

Learning Bayesian Networks With Hidden Variables for User Modeling

Frank Wittig*

Department of Computer Science, University of Saarbrücken, Germany

Abstract. The goal of the research summarized here is to develop methods for learning Bayesian networks on the basis of empirical data, focusing on issues that are especially important in the context of user modeling. These issues include the treatment of theoretically interpretable hidden variables, ways of learning partial networks and combining them into one single compound network, and ways of taking into account the special properties of datasets acquired through psychological experiments.

1 Goals

Although Bayesian networks have frequently been employed for user modeling tasks, techniques for learning the networks from data have so far rarely been applied in this context. Yet these learning techniques offer a promising way of deriving satisfactory answers to the question “Where do the numbers come from?”. Possible sources of data include (a) records of naturally occurring user behavior and (b) the raw data of psychological experiments. The latter type of data is especially interesting for user modeling research in that it raises some issues that are more or less specific to situations in which a human agent is being modeled.

I aim to achieve the following specific goals:

1. Identify, develop, and extend existing methods for gathering empirical data and for learning appropriate (perhaps very complex) Bayesian networks for user modeling. These networks should be not only technically correct but also interpretable.
2. Offer explicit guidance to future user modelers who want to apply such learning methods.
3. In the specific example domain that I am working in, strengthen the empirical basis of the Bayesian networks for the READY system (see, e.g., Berthold and Jameson, 1999).

2 Issues and Current Status

First, appropriate algorithms for the learning task have to be identified. For user modeling, an especially important case of learning Bayesian networks is the case where the structure—which may include hidden variables such as the user’s working memory load—is known, whereas the entries in the conditional probability tables (CPTs) are unknown. There exist several methods for handling this case, including Gibbs sampling (see Heckerman, 1995) and the APN approach

* This research is being supported by the German Science Foundation (DFG) in its Collaborative Research Center on Resource-Adaptive Cognitive Processes, SFB 378, Project B2, READY.

(Russell et al., 1995). Currently, I'm learning rather simple networks on the basis of data from a psychological experiment that was recently conducted in the READY project, in which a user's execution of different types of instructions with or without a concurrent secondary task was studied. So far, I have received the best results with an implementation of the APN algorithm with directions chosen by the Polak-Ribière method as described by Russell et al. (1995).

A second issue concerns the problem of learning probabilities in such a way that the learned networks will be interpretable. Normally, learning algorithms do not take into account the meaning of the variables in question. They yield a network that performs more or less well when applied to particular cases. But the CPTs associated with the nodes—especially with those for the hidden variables—are sometimes inconsistent with the intended theoretical interpretation of the variables in question. I am looking for ways of influencing the learning process so that the theoretical meaning of the key variables is maintained. One possible solution to this problem is to introduce constraints on the CPT entries before learning takes place. (The use of such constraints may also greatly increase the efficiency of the learning.)

The known learning algorithms do not consider the special situation of learning from a dataset that contains data from a limited number of persons. The typically large individual differences between experimental subjects have to be taken into account when the results of the learning are generalized to other persons. A first approach that I have tried is to introduce variables (e.g., "average speed of execution") that capture important individual differences. Since the values of these variables can be computed straightforwardly on the basis of the raw data, these variables can serve as observable variables for the purpose of learning, although they would not be observable in an application situation.

In most realistic scenarios, it will be impossible to gather data for entire networks at once, because the complete networks are in general too complex. This is an especially important problem in the present context, because psychological experiments are typically designed to investigate only a few variables at a time. Therefore, an interesting problem is that of learning partial networks and putting these parts together into one single compound Bayesian network. During the learning phase it should be possible to fix some CPTs in advance and to restrict the learning process to the other CPTs.

The evaluation of the learned Bayesian networks can be accomplished in at least two ways: In a traditional cross-validation, the dataset is partitioned into a learning set and a test set and the learned Bayesian networks are evaluated in terms of their accuracy on the test cases. A complementary type of evaluation checks whether a learned network can actually make useful inferences about a user on the basis of the limited input data that is typically available (see Berthold and Jameson, 1999).

References

- Berthold, A., and Jameson, A. (1999). Interpreting symptoms of cognitive load in speech input. In Kay, J., ed., *User Modeling: Proceedings of the Seventh International Conference, UM99*. Vienna, New York: Springer Wien New York.
- Heckerman, D. (1995). A tutorial on learning with Bayesian networks. Technical Report MSR-TR-95-06, Microsoft Research. Revised November 1996.
- Russell, S., Binder, J., Koller, D., and Kanazawa, K. (1995). Local learning in probabilistic networks with hidden variables. In Mellish, C. S., ed., *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*. San Mateo, CA: Morgan Kaufmann. 1146–1152.