International Journal of Approximate Reasoning 54 (2013) 1168–1181

Contents lists available at SciVerse ScienceDirect

International Journal of Approximate Reasoning

journal homepage: www.elsevier.com/locate/ijar

# An interactive approach for Bayesian network learning using domain/expert knowledge

Andrés R. Masegosa *, Serafín Moral

*Department of Computer Science and Artificial Intelligence, University of Granada, Spain*

ARTICLE INFO

*Article history:*
Received 13 February 2012
Revised 13 March 2013
Accepted 18 March 2013
Available online 8 April 2013

*Keywords:*
Probabilistic graphical models
Bayesian networks
Interactive structure learning
Domain expert knowledge
Stochastic search

ABSTRACT

Using domain/expert knowledge when learning Bayesian networks from data has been considered a promising idea since the very beginning of the field. However, in most of the previously proposed approaches, human experts do not play an active role in the learning process. Once their knowledge is elicited, they do not participate any more. The interactive approach for integrating domain/expert knowledge we propose in this work aims to be more efficient and effective. In contrast to previous approaches, our method performs an active interaction with the expert in order to guide the search based learning process. This method relies on identifying the edges of the graph structure which are more unreliable considering the information present in the learning data. Another contribution of our approach is the integration of domain/expert knowledge at different stages of the learning process of a Bayesian network: while learning the skeleton and when directing the edges of the directed acyclic graph structure.



## 1. Introduction

Bayesian networks (BN) [26] are a state-of-the-art model for reasoning under uncertainty in the machine learning field. They are especially useful in real-world problems composed by many different variables with a complex dependency structure. Examples of areas where these models have been successfully applied include genomics, text classification, automatic robot control, fault diagnostic, etc. (see [28] for a good review of practical applications).

Every Bayesian network (BN) has a qualitative part and a quantitative part. The qualitative part (i.e., the structure of the BN) consists of a directed acyclic graph (DAG) where the nodes correspond to the variables in the domain problem and the edges between two variables correspond to direct probabilistic dependencies. On the other hand, the quantitative part consists of the specification of the conditional probability distributions that are stored in the nodes of network.

One of the main challenges in this research field is the problem of learning the structure (the qualitative part) of a Bayesian network from a previously given set of observational data. This problem has been the subject of a great deal of research [13,24,32]. In many of these approaches, humans only participate in the definition of the problem and the structural learning is carried out automatically, without human intervention, like usually happens with most of the machine learning models [11]. However, Bayesian networks provide a graphical representation of the dependencies among the variables that can be easily interpreted by humans [26]. This key property opened the possibility of human intervention during the learning process. In fact, from the very beginning of the field [13,6] and until recent years [10,3,2], many approaches have been proposed to introduce domain or expert (d/e) knowledge for boosting the reliability of the automatic learning methods. This is also becoming a new emerging trend in other relevant fields, like gene expression data mining [3] where there is an

* Corresponding author.
*E-mail addresses:* andrew@decsai.ugr.es (A.R. Masegosa), smc@decsai.ugr.es (S. Moral).

0888-613X/$ - see front matter © 2013 Elsevier Inc. All rights reserved.
http://dx.doi.org/10.1016/j.ijar.2013.03.009

growing interest in exploiting the large amount of domain knowledge available in the literature, but especially in knowledge repositories such as KEGG [25] or MIPS [22].

However, the methodologies for introducing d/e knowledge that have been proposed so far demand or request this knowledge *a priori*. In some cases, this knowledge is provided pursuing a Bayesian approach with the definition of informative prior distributions over the model space [6,13,15]. For example, in a recent work [2], authors employ stochastic logic programs to define these priors. In other cases the knowledge is given by defining fixed restrictions in the model space by explicitly enumerating the existence/absence of some edges between the variables [10] or, as proposed in [5], by defining directed path constraints which encode some previously known causal relationships between variables of the domain problem. This kind of knowledge is extracted from experimental data or from some previously built ontology [4]. Other works [29] also explore the problem of unreliable d/e knowledge and propose methods which combine (unreliable) knowledge of independent sources as an effective way to improve the overall quality of the elicited information. But in all of these works the role of the expert in the learning process ends once s/he has provided the required knowledge.

The problem we find in these approaches is that human experts do not receive any support from the learning system when they introduce their d/e knowledge. As mentioned before, this knowledge is usually introduced by giving information about particular edges. But the number of possible edges in a domain problem with many variables is very large, and the elicitation of this knowledge could be quite costly [3] (for any elicited edge, a direct dependency relationship between two variables needs to be asserted). Furthermore, the presence of a particular link is not an isolated event that can be asserted separately from the rest of the graph structure. The simplest restriction is that it is not allowed to create directed cycles, but it could also happen that the existence of a link between two variables depends on the absence of an alternative path joining them. So, context information (what it is already known about the graph structure) can be useful in order to introduce d/e knowledge. What we show in this paper is that many edges can be reliably inferred by simply analyzing the learning data; we do not need to elicit any prior d/e knowledge about them, while other edges remain very uncertain using only the information present in the data sample and introduce a lot of noise in the inferred DAG structure. In this work we argue that the efforts of experts should be focused on these *conflictive* edges in order to boost the quality of the learnt Bayesian networks, and that the certain information already extracted from the data can be useful in this interactive phase.

Following similar ideas, the so-called active learning approaches use experimental data as a complementary source of information to the given observational data [23,33,21]. These methods assume that some variables in the domain problem can be intervened (i.e. their value can be fixed to a predetermined value) when collecting data samples. Hence, the collected data are experimental, not observational. The above reference proposes alternative strategies to decide how to perform these interventions and the number of experimental samples that must be collected. They show that the request of experimental data can be minimized by firstly analyzing the observational data that we have available, and conclude that using experimental data is only worthy if it contains information that is not already present in the observational data. For example, [23] proposed a decision theoretic approach for deciding which interventions should be performed. This decision approach translates to selecting in each step the intervention which most reduces the conditional entropy of the posterior over the graph structures given the experimental data. They apply an online Markov Chain Monte Carlo (MCMC) method to estimate the posterior over the alternative graph structures, and use importance sampling to find the best action to perform in each step.

Our work is along the lines of these last mentioned works. We propose an interactive methodology to identify which edges of the DAG model cannot be reliably inferred with the information present in the given observational data. We then assume that this information can be obtained from an expert (who might not be fully reliable) and integrated in our data learning process. Under this methodology, there is a close direct interaction between the human and the learning system, since the human answers questions submitted by the system, and the system performs the structure learning guided by the information provided by the expert. As mentioned before, one of the main advantages of this interactive procedure is that the system only requests information about those edges whose presence in the inferred model cannot be discerned with the information present in the data. Therefore, this procedure reduces the amount of d/e knowledge that must be requested.

This paper also shows that the integration of d/e knowledge can be carried out at different levels of the model space. In the first level, a skeleton (i.e., an undirected graph which may contain cycles) is learnt with the help of d/e knowledge and then, with the constraint of this initial skeleton, a BN model is inferred, using d/e knowledge as well. We will show that the integration of d/e knowledge in every level boosts the quality of the learnt BN w.r.t. the model inferred using only the information present in the data or integrating d/e knowledge in only one of the levels. This way, we extend the ideas previously presented in [7,8] for this problem. In particular, we remove the restriction imposed by the previously presented methods, where the BN learning process has to be carried assuming that a total order of the variables is previously given. This change is quite relevant, since now the model space is much larger (i.e., from an exponential size space to a super-exponential size space) and the learning problem is much more challenging. In addition to this, the methodology to integrate d/e knowledge is also extended to refine the initial skeleton structure that is inferred to constrain the BN model space. The presented methodology is only developed for multinomial Bayesian networks, but it can be extended to deal with continuous Gaussian variables.

The paper is structured as follows. Section 2 exposes the previous knowledge. Our approach is presented in Section 3 and experimentally validated in Section 4. Finally, Section 5 includes the main conclusions and future work.