algorithm SELECTIVE-GENERATE-OPS applied to all edges using clique-size bound  $s$  generates all  $\Pi$ -consistent delete operators for  $\Pi \in \{\Pi_{PS}^{s+1}, \Pi_{CL}^s, \Pi_{VW}^s\}$ .*

We now argue that running SBES on a domain of  $n$  variables when using Algorithm SELECTIVE-GENERATE-OPS with a bound  $s$  requires only a polynomial number in  $n$  of calls to the scoring function. Each clique in the inner loop of the algorithm can contain at most  $n$  nodes, and therefore we generate and score at most  $(n + 1)^s$  operators, requiring at most  $2(n + 1)^s$  calls to the scoring function. Because the cliques are maximal, there can be at most  $n$  of them considered in the outer loop. Because there are never more than  $n^2$  edges in a CPDAG, and we will delete at most all of them, we conclude that even if we decided to rescore ev-

ery operator after every edge deletion, we will only make a polynomial number of calls to the scoring function.

From the above discussion and the fact that SBES completes using at most a polynomial number of calls to the scoring function, we get the following result for the full SGENS algorithm.

**Proposition 2.** *The SGENS algorithm, when run over a domain of  $n$  variables and given  $\Pi \in \{\Pi_{PS}^{s+1}, \Pi_{CL}^s, \Pi_{VW}^s\}$ , runs to completion using a number of calls to the DAG scoring function that is polynomial in  $n$  and exponential in  $s$ .*




Figure 7: An example CPDAG  $\mathcal{C}$  and the resulting operators generated by SELECTIVE-GENERATE-OPS( $\mathcal{C}, X, Y, s$ ) for various values of  $s$ .

## 5 EXPERIMENTS

In this section, we present a simple synthetic experiment comparing SBES and BES that demonstrates the value of pruning operators. In our experiment we used an *oracle* scoring function. In particular, given a generative model  $\mathcal{G}$ , our scoring function computes the minimum-description-length score assuming a data size of five billion records, but without actually sampling any data: instead, we use exact inference in  $\mathcal{G}$  (i.e., instead of counting from data) to compute the conditional probabilities needed to compute the expected log loss. This allows us to get near-asymptotic behavior without the need to sample data. To evaluate the cost of running each algorithm, we counted the number of times the scoring function was called on a unique node and parent-set combination; we cached these scores away so that if they were needed multiple times during a run of the algorithm, they were only computed (and counted) once.

In Figure 8, we show the average number of scoring-function calls required to complete BES and SBES when starting from a complete graph over a domain of  $n$  binary variables, for varying values of  $n$ . Each average is taken over ten trials, corresponding to ten random generative models. We generated the structure of each generative model as follows. First, we ordered all node pairs by randomly permuting the nodes and taking each node in turn with each of its predecessors in turn. For each node pair, we chose to attempt an edge insertion with probability one half. For each attempt, we added an edge if doing so (1)

did not create a cycle and (2) did not result in a node having more than two parents; if an edge could be added in either direction, we chose the direction at random. We sampled the conditional distributions for each node and each parent configuration from a uniform Dirichlet distribution with equivalent-sample size of one. We ran SBES with  $\Pi_{PS}^2$ .




Figure 8: Number of score evaluations needed to run BES and SBES, starting from the complete graph, for a range of domain sizes.

Our results show clearly the exponential dependence of BES on the number of nodes in the clique, and the increasing savings we get with SBES by leveraging the fact that  $\Pi_{PS}^2$  holds in the generative structure.

Note that to realize large savings in practice, when GES runs FES instead of starting from a dense graph, a (relatively sparse) generative distribution must lead FES to an equivalence class containing a (relatively dense) undirected clique that is subsequently “thinned” during BES. We can synthesize challenging grid distributions to force FES into such states, but it is not clear how realistic such distributions are in practice. When we re-run the clique experiment above, but where we instead start both BES and SBES from the model that results from running FES (i.e., with no polynomial-time guarantee), the savings from SBES are small due to the fact that the subsequent equivalence classes do not contain large cliques.

## 6 CONCLUSION

We introduced a restricted version of the GES algorithm, SGES, that leverages graph-theoretic complexity properties to prune the backward-phase operators. We showed that for a particular class of properties—which includes maximum-clique size, maximum number of parents, and  $v$ -width—we can guarantee that the number of score evaluations is polynomial in the number of nodes in the domain.

The fact that we can use our approach to selectively choose operators for any hereditary and equivalence invariant graph-theoretic property provides the opportunity to explore alternative complexity measures. Another candidate complexity measure is the maximum number of  $v$ -structures. Although the corresponding property does not limit the maximum size of a *clique* in  $\mathbf{H}$ , it limits directly the size  $|\mathbf{H}|$  for every operator and thus it is easy to enumerate these operators efficiently. Another complexity measure of interest is *treewidth*, due to the fact that exact inference in a Bayesian-network model is takes time exponential in this measure.

The results we have presented are for the general Bayesian-network learning problem. It is interesting to consider the implications of our results for the problem of learning particular subsets of Bayesian networks. One natural class that we discussed in Section 2 is that of polytrees. If we assume that the generative distribution is perfect with respect to a polytree then we know the  $v$ -width of the generative graph is one. This implies, in the limit of large data, that we can recover the structure of the generative graph with a polynomial number of score evaluations. This provides a score-based recovery algorithm analogous to the constraint-based approach of Geiger et al. (1990).

We presented a simple complexity analysis for the purpose of demonstrating that SGES uses a only polynomial number of calls to the scoring function. We leave as future work a more careful analysis that establishes useful constants in this polynomial. In particular, we can derive tighter bounds on the total number of node-and-parent-configurations that are needed to score all the operators for each CPDAG, and by caching these configuration scores we can further take advantage of the fact that most operators remain valid (i.e., the preconditions still hold) and have the same score after each transformation.

Finally, we plan to investigate practical implementations of poly-FES that have the polynomial-time guarantees needed for SGES.



## References

- [1] Pieter Abbeel, Daphne Koller, and Andrew Y. Ng. Learning factor graphs in polynomial time and sample complexity. *JMLR*, 7:1743–1788, 2006.
- [2] Steen A. Andersson, David Madigan, and Michael D. Perlman. A characterization of Markov equivalence classes for acyclic digraphs. *Annals of Statistics*, 25:505–541, 1997.
- [3] Jean R. S. Blair and Barry W. Peyton. An introduction to chordal graphs and clique trees. In *Graph Theory and Sparse Matrix Computations*, pages 1–29, 1993.
- [4] David Maxwell Chickering. A transformational characterization of Bayesian network structures. In *UAI’95*, pages 87–98. 1995.
- [5] David Maxwell Chickering. Learning Bayesian networks is NP-complete. In *Learning from Data: Artificial Intelligence and Statistics V*, pages 121–130. Springer-Verlag, 1996.
- [6] David Maxwell Chickering. Optimal structure identification with greedy search. *JMLR*, 3:507–554, November 2002.
- [7] David Maxwell Chickering and Christopher Meek. Finding optimal Bayesian networks. In *UAI’02*, pages 94–102. 2002.
- [8] David Maxwell Chickering and Christopher Meek. Selective greedy equivalence search: Finding optimal Bayesian networks using a polynomial number of score evaluations. MSR-TR-2015-45, 2015.
- [9] David Maxwell Chickering, Christopher Meek, and David Heckerman. Large-sample learning of Bayesian networks is NP-hard. *JMLR*, 5:1287–1330, October 2004.
- [10] C. Chow and C. Liu. Approximating discrete probability distributions with dependence trees. *IEEE Transactions on Information Theory*, 14:462–467, 1968.
- [11] Sanjoy Dasgupta. Learning polytrees. In *UAI’99*, pages 131–141. 1999.
- [12] Jack Edmonds. Optimum branching. *J. Res. NBS*, 71B:233–240, 1967.
- [13] Nir Friedman, Iftach Nachman, and Dana Peer. Learning Bayesian network structure from massive datasets: The “sparse candidate” algorithm. In *UAI’99*. 1999.
- [14] Serge Gaspers, Mikko Koivisto, Mathieu Liedloff, Sebastian Ordyniak, and Stefan Szeider. On finding optimal polytrees. In *AAAI’12*. 2012.
- [15] Dan Geiger, Azaria Paz, and Judea Pearl. Learning causal trees from dependence information. In *AAAI’90*, pages 770–776. AAAI Press, 1990.
- [16] Steven B. Gillispie and Michael D. Perlman. Enumerating Markov equivalence classes of acyclic digraph models. In *UAI’01*, pages 171–177. 2001.
- [17] Markus Kalisch and Peter Buhlmann. Estimating high-dimensional directed acyclic graphs with the PC algorithm. *JMLR*, 8:613–636, 2007.
- [18] David Karger and Nathan Srebro. Learning Markov networks: Maximum bounded tree-width graphs. In *SODA’01*, pages 391–401, January 2001.
- [19] Mikko Koivisto and Kismat Sood. Exact Bayesian structure discovery in Bayesian networks. *JMLR*, 5:549–573, December 2004.
- [20] Kaname Kojima, Eric Perrier, Seiya Imoto, and Satoru Miyano. Optimal search on clustered structural constraint for learning Bayesian network structure. *JMLR*, 11:285–310, 2010.
- [21] Christopher Meek. Finding a path is harder than finding a tree. *JAIR*, 15:383–389, 2001.
- [22] Mukund Narasimhan and Jeff Bilmes. PAC-learning bounded tree-width graphical models. In *UAI’04*, *UAI’04*, pages 410–417, 2004.
- [23] Sebastian Ordyniak and Stefan Szeider. Parameterized complexity results for exact Bayesian network structure learning. *JAIR*, 46:263–302, 2013.
- [24] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 1988.
- [25] Dafna Shahaf, Anton Chechetka, and Carlos Guestrin. Learning thin junction trees via graph cuts. In *AISTATS’09*, 2009.
- [26] Tomi Silander and Petri Myllymäki. A simple approach for finding the globally optimal Bayesian network structure. In *UAI’06*, pages 445–452, 2006.
- [27] Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, Prediction, and Search*. Springer-Verlag, New York, 1993.
- [28] Ioannis Tsamardinos, Laura E. Brown, and Constantin F. Aliferis. The max-min hill-climbing Bayesian network structure learning algorithm. *Machine Learning*, 2006.
- [29] Thomas Verma and Judea Pearl. Equivalence and synthesis of causal models. In *UAI’91*, pages 220–227, 1991.