Verifying monotonicity of Bayesian networks with domain experts

Linda C. van der Gaag *, Hermina J.M. Tabachneck-Schijf, Petra L. Geenen

Department of Information and Computing Sciences, Utrecht University, P.O. Box 80.089, 3508 TB Utrecht, The Netherlands

ARTICLE INFO

ABSTRACT

In many realistic problem domains, the main variable of interest behaves monotonically in the observable variables, in the sense that higher values for the variable of interest become more likely with higher-ordered observations. This type of knowledge appears to naturally emerge from experts during knowledge elicitation, without explicit prompting from the knowledge engineer. The experts' concept of monotonicity, however, may not correspond to the mathematical concept of monotonicity in Bayesian networks. We present a method that provides both for verifying whether or not a network exhibits the properties of monotonicity suggested by the experts and for studying the violated properties with the experts. We illustrate the application of our method for a real Bayesian network in veterinary science.

© 2008 Elsevier Inc. All rights reserved.

1. Introduction

For many diagnostic problems, the relation between the output variable of interest and the observable input variables is isotone in the sense that higher values for the input variables give rise to a higher-ordered output for the variable of interest. In a medical diagnostic application, for example, observing symptoms and signs that are more severe will result in a more severe disease becoming a more likely explanation. Such sensations of isotonicity are commonly shared in a domain of application. Upon developing models to support experts in their diagnostic reasoning tasks, it is important that the resulting model reflects any commonly acknowledged sensations of isotonicity. If a model violates one of these sensations, it will exhibit a reasoning behaviour that is counterintuitive to the experts, which is likely to result in a dip in acceptance of the model in daily practice. For many types of model, therefore, techniques have been developed for studying properties of isotonicity, for example for neural networks [1], for classification trees [11], and for regression models [12]. Also for Bayesian networks a mathematical concept of isotonicity has been formulated [4]. A network is said to be isotone in its observable input variables if the probability distribution computed for the output variable given specific observations is stochastically dominated by any such distribution given higher-ordered observations.

Our experiences with developing Bayesian networks in many fields of medicine and veterinary science show that, if sensations of isotonicity are commonly acknowledged in a domain, then experts will naturally produce statements during knowledge elicitation that suggest properties of isotonicity. They in fact will do so without explicit prompting. These experiences corroborate results from educational research which suggest the existence of an intuitive reasoning rule that underlies people's tendency to recognise isotonicity, namely 'more of the input implies more of the output' [13]. As the experts' sense of isotonicity is likely to come from experiential knowledge, however, it is probably heuristic and may not hold for all situations that can possibly be encountered. As a consequence, the mathematical translation of their concept of isotonicity almost inherently differs from the mathematical concept of isotonicity formulated for Bayesian networks. Although the

---

* Corresponding author.
E-mail addresses: linda@cs.uu.nl (L.C. van der Gaag), hermi@cs.uu.nl (H.J.M. Tabachneck-Schijf), petrag@cs.uu.nl (P.L. Geenen).

experts' statements appear to imply properties of isotonicity, therefore, any such properties have to be carefully verified before they can be exploited in the engineering of a network.

Based upon the above considerations, we developed a method for verifying isotonicity of Bayesian networks with domain experts. Given the properties of isotonicity that have been suggested by the experts during knowledge elicitation, the method focuses on a relevant subset of the observable variables. For these variables, a lattice of all possible joint value assignments is constructed, which subsequently is enhanced with probabilistic information about the effects of these assignments on the probability distribution over the output variable of interest. The enhanced lattice then is used for identifying any violations of the suggested properties of isotonicity within the network at hand. The experts subsequently are presented with these violations by means of pairs of vignettes stated in their domain's terminology, and are asked to carefully study the properties of isotonicity that are not matched by the network. Our method thereby provides both for verifying any implied isotonicities in the network and for verifying the experts' statements suggesting isotonicity.

We applied our method for verifying monotonicity to a Bayesian network in veterinary science. In recent years, we developed a network for the early detection of classical swine fever in individual animals. Both the network's structure and its associated probabilities were elicited from two domain experts. During the elicitation interviews, the experts had produced several statements that reflected their sensations of isotonicity. We verified the isotonicities that were suggested for some of the observable input variables with our method. We found a relatively small number of violations of the implied properties of isotonicity in our network and presented these violations to two independent veterinarians. The results from the interviews showed that the network should indeed have been isotone in the observable variables under consideration and that the identified violations were indicative of modelling inadequacies.

The paper is organised as follows. In Section 2, we briefly describe our Bayesian network for classical swine fever. In Section 3, we review the mathematical concept of monotonicity defined for Bayesian networks. In Section 4, we present our method for verifying properties of monotonicity with domain experts. We report on the application of our method in Section 5. The paper ends with our concluding observations in Section 6.

2. A Bayesian network for classical swine fever

In close collaboration with two experts from the Central Institute of Animal Disease Control in the Netherlands, we are developing a Bayesian network for the early detection of classical swine fever in individual pigs. Classical swine fever is an infectious disease of pigs, which has serious socio-economical consequences upon an outbreak. As the disease has a potential for rapid spread, it is imperative that its occurrence is detected in the early stages. The Bayesian network under construction is aimed at supporting veterinary practitioners in the diagnosis of the disease when visiting pig farms because of disease problems of unknown cause.

Classical swine fever is a viral disease. The virus causing the disease is transmitted mainly by direct contact between pigs, yet transmission by farmers is also known to occur. When a pig is infected, the virus first invades the lymphatic system. It subsequently affects the blood vessels and the immune system, which may give rise to haemorrhaging and diminished resistance to secondary infections. The virus will ultimately affect several organs and the pig will die. As a consequence of the infection, a pig will show different disease symptoms, among which are fever, neurological disorders, and skin haemorrhages. Clinical symptoms seen by the farmer or by the veterinarian are usually the first indications of the presence of classical swine fever in a herd. The disease is hard to detect, however, since its early symptoms are rather atypical and are shared to a large extent by common respiratory and gastro-intestinal infections. The disease moreover has a low incidence.

Our Bayesian network for classical swine fever currently includes 42 variables for which over 2400 parameter probabilities have been assessed. The variables in the network model the risk factors and the pathogenesis of the disease. More specifically, the network also models the clinical signs observed in a pig, to provide for diagnosis at a farm site. For the construction of the network, we held one unstructured interview in which the experts were asked to describe the domain, and 11 structured interviews in which the experts were asked detailed questions. In six of these structured interviews, the probabilities required for the network were obtained using standardised forms with questions accompanied by a probability scale containing verbal and numerical anchors [6]. Both experts were present at all interview sessions and consensus was always reached. An initial version of our network has now been completed. The graphical structure of this network is shown in Fig. 1; in the sequel, we will refer to this network as the CSF network. We currently are in the process of studying the network's performance, both by evaluating its output given real data gathered in a field test and by analysing reasoning patterns with artificial data.

3. The concept of monotonicity

Upon reviewing the mathematical concept of monotonicity for Bayesian networks, we assume that the variables of a network have different roles. We assume more specifically that the network includes a single output variable C and one or more observable input variables E; in addition, it may include an arbitrary number of intermediate variables which serve to correctly model the domain's knowledge yet cannot be observed in practice. In the CSF network, for example, the output variable models whether or not a pig has a viraemia of classical swine fever, that is, whether or not the virus has entered into the pig's system of blood vessels. Among other clinical symptoms and signs, the observable input variables include the presence