
Learning Equivalence Classes of Bayesian Network Structures

David Maxwell Chickering*
Computer Science Department
University of California at Los Angeles
dmax@cs.ucla.edu

Abstract

Approaches to learning Bayesian networks from data typically combine a scoring metric with a heuristic search procedure. Given a Bayesian network structure, many of the scoring metrics derived in the literature return a score for the entire equivalence class to which the structure belongs. When using such a metric, it is appropriate for the heuristic search algorithm to search over equivalence classes of Bayesian networks as opposed to individual structures. We present the general formulation of a search space for which the states of the search correspond to equivalence classes of structures. Using this space, any one of a number of heuristic search algorithms can easily be applied. We compare greedy search performance in the proposed search space to greedy search performance in a search space for which the states correspond to individual Bayesian network structures.

1 INTRODUCTION

Recently, many researchers have developed methods for learning Bayesian networks from data. The available techniques include Bayesian methods [Cooper and Herskovits, 1991, Buntine, 1991, Spiegelhalter et al., 1993, Heckerman et al., 1995], quasi-Bayesian methods [Lam and Bacchus, 1993, Bouckaert, 1993], and non-Bayesian methods [Pearl and Verma, 1991, Spirtes et al., 1993]. Much of the work in learning Bayesian networks has been devoted to the derivation of a *scoring metric*. Given a candidate Bayesian network structure, the scoring metric evaluates how well the structure “fits” the observed data and prior

knowledge. Once the scoring metric has been defined, learning Bayesian networks reduces to a search for one or more structures that have a high metric score. Chickering (1995a) shows that this search problem is NP-hard when the Bayesian scoring metric derived by Heckerman et al. (1995) is used. Consequently, it is appropriate to apply heuristic search algorithms in this domain.

Before any search algorithm can be applied to a problem, we must define the three components of a *search space*. First, we need to identify the states of the search, or equivalently, the set of all potential solutions to the search problem. Second, we need a representation for the states of the search. Third, we need a set of operators that transform the representation of one state to another so that the algorithm can traverse the space in a systematic way. Once the search space has been defined, any one of a number of well-known search algorithms can easily be applied to that space.

In perhaps the simplest formulation of a search space for learning Bayesian networks, the states of the search are defined to be individual Bayesian networks, the representation of a state is simply an acyclic directed graph, and the operators are defined to be local changes to those graphs. For example, Chickering et al. (1995) compare various search procedures in the search space of Bayesian networks, using the following operators: for any pair of nodes x and y , if x and y are adjacent, the edge connecting them can be either deleted or reversed. If x and y are not adjacent, an edge can be added in either direction. All operators are subject to the constraint that a cycle can not be formed. We shall use *Bayesian-network space*, or *B-space* for short, to denote the search space defined in this way. Figure 1 shows an example of each operator in B-space.

When two Bayesian networks assert the same set of independence assumptions among the variables in the domain, we say that those networks are *equivalent*. The relation of network equivalence imposes a set of

*Supported by NSF Grant No. IRI-9119825, and a grant from Rockwell International.

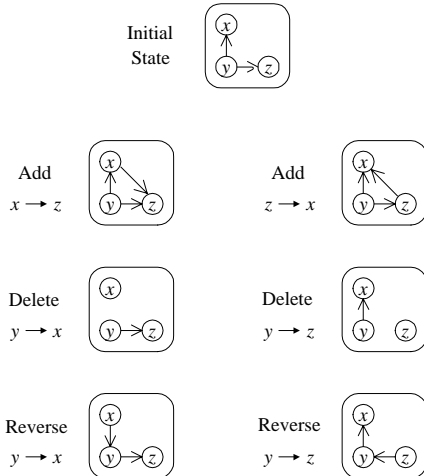


Figure 1: States resulting from the application of a single operator in B-space

equivalence classes over Bayesian network structures. Chickering (1995b) shows that many of the scoring metrics derived in the literature actually return the measure of fit for the entire equivalence class to which the candidate structure belongs. We call any such metric *score equivalent*.

If B-space is used by a search algorithm in conjunction with a score-equivalent metric, any particular equivalence class may be represented by a huge number of search states. Furthermore, many of the operators defined for this space move between states corresponding to the same equivalence class. Because the objective of the search algorithm is to identify the state with highest metric score, it seems natural that the states of the search space should correspond to equivalence classes of Bayesian networks whenever a score-equivalent metric is used.

In this paper we provide the specification for a search space in which the states are defined to be equivalence classes of Bayesian networks. In Section 2 we introduce our notation and describe previous relevant work. In Section 3, we define the search space. In particular, we describe an efficient representation for equivalence classes, and introduce a set of simple operators that can be applied to this representation. Finally, in Section 4 we compare greedy search performance in B-space to greedy search performance in the search space that we define in Section 3. We should emphasize that the search space we propose is suitable for any number of search algorithms, and that we have chosen to compare this space to B-space using greedy search because greedy search is very prevalent in the literature on learning Bayesian networks.

2 BACKGROUND

In this section we introduce our notation and describe previous relevant work.

A Bayesian network B for a set of variables $U = \{x_1, \dots, x_n\}$ is a pair $(\mathcal{G}, \theta_{\mathcal{G}})$ where $\mathcal{G} = (U, E_{\mathcal{G}})$ is a dag, and $\theta_{\mathcal{G}}$ is the set of conditional probability distributions that correspond to \mathcal{G} . We now present a formal definition of equivalence.

Definition *Two dags \mathcal{G} and \mathcal{G}' are equivalent if for every Bayesian network $B = (\mathcal{G}, \theta_{\mathcal{G}})$, there exists a Bayesian network $B' = (\mathcal{G}', \theta_{\mathcal{G}'})$ such that B and B' define the same probability distribution, and vice versa.*

We use $\mathcal{G} \approx \mathcal{G}'$ to denote that \mathcal{G} and \mathcal{G}' are equivalent. As was stated earlier, the relation \approx defines a set of equivalence classes over the network structures. A directed edge $x_i \rightarrow x_j \in E_{\mathcal{G}}$ is *compelled* in \mathcal{G} if for every dag $\mathcal{G}' \approx \mathcal{G}$, $x_i \rightarrow x_j \in E_{\mathcal{G}'}$. For any edge $e \in E_{\mathcal{G}}$, if e is not compelled in \mathcal{G} , then e is *reversible* in \mathcal{G} , that is, there exists some dag \mathcal{G}' equivalent to \mathcal{G} in which e has opposite orientation.

The *skeleton* of any dag is the undirected graph resulting from ignoring the directionality of every edge. A *v-structure* in dag \mathcal{G} is an ordered triple of nodes (x, y, z) such that (1) \mathcal{G} contains the arcs $x \rightarrow y$ and $z \rightarrow y$, and (2) x and z are not adjacent in \mathcal{G} . Verma and Pearl (1990) derive the following characterization of equivalent structures:

Theorem 1 [Verma and Pearl, 1990] *Two dags are equivalent if and only if they have the same skeletons and the same v-structures.*

Figure 2 shows the set of all 27 states of B-space defined over three-node Bayesian networks. All dags contained within the same dashed rectangle are equivalent — there are 11 equivalence classes for this space.

A consequence of Theorem 1 is that for any edge e participating in a v-structure in some dag \mathcal{G} , if that edge is reversed in some other dag \mathcal{G}' , then \mathcal{G} and \mathcal{G}' are not equivalent. Thus any edge participating in a v-structure is compelled. However, not every compelled edge necessarily participates in a v-structure — we leave as an exercise to the reader to prove that the edge from z to w is compelled in the dag shown in Figure 3.

Acyclic partially directed graphs, or *pdags* for short, are graphs that contain both directed and undirected edges, and are commonly used to represent equivalence classes of Bayesian networks¹. Let \mathcal{P} denote an

¹pdags are sometimes called *patterns* in the literature

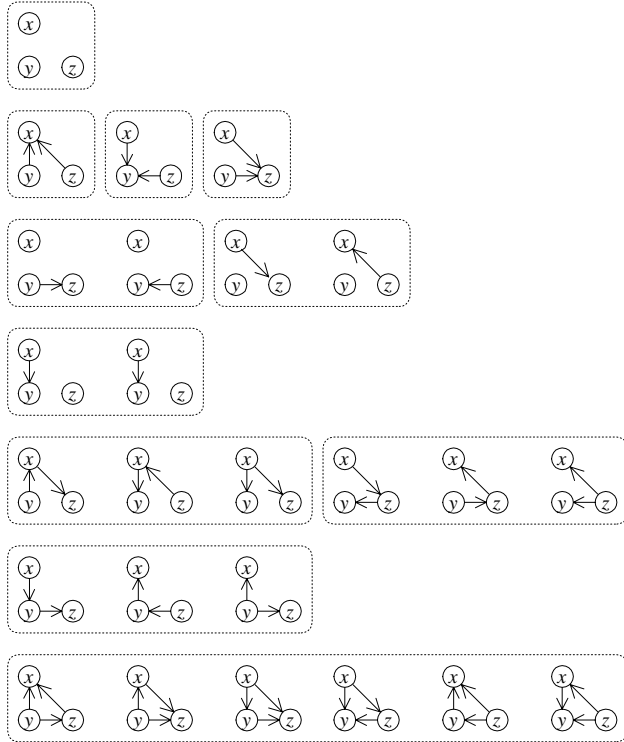


Figure 2: The 27 states of B-search for networks containing three nodes. The equivalent networks in this space are grouped together with a dashed rectangle.

arbitrary pdag. We define the equivalence class of dags corresponding to \mathcal{P} — denoted $Class(\mathcal{P})$ — as follows: $\mathcal{G} \in Class(\mathcal{P})$ if and only if \mathcal{G} and \mathcal{P} have the same skeleton and the same set of v-structures². From Theorem 1, it follows that a pdag containing a directed edge for every edge participating in a v-structure, and an undirected edge for all other edges, uniquely identifies an equivalence class of dags. However, there may be many other pdags that correspond to the same equivalence class. For example, any dag interpreted as a pdag can be used with our definition of $Class$ to represent its own equivalence class. It follows that for an equivalence class containing k dags, there are *at least* k different pdags that can represent that class. In the next section, we introduce a subclass of pdags that have a one to one correspondence with equivalence classes.

If a dag \mathcal{G} has the same skeleton and the same set of *directed edges* as a pdag \mathcal{P} , we say that \mathcal{G} is a *consistent extension* of \mathcal{P} . Note that any dag that is a consistent extension of \mathcal{P} must also be contained in $Class(\mathcal{P})$, but that not every dag in $Class(\mathcal{P})$ is a

²The definitions for the skeleton and set of v-structures for a pdag are the obvious extensions to these definitions for dags.




Figure 3: Example of a dag containing a compelled edge that does not participate in a v-structure.

consistent extension of \mathcal{P} . If there is at least one consistent extension of a pdag \mathcal{P} , we say that \mathcal{P} *admits a consistent extension*. Figure 4a shows a pdag that admits a consistent extension, and Figure 4b shows one such consistent extension. Figure 4c shows a pdag that does not admit a consistent extension.




Figure 4: (a) a pdag that admits a consistent extension, (b) a consistent extension of the pdag in (a) and (c) a pdag that does not admit a consistent extension.

3 DEFINING THE SEARCH SPACE

As was discussed in Section 1, a search space has three components: (1) a set of states, (2) a representation for the states, and (3) a set of operators. Given that a score-equivalent metric is being used, we have argued that the states of the search should be equivalence classes of Bayesian network structures. In this

section, we define an efficient representation for equivalence classes, and a simple set of operators that can be applied to this representation. We call the resulting search space *equivalence-class space*, or *E-space* for short.

Recall the definition of compelled and reversible edges given in Section 2. Every dag in a particular equivalence class has the same set of compelled and reversible edges, and consequently we can associate compelled and reversible edges with equivalence classes as well as with dags. We define the *completed pdag representation* of an equivalence class to be the pdag consisting of a directed edge for every compelled edge in the equivalence class, and an undirected edge for every reversible edge in the equivalence class. Figure 5a shows a dag \mathcal{G} , and Figure 5b shows the completed pdag representation for $Class(\mathcal{G})$.




Figure 5: (a) a dag \mathcal{G} and (b) the completed pdag representation for $Class(\mathcal{G})$

We use completed pdags to represent the states of the search in E-space. Using completed pdags instead of general pdags (or dags in the case of B-space) eliminates the problem of having multiple representations for the same equivalence class. We emphasize this result with a lemma:

Lemma 2 *Let \mathcal{P}^{c_1} and \mathcal{P}^{c_2} denote two completed pdags that both admit a consistent extension. Then $\mathcal{P}^{c_1} = \mathcal{P}^{c_2}$ if and only if $Class(\mathcal{P}^{c_1}) = Class(\mathcal{P}^{c_2})$.*

Proof: Follows immediately by Theorem 1 and by definition of compelled and reversible edges. \square

To complete the specification of E-space, we define a set of simple operators that can be applied to completed pdags. The operators make extensive use of two algorithms. The first algorithm, which we refer to as PDAG-TO-DAG, takes as input an arbitrary pdag and returns a consistent extension if one exists. PDAG-TO-DAG returns an error message if the given pdag does not admit a consistent extension. Dor and Tarsi (1992) present an efficient implementation of PDAG-TO-DAG that we used for our experimental results. The second algorithm, which we

refer to as DAG-TO-PDAG, takes as input an arbitrary dag \mathcal{G} and returns the completed pdag representation for $Class(\mathcal{G})$. Chickering (1995b) and Meek (1995) have both derived implementations of DAG-TO-PDAG. For our experimental results, we used the algorithm derived by Chickering (1995b), which is asymptotically optimal in the average case.

Given a completed pdag \mathcal{P}^c , we define the following four types of operators:

1. For any undirected edge $x-y$ in \mathcal{P}^c , we can delete the edge
2. For any directed edge $x \rightarrow y$ in \mathcal{P}^c , we can either delete the edge or reverse the edge
3. For any pair of nodes x and y that are not adjacent in \mathcal{P}^c , we can either insert an undirected edge between x and y , or insert a directed edge in either direction
4. For any triple of nodes x, y and z in \mathcal{P}^c , if x and z are not adjacent, and either x and y or z and y are adjacent, we can insert the v-structure (x, y, z)

All operators are subject to the constraint that the resulting pdag is acyclic and admits a consistent extension.

For a given completed pdag, let \mathcal{P} denote the pdag, not necessarily completed, that results after directly applying one of the operators. The completed pdag that results from the operator is obtained as follows. First, the algorithm PDAG-TO-DAG is called with input \mathcal{P} to extract a consistent extension \mathcal{G} . If \mathcal{P} does not admit a consistent extension, then the given operator is not legal. To complete the application, the algorithm DAG-TO-PDAG is called with input \mathcal{G} to build the resulting completed pdag representation. The process of applying an operator is depicted schematically in Figure 6, and the application of each operator type is illustrated in Figure 7.

Note that the consistent extension obtained during the application of an operator can be used to score the resulting state if the scoring metric takes a dag as input, which is typically the case. Alternatively, PDAG-TO-DAG can be called with the completed pdag representation of a state whenever that state needs to be scored.

The proposed operators are all simple and local changes to the edges in a pdag, but as we see in the example of Figure 6 an operator can have “cascading” effects on the pdag. Furthermore, a local change in a completed pdag may not correspond to a local change in the consistent extension used to score the equivalence class. This is unfortunate, because most scoring

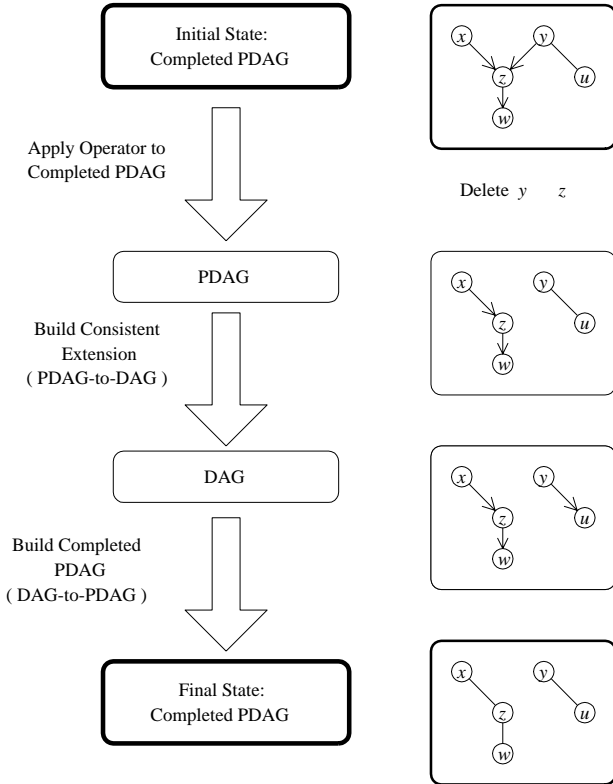


Figure 6: Diagram depicting how operators are applied. The pdags on the right give an example for every stage of the process.

metrics can exploit the locality of changes to dags to efficiently update the corresponding score. Fortunately, we have found that the cascading effects are not very common in practice. For most of the operator applications, the resulting state can be scored by applying a small number of local changes to a consistent extension of the current state.

4 EXPERIMENTAL RESULTS

In this section we compare the performance of greedy search when applied to E-space to the performance of greedy search when applied to B-space. Greedy search is a simple algorithm that, given the current state, always moves to the adjacent state that increases the metric score the most. If no adjacent state has a higher metric score, the algorithm terminates. In all of our experiments, we used greedy search to identify a single equivalence class with a high metric score.

Many researchers comparing search algorithms measure search performance in terms of how close the learned network (or equivalence class) is to the *gold standard network*, or the model from which the train-

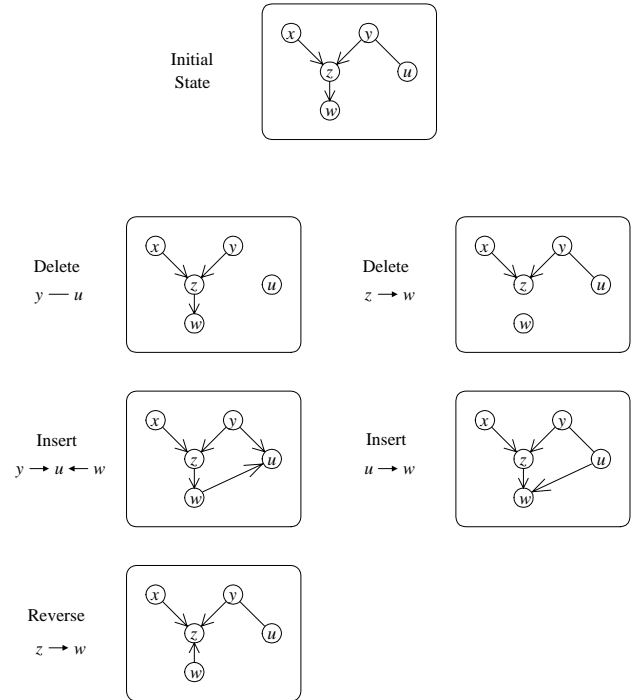


Figure 7: Example of each type of operator

ing data was generated. Typical choices for this comparison include structural difference and cross entropy. Although we report structural differences in our results, we emphasize that the goal of the search algorithm is to identify the structure with highest metric score, and consequently metric score is the best indicator of search performance in this domain.

In our experiments we compared the structure of the learned equivalence class to the structure of the equivalence class to which the gold standard belongs. To make this comparison, we first generated the completed pdag representation for both equivalence classes. Then, for each pair of nodes in the pdags, if the edge between those nodes were different in the two structures, we added one to the structural difference.

For all of our experiments, we used a score-equivalent variant of the *BDe* metric derived by Heckerman et al. (1995). Given data and prior knowledge, the BDe metric evaluates the relative posterior probability of the candidate structure using some reasonable assumptions. The version we used corresponds to an empty prior network and an equivalent sample size of either 8 (for the random graphs) or 16 (for the Alarm network). For a detailed description of the BDe metric we refer the reader to Heckerman et al. (1995). In all of the results reported, the BDe scores are expressed in log-space.

Table 1: Effect of gold-standard size on search performance. Each metric score and structural difference is the average over 9 total random databases consisting of 500 cases each.

Number of Nodes	E-space Score	B-space Score	Score Diff.	E-space Struct	B-space Struct	Struct Diff	E-space Time	B-space Time	Time Ratio
5	-1326.64	-1327.06	0.42	0.78	1.44	0.66	1	0	—
10	-2745.55	-2764.05	18.5	4.44	10.56	6.12	18.11	1.67	10.84431
15	-3665.29	-3677.17	11.88	17.67	21.89	4.22	70.44	6.22	11.32476
20	-5372.94	-5408.67	35.73	25.11	30.78	5.67	184.67	11.78	15.67657
25	-6786.83	-6860.24	73.41	32.67	47.11	14.44	487.33	22.56	21.60151

For each experiment we present (1) the difference in metric score, (2) the difference in structural difference from the gold standard, and (3) the ratio of learning times. The difference in metric score is the score in E-space minus the score in B-space, and hence any positive number indicates that search in E-space is outperforming search in B-space. The difference in structural difference, however, is the structural difference in B-space minus the structural difference in E-space. We chose to present the results this way so that a positive difference in structural score also denotes a win for E-space. The ratio of times denotes how many times longer the search in E-space took than the search in B-space. This ratio was always greater than one.

For many of the experiments we used randomly-generated gold standards. All random gold standards contained binary variables, and were generated as follows: for every pair of nodes in the graph, a directed edge was inserted with probability 0.3, subject to the constraint that no node could have more than 4 parents. Each conditional parameter set of the resulting network was drawn from a uniform Dirichlet distribution.

Our first experiment investigated how search performance in E-space compares to search performance in B-space as the size of the graphs increase. We used five sizes for the gold-standard in this experiment: 5, 10, 15, 20, and 25 nodes. For each graph size, we generated three random gold standards. For each gold standard, we generated three random databases, consisting of 500 cases each. We ran greedy search in both E-space and B-space, using each of the databases generated. The results are summarized in Table 1. For every gold-standard size, greedy search in E-space outperformed greedy search in B-space, both in terms of the average metric score and in terms of the average structural difference. Furthermore, as the complexity of the gold standard increased, the difference between the quality measures tended to increase as well.

Our next experiment compares the search spaces when the size of the database increases. We used random

gold standards for this experiment, where each gold standard contained 10 binary nodes. We used six database sizes: 500, 1000, 1500, 2000, 2500, and 3000 cases. For each database size, three gold standards were generated, and for each gold standard, three databases containing of the given number of cases were generated. For each database, greedy search was run in both E-space and B-space. The results are summarized in Table 2. Again, greedy search in E-space outperformed greedy search in B-space, in both quality measures, for every database size. The difference in the measures increased at first, but leveled off as the size of the databases increased. We conjecture that as the database size grows indefinitely, this difference will *decrease* and eventually greedy search in either space will obtain the same equivalence class. We base this conjecture on the assumption that as the data grows to infinity, the both search spaces will become smoother.

The final experiment was run using data generated from the Alarm network (Beinlich et al., 1989). The network, which contains 37 nodes and 46 edges, is an expert system for the problem of ICU ventilator management and has become a standard benchmark for learning algorithms. We generated 10 databases from the Alarm network, where each database contained 500 cases. For each database, greedy search was run in both E-space and B-space. The results are summarized in Table 3. In this table, we have omitted the search time because we were unable to finish the experiment using some obvious optimizations for E-space (which were implemented for all the other experiments). Without these optimizations, the search time in E-space was approximately 130 times longer than the search time in B-space. We are confident that this ratio will be significantly reduced.

5 DISCUSSION

Although our empirical results show that the greedy algorithm applied to E-space consistently outperforms the greedy algorithm applied to B-space, the time to complete the search in E-space was significantly longer.

Table 2: Effect of database size on search performance. Each metric score and structural difference entry is the average over 9 total databases generated from gold standards containing 10 nodes.

DB Size	E-space Score	B-space Score	Score Diff.	E-space Struct	B-space Struct	Struct Diff	E-space Time	B-space Time	Time Ratio
500	-2745.55	-2764.05	18.5	4.44	10.56	6.12	18.11	1.67	10.84431
1000	-5399.82	-5449.35	49.53	2.67	9.22	6.55	21.67	3.33	6.507508
1500	-8092.83	-8148.43	55.6	3.56	8.22	4.66	25.67	4.33	5.928406
2000	-10724.9	-10825.1	100.2	2	7	5	34.56	6.11	5.656301
2500	-13386.5	-13416.6	30.1	2.78	10.78	8	36.89	8.22	4.487835
3000	-16050.7	-16146	95.3	2.78	12.22	9.44	43	9.33	4.608789

Table 3: Greedy search performance for the Alarm metric. Metric scores and structural differences are averages over 10 databases, where each database contains 500 cases.

E-space Score	B-space Score	Score Diff.	E-space Struct	B-space Struct	Struct Diff
-5599.72	-5625.88	26.16	72.9	82	9.1

There are two reasons for this. Consider a search where all networks have n nodes. In B-space, there are $\frac{n(n-1)}{2}$ operators for greedy search to consider at each step. In E-space, however, there are $\frac{n(n-1)}{2} + 2e(n-2)$ operators at each step, where e denotes the number of edges in the current completed pdag (the extra operators correspond to v-structure insertions). In addition to the extra operators, there is also additional overhead to apply each operator. In particular, the algorithm PDAG-TO-DAG runs in time $O(n^2)$, and the algorithm DAG-TO-PDAG runs in time $O(e)$. Note that the extra overhead from these algorithms is a function of the network and does not depend on the size of the database. Consequently, as the data-sets grow large this overhead becomes less significant, as is demonstrated in Table 2.

In addition to the optimizations alluded to in the previous section, there are many tricks that can be applied to speed up search in E-space that we did not have time to complete. For the final version of this paper, we anticipate a significant reduction in running time for E-space.

As with most search problems, performance can be improved by customizing the algorithms for the particular domain. We believe that by restricting the number of v-structure insertions allowed to be applied to any given state, the time to complete greedy search in E-space can be reduced to roughly the same time as in B-space, without significant reduction in solution quality.

An interesting extension to the work presented here is

to combine E-space and B-space. One approach would be to run the greedy search in B-space until a local maximum is reached. Next, generate the completed pdag representation and see if the score can be increased in E-space. If it can, make one step in E-space and then switch back to B-space and go to the next local maximum. This approach will be fast because each local maximum is reached in B-space. Furthermore, by using E-space to get out of these local maxima, we hope that the resulting search performance will be good.

References

- [Beinlich et al., 1989] Beinlich, I., Suermondt, H., Chavez, R., and Cooper, G. (1989). The ALARM monitoring system: A case study with two probabilistic inference techniques for belief networks. In *Proceedings of the Second European Conference on Artificial Intelligence in Medicine*, London, pages 247–256. Springer Verlag, Berlin.
- [Bouckaert, 1993] Bouckaert, R. (1993). Probabilistic network construction using the minimum description length principle. In *ECSQARU*, pages 41–48.
- [Buntine, 1991] Buntine, W. (1991). Theory refinement on Bayesian networks. In *Proceedings of Seventh Conference on Uncertainty in Artificial Intelligence*, Los Angeles, CA, pages 52–60. Morgan Kaufmann.
- [Chickering et al., 1995] Chickering, D., Geiger, D., and Heckerman, D. (1995). Learning Bayesian net-

- works: Search methods and experimental results. In *Proceedings of the Fifth International Workshop on Artificial Intelligence and Statistics*.
- [Chickering, 1995a] Chickering, D. M. (1995a). Learning Bayesian networks is NP-Complete. *Submitted to: Lecture Notes in Statistics*.
- [Chickering, 1995b] Chickering, D. M. (1995b). A transformational characterization of Bayesian network structures. In *Proceedings of Eleventh Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann. to appear.
- [Cooper and Herskovits, 1991] Cooper, G. and Herskovits, E. (1991). A Bayesian method for constructing Bayesian belief networks from databases. In *Proceedings of Seventh Conference on Uncertainty in Artificial Intelligence*, Los Angeles, CA, pages 86–94. Morgan Kaufmann.
- [Dor and Tarsi, 1992] Dor, D. and Tarsi, M. (1992). A simple algorithm to construct a consistent extension of a partially oriented graph. Technical Report R-185, Cognitive Systems Laboratory, UCLA Computer Science Department.
- [Heckerman et al., 1995] Heckerman, D., Geiger, D., and Chickering, D. (1995). Learning discrete Bayesian networks. *Machine Learning*. to appear.
- [Lam and Bacchus, 1993] Lam, W. and Bacchus, F. (1993). Using causal information and local measures to learn Bayesian networks. In *Proceedings of Ninth Conference on Uncertainty in Artificial Intelligence*, Washington, DC, pages 243–250. Morgan Kaufmann.
- [Meek, 1995] Meek, C. (1995). Causal inference and causal explanation with background knowledge. In *Proceedings of Eleventh Conference on Uncertainty in Artificial Intelligence*, To Appear. Morgan Kaufman.
- [Pearl and Verma, 1991] Pearl, J. and Verma, T. (1991). A theory of inferred causation. In Allen, J., Fikes, R., and Sandewall, E., editors, *Knowledge Representation and Reasoning: Proceedings of the Second International Conference*, pages 441–452. Morgan Kaufmann, New York.
- [Spiegelhalter et al., 1993] Spiegelhalter, D., Dawid, A., Lauritzen, S., and Cowell, R. (1993). Bayesian analysis in expert systems. *Statistical Science*, 8:219–282.
- [Spirtes et al., 1993] Spirtes, P., Glymour, C., and Scheines, R. (1993). *Causation, Prediction, and Search*. Springer-Verlag, New York.
- [Verma and Pearl, 1990] Verma, T. and Pearl, J. (1990). Equivalence and synthesis of causal models. In *Proceedings of Sixth Conference on Uncertainty in Artificial Intelligence*, pages 220–227.