

Learning the DAG of Bayesian Belief Networks by Asking (Conditional) (In-)Dependence Questions

Claus Möbus

Learning and Cognitive Systems
 Department of Computing Science
 University of Oldenburg, D-26111 Germany
<http://www.lks.uni-oldenburg.de/>
claus.moebus@uni-oldenburg.de

Hilke Garbe

Learning and Cognitive Systems
 Department of Computing Science
 University of Oldenburg, D-26111 Germany
<http://www.lks.uni-oldenburg.de/>
hilke.garbe@uni-oldenburg.de

ABSTRACT

Bayesian belief networks (BBNs) have become the de facto standard for the representation of uncertain knowledge. They consist of a qualitative and of a quantitative part describing the (in-)dependencies between the variables of interest as a directed acyclic graph (DAG) and the decomposition of the joint probability distribution (JPD) as a product of conditional probability distributions constrained by the structure of the DAG. In this paper we present a new constraint-based query procedure: *Query-an-Oracle* (QAO). We assume that an *oracle* – preferable a human domain expert – is at hand which is competent and willing to answer questions generated by QAO concerning the directed (causal) dependence and (conditional) independence of the relevant random variables in the domain. Compared to other structure learning methods (e.g. the PC-Algorithm of Peter Spirtes and Clark Glymour and the IC-Algorithm of Pearl) QAO has a number of advantages. It derives the DAG of the BBN with less computational complexity, with no redundant questions, and is able to exploit *directed* dependence information without urging oracles to differentiate between *direct* and *indirect* influence.

Categories and Subject Descriptors

I.2.4 Knowledge Representation Formalisms and Methods

General Terms

Algorithms, Measurement, Design, Experimentation, Human Factors

LEARNING THE STRUCTURE OF BBNs

We were looking for an alternative method which is formal, simple to use for domain experts, not restrictive in its assumptions, and not suffering from the same or similar drawbacks as the Bounded Strata Method [1] or the PC-/IC-Algorithm [2, 3].

The Query-an-Oracle (QAO) Algorithm

This led to the development of a *greedy* knowledge acquisition

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright is held by the author/owner(s).
 K-CAP'09, September 1–4, 2009, Redondo Beach, California, USA.
 ACM 978-1-60558-658-8/09/09.

method for the construction of the transitive closure of the *precedes/causes*(X,Y)-relation[4,5].




Figure 1. Model Space of Structure Learning Algorithms PCA, ICA, and QAO

First step of QAO

In the *first* step of QAO this algorithm controls the non-redundant pair comparison of variables, and generates the Hasse and the transitive closure diagram of the partial-order relation *precedes/causes*(X, Y). The *greedy* behaviour is controlled by 13 production or propagation rules [4, 5]. When a pair (i, j) of variables is presented the oracle has to select a rating from a set of alternatives { *i causes/precedes j*, *i follows j*, *i neither causes/precedes nor follows j* } internally abbreviated as {+(i, j), -(i, j), 0(i, j)}. The QAO selects a special order of pair comparisons along the main diagonal of the adjacency matrix. When the variables are ancestral ordered this minimizes the number of pair comparisons queried from the oracle and maximizes the number of transitive inferences that are automatically generated by the algorithm. Because each transitive edge has to be confirmed by a (directional) conditional independence test in the *second* step of QAO the sequence of the query process is unimportant when we take the query complexity of *both* steps into account. Taking only the +(i,j) markings from the transitive closure we can reconstruct the Hasse diagram.

Second step of QAO

After termination of step 1 the semantics of the “+” edges has become unambiguous. The propagation process in step 1 denotes “+”-edges as *not transitive*. The semantics of the “+”-edge between X and Y is that X is a *direct influencer* of Y. In contrast to that the semantics of the “++”-edge is still ambiguous. The oracle has to decide whether “++”-edges denote a *direct influencer* or an *indirect* one. The *true DAG* (only known to nature) lies somewhere between the Hasse model and the transitive closure model (Fig. 1). In [4] we recommended a Markov blanket test for *every* variable for the decision in the second step. The improved solutions we propose here are much simpler. There are two alternative solutions for the *second* step.

- (1) The *first* is a Markov blanket test only for “++”-edges.
- (2) The *other* is a new kind of query to the oracle: Instead of the conventional *non directional* conditional independence $I(X, Y | Z)$ QAO asks for each transitive “++(X,Y)”-edge in the transitive closure model a directional conditional independence $I_{dir}(Y, X | Z)$: *Does X causally influence Y, when Z is known? Yes or No?*

Complexity of QAO

A rough calculation of QAO’s complexity for the *first* step gives following results. In the best case the true graph describes a total order. When we assume that the variables are ancestral ordered. QAO acquires the Hasse and the transitive closure diagrams by $(n-1)$ questions only. Due to transitivity $(n-2)(n-1)/2$ instances of a precedes/causes relations are inferred by QAO and need not be queried from the oracle. Thus, the *first* step of QAO has a best case complexity of $O(n)$ for the query process and a computational complexity of $O(n^4)$ because the computation of transitive edges requires at most $[(n-1)(n-2)/2]^2$ edge tests. In the worst case QAO needs $n(n-1)/2$ ratings. Then the DAG has no transitivity edges and the *second* step is not required. In this case the complexity of the query process is $O(n^2)$ and the computational complexity for the transitivity calculation is $O(c)$.

A Comparison of PCA and QAO

It is interesting to benchmark the improved QAO asking the directional conditional independence questions in the second step with the conventional constraint-based structure learning methods PCA [2] and ICA [3].

Example 7.1.3 from Jensen and Nielsen

The first example for a comparison is from [6, p.235f] (Figure 2). We wrote a computer program according to the Jensen and Nielsen’s pseudo code. In the *first* step PCA generated 40 independence queries (10 independence queries $I(X,Y)$ of order zero, 21 queries $I(X,Y|{Z})$ of order one and 9 queries $I(X,Y|{Z1,Z2})$ of order two). The independence queries $I(A, B)$, $I(B, C)$, $I(C, D | \{A\})$, $I(A, E | \{C, D\})$, and $I(E, B | \{C, D\})$ were answered with “Yes”. The result is the undirected skeleton graph (Figure 2). Then in the *second* step three production rules are applied to the skeleton. They mainly introduce directed v-structures (here two). Then in the *third* step the oracle is queried again to

direct the remaining undirected edges. The total number of queries is 41.

What is the behaviour of QAO? In the *first* step QAO asks 8 pair comparison queries. In the *second* step 2 directional conditional independence questions are generated $I_{dir}(E, A | \{C, D\})$ and $I_{dir}(E, B | \{C, D\})$ (Figure 2) and answered with “Yes”. The total number of queries is 10. Compared to PCA, this is a reduction of almost 75%.

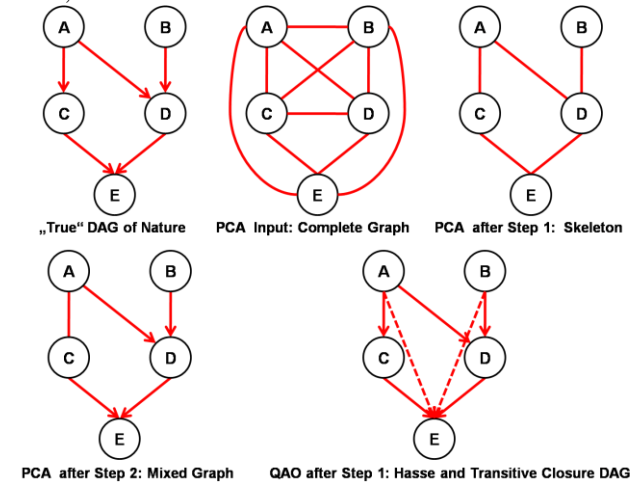


Figure 2. Steps 1 and 2 of PCA and Step 1 of QAO

SUMMARY

We presented a new structure learning algorithm Query-an-Oracle (QAO). This is in the *second* step an improved version of [4]. Compared to the well-known PC-algorithm of Spirtes et al. [2] and to the IC-algorithm of Pearl [3] it has a considerably smaller query complexity which is exactly $O(n^2)$. QAO is the first structure learning algorithm which exploits *directional dependence* ratings without urging oracles to give ratings of *direct* influence or *direct* control.

REFERENCES

- [1] Pearl, J., 1998, Probabilistic Reasoning in Intelligent Systems, San Mateo, CA., ISBN 0-934613-73-7
- [2] Spirtes, P., Glymour, C., and R. Scheines, Causation, Prediction, and Search, Camb, Mass.: MIT Press, 2000
- [3] Pearl, J., 2001, Causality: Models, Reasoning, and Inference, Cambridge University Press
- [4] Möbus, C., and Seebold, H., A Greedy Knowledge Acquisition Method for the Rapid Prototyping of Bayesian Belief Networks, in: Ch-K. Looi, G. McCalla, B. Bredeweg & J. Breuker (eds), Artificial Intelligence in Education, Amsterdam: IOS Press, 2005, p. 875-877, ISBN 1-58603-530-4
- [5] Möbus, C. & Seebold, H., and Garbe, H., A Greedy Knowledge Acquisition Method for the Rapid Prototyping of Knowledge Structures, p.211–212, in: Clark, P.; Schreiber, G.; Proc. 3rd Int. Conf. on Knowledge Capture, 2005, NY: ACM Press, ISBN 1-59593-163-5;
- [6] Jensen, F.V. & Nielsen, Th. D, 2007, Bayesian Networks and Decision Graphs, Berlin: Springer