

Incorporating expert priors when learning Bayesian networks: Heart failure as a case study[☆]

M. Julia Flores^{*a}, Ann Nicholson^b, Andrew Brunskill^c, Kevin Korb^b, Steven Mascaro^d

^a*Departamento de Sistemas Informáticos SIMD i³A, Universidad de Castilla-La Mancha, Campus Universitario s/n, Albacete, 02071, Spain.*

^b*Information Technology, Monash University, Australia*

^c*University of Washington, USA.*

^d*Bayesian Intelligence Pty Ltd, Australia*

Abstract

Bayesian networks (BNs) [1, 2] are rapidly becoming a leading technology in applied Artificial Intelligence (AI). By combining a graphical representation of the dependencies between variables with probability theory and efficient inference algorithms, BNs provide a powerful and flexible tool for reasoning under uncertainty.

Medicine has undoubtedly been the most popular application area for Bayesian networks to date. It is a complex domain where experienced medical practitioners implicitly hold much knowledge, and as such, has long been a target of expert systems. The popularity of BNs for medical application is based on their ability to explicitly model causal interventions, to reason both diagnostically and predictively and the visualisation of the graphical representation, which assists their use in explanation.

BNs may be built by eliciting expert knowledge or learned via causal discovery programs such as CaMML [3]. Both approaches to building BNs have limitations: expert elicitation is expensive, time-consuming and relies on experts having full domain knowledge, while automated learning is often ineffective given small or noisy datasets. This has led to hybrid approaches which incorporate prior information elicited from experts into the causal discovery process; here we use such a hybrid approach for a medical application case study, namely modeling heart failure.

This paper emphasises the knowledge engineering process, describing all the steps required to produce a useful predictive BN. These include the data pre-processing, attribute selection, the methods used to elicit different kinds of structural information from experts, and how to provide this information to the learning software. We present experimental results showing how the weight given to the expert priors impacts on the resulted learned BNs. We also looked at different ways of analysing and visualising the resultant structures so that these can then be presented back to the experts as part of an iterative development methodology.

Key words: Bayesian networks, machine learning, medicine applications, expert systems, expert priors

References

- [1] J. Pearl, Probabilistic Reasoning in Intelligent Systems, Morgan Kaufmann, San Mateo, Ca., 1988.
- [2] K. B. Korb, A. E. Nicholson, Bayesian Artificial Intelligence, Computer Science and Data Analysis, Chapman & Hall / CRC, Boca Raton, 2004.
- [3] C. Wallace, K. Korb, Learning linear causal models by MML sampling, in: Gammernan (Ed.), Causal Models and Intelligent Data Management, Springer, 1999.

[☆]This work has been partially supported by the Spanish Ministerio de Educación y Tecnología under Project TIN2007-67418-C03-01 and UCLM under Project PL20091291.

*Corresponding author. Tel: +34967599284. Fax: +34967599224

Email addresses: Julia.Flores@uclm.es (M. Julia Flores), Ann.Nicholson@infotech.monash.edu.au (Ann Nicholson), jaberwok@u.washington.edu (Andrew Brunskill), Kevin.Korb@infotech.monash.edu.au (Kevin Korb), Steven.Mascaro@bayesian-intelligence.com (Steven Mascaro)

Preprint submitted to Artificial Intelligence in Medicine