
Verifying Monotonicity of Bayesian Networks with Domain Experts

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

Department of Information and Computing Sciences, Utrecht University,

P.O. Box 80.089, 3508 TB Utrecht, the Netherlands.

e-mail: {linda,petrag,hermi}@cs.uu.nl

Abstract

In many real 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 arise naturally during knowledge elicitation, without explicit prompting. The monotonicity properties declared by experts, however, may not correspond to the mathematical concept of monotonicity in Bayesian networks. We present a method for verifying, with the help of the experts, whether or not a network exhibits the implied properties of monotonicity. We illustrate the application of our method for a real Bayesian network in veterinary science.

1 INTRODUCTION

In many problem domains, the variables of importance have different roles. Often, a number of observable input variables are distinguished and a single output variable. In a biomedical diagnostic application, for example, the input variables capture the findings from different diagnostic tests and the output variable models the possible diseases. Multiple input variables and a single output variable in fact are typically found in any type of diagnostic problem.

For many diagnostic problems, the relation between the output variable 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 main variable of interest. In a biomedical diagnostic application, for example, observing symptoms and signs that are more severe will result in a more severe disease becoming a more likely explanation. The mathematical concept of *monotonicity in distribution* has been introduced to capture this type of knowledge for Bayesian networks [1]. A network is *isotone in distribution* if the probability distribution computed for the output variable given specific observations is stochastically

dominated by any such distribution given higher-ordered observations.

Experience shows that, if properties of monotonicity are commonly acknowledged in a domain, then experts will naturally produce statements during knowledge elicitation that suggest monotonicity. They do so without explicit prompting, yet also without explicitly using the word ‘monotone’. Although their statements thus appear to imply monotonicity in distribution, it may very well be that their conception of the stated properties differs from the concept of monotonicity in the mathematical sense. Any suggested properties of monotonicity therefore have to be carefully verified.

We have developed a method for verifying monotonicity of Bayesian networks with the help of domain experts. The method focuses on a subset of the observable variables and builds upon a lattice of all possible joint value assignments to these variables. This lattice is enhanced with probabilistic information about the effects of these assignments on the probability distribution over the main variable of interest. The enhanced lattice then is used for identifying any violations of the properties of monotonicity in distribution. The experts subsequently are presented with these violations by means of pairs of vignettes stated in the domain’s terminology, for their careful consideration. The method has been designed so as to ask little time as well as little cognitive effort from the experts in the verification of their statements.

We applied our method for verifying monotonicity to a Bayesian network in veterinary science. In recent years, we developed a network for the detection of classical swine fever in individual animals. Both the network’s structure and its associated probabilities were elicited from two experts. During the elicitation interviews, the experts made several statements that suggested monotonicity. We verified the implied monotonicities with our method. We found a relatively small number of violations of the mathematical properties of monotonicity in our network and presented these violations to two veterinarians. The results from the interviews showed that the network should indeed

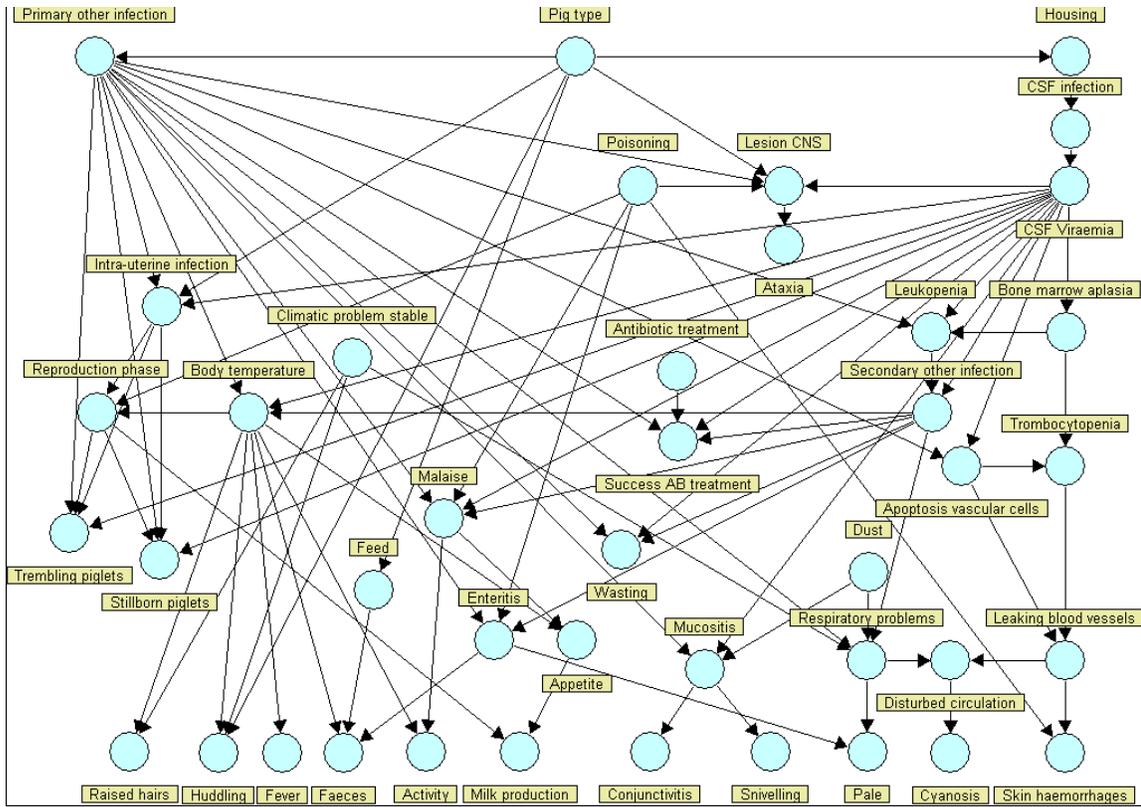


Figure 1: The graphical structure of our Bayesian network for classical swine fever in individual pigs.

have been monotone in distribution 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 in distribution. 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 vis-

iting 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 for example 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 bleedings 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 symptoms are rather atypical and are shared to a large extent by common airways 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 clin-

ical signs observed in a pig, to provide for diagnosis at a farm site. Figure 1 depicts the graphical structure of the current network. In the one and a half years since the beginning of 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 words and numbers [2]. Both experts were present at all interview sessions and consensus was always reached.

An initial version of our network is now completed and we currently are in the process of studying its performance, both by evaluating its output given real data and by analysing reasoning patterns with artificial data.

3 THE CONCEPT OF MONOTONICITY

Upon reviewing the mathematical concept of monotonicity, we assume that the variables of a Bayesian network have different roles. We assume more specifically that the network includes a single output variable C and one or more observable 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. Bayesian networks with multiple input variables and a single output variable are typically found in any type of diagnostic application. Given a joint value assignment e to their set of observable variables, these networks are used for computing the posterior probability distribution $\Pr(C | e)$ for their main variable of interest.

Each variable can adopt one of a set of discrete values. We assume that there exists a total ordering \leq on the set of values for the variable. The total orderings per variable induce a partial ordering \preceq on the set of all joint value assignments to any subset of the network's variables. The concept of *monotonicity in distribution* now builds upon the posterior probability distributions over the output variable given the possible joint value assignments to the observable variables [1]. It is defined in terms of stochastic dominance. For a probability distribution $\Pr(C)$ over the output variable C , the cumulative distribution function F_{\Pr} is defined as $F_{\Pr}(c) = \Pr(C \leq c)$ for all values c of C . For two distributions $\Pr(C)$ and $\Pr'(C)$ over C , associated with $F_{\Pr}(C)$ and $F_{\Pr'}(C)$ respectively, we say that $\Pr'(C)$ is *stochastically dominant* over $\Pr(C)$, denoted $\Pr(C) \leq \Pr'(C)$, if $F_{\Pr'}(c) \leq F_{\Pr}(c)$ for all values c . We now say that a Bayesian network is *isotone in distribution* for its observable variables E if

$$e \preceq e' \rightarrow \Pr(C | e) \leq \Pr(C | e')$$

for all joint value assignments e, e' to E . If

$$e \preceq e' \rightarrow \Pr(C | e) \geq \Pr(C | e')$$

for all e, e' , then the network is said to be *antitone in distribution* for E . Informally speaking, we have that a Bayesian network is isotone in distribution if entering a higher-ordered value assignment to the observable variables cannot make higher-ordered values of the output variable less likely.

Establishing whether or not a Bayesian network is monotone in distribution amounts to verifying that entering a higher-ordered value assignment to the observable variables results in a stochastically dominant probability distribution over the main variable of interest. Unfortunately, the problem of establishing monotonicity is highly intractable [1]. Although an approximate algorithm is available, its complexity tends to forestall use in a practical setting.

4 A METHOD FOR VERIFYING MONOTONICITY

As argued above, experts will naturally produce statements during knowledge elicitation that suggest monotonicity. They typically do so without explicit prompting from the knowledge engineer, yet also without explicitly using the word 'monotone'. An example statement suggesting monotonicity in our domain of application is

Observing ataxia is always indicative of classical swine fever.

Although the experts' statements thus appear to imply monotonicity, it may very well be that their conception of the stated properties differs from the mathematical concept reviewed in the previous section. Based upon the mathematical concept of monotonicity in distribution, we have designed a method for verifying, with the help of domain experts, whether or not a Bayesian network exhibits the implied properties of monotonicity. The method focuses on a specific subset of the observable variables and constructs an *assignment lattice* for all joint value assignments to these variables. The lattice subsequently is enhanced with probabilistic information computed from the Bayesian network under study. From the enhanced lattice, the method identifies any violation of the properties of monotonicity of the output. These violations then are presented to the experts by means of pairs of vignettes that are stated in the experts' terminology, for their careful consideration.

4.1 THE ASSIGNMENT LATTICE AND ITS USE

In detailing the assignment lattice and its use, we restrict the discussion to binary variables, each of which adopts one of the values *true* and *false*. If a variable V has adopted the value *true*, we will write v ; we use \bar{v} to denote $V = \textit{false}$. We take the total ordering \leq with *false* \leq *true* on the two values. The *assignment lattice* for a set X of n observable variables now in essence encodes all joint value assignments to X , along with their partial ordering. For each joint

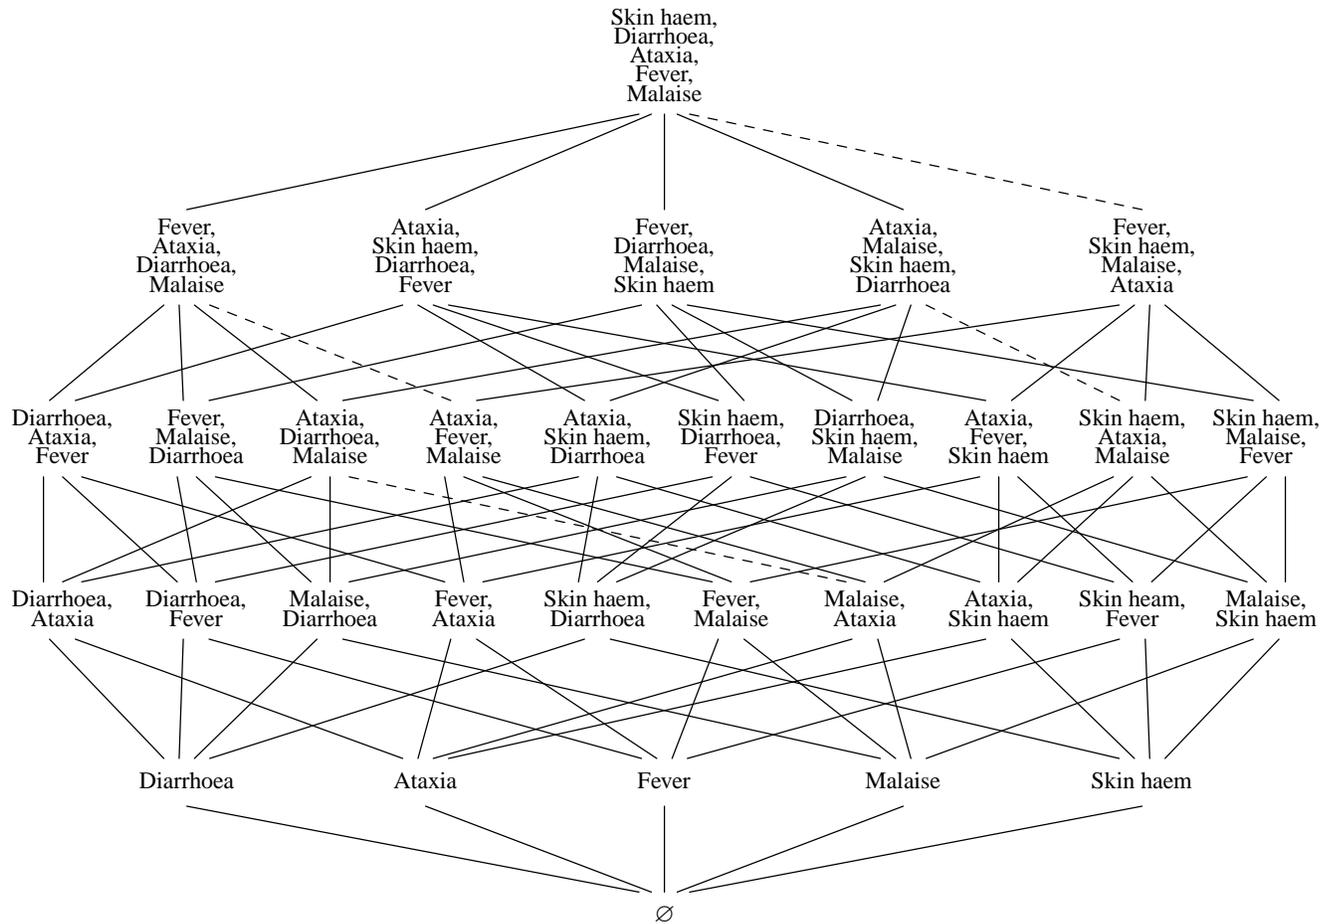


Figure 2: An assignment lattice for our Bayesian network for classical swine fever; the violated monotonicity properties are indicated by dashed lines.

value assignment x to X , we construct a set $L(x) \subseteq X$ such that $X_i \in L(x)$ if and only if $X_i = \text{true}$ occurs in x . From the 2^n possible value assignments to X , 2^n subsets of X are obtained, which with each other constitute the power set of X . From these subsets, we construct a standard lattice: the elements of the lattice are the various subsets of X and the links in the lattice capture the set-inclusion relation between them. We say that a set $L(x)$ *directly precedes* a set $L(x')$ in the lattice if $L(x) \subset L(x')$ and there is no set $L(x'')$ with $L(x) \subset L(x'')$ and $L(x'') \subset L(x')$. Note that the set-inclusion relation of the lattice coincides with the partial ordering \preceq on the joint value assignments to X . The bottom of the assignment lattice is the empty set, denoting the joint value assignment to X in which all variables have adopted the value *false*; the top of the lattice equals the set X , encoding the joint value assignment in which all observable variables have adopted the value *true*. Figure 2 depicts, as an example, the assignment lattice that is constructed for the five observable variables *Diarrhoea*, *Ataxia*, *Fever*, *Malaise*, and *Skin haemorrhages* from our Bayesian network for classical swine fever.

The assignment lattice for the set of observable variables E of a Bayesian network as described above, captures all possible joint value assignments to E along with the partial ordering between them. To describe the effects of the various assignments on the probability distribution over the output variable C , we enhance the lattice with probabilistic information. For each element $L(e)$ of the lattice, the conditional probability $\Pr(c | e)$ is computed from the Bayesian network under study; this probability is associated with the element $L(e)$. We now recall that the network is isotone in distribution for E if entering a higher-ordered value assignment to E results in a stochastically dominant probability distribution over the main variable of interest or, for the binary variable C , in a higher probability of C being *true*. Since the partial ordering \preceq on the value assignments coincides with the set-inclusion relation of the assignment lattice for E , we thus have that the network is isotone in distribution if for each pair of elements $L(e)$ and $L(e')$ in the lattice where $L(e)$ directly precedes $L(e')$, we have for the associated probabilities that $\Pr(c | e) \leq \Pr(c | e')$. We say that the pair of probabilities $\Pr(c | e)$ and $\Pr(c | e')$ *violates* the isotonicity if $\Pr(c | e) > \Pr(c | e')$. Similar observations hold for antitonicity in distribution.

From the above considerations, we have that the enhanced assignment lattice provides directly for verifying monotonicity in distribution of a Bayesian network. We note however, that the lattice encodes an exponential number of value assignments to the set of observable variables. Constructing the lattice and computing the various probabilities to be associated with its elements, therefore, takes exponential time and will be inhibitive for most real networks. The assignment lattice fortunately also provides for studying monotonicity properties of a network for a particular

subset of the set of observable variables. To this end, a subset X of observable variables is selected from the network. An assignment lattice then is constructed from these variables as described above. The probabilities associated with the elements of this lattice are conditioned on a *fixed* joint value assignment e^- to the variables $E^- = E \setminus X$, that is, with each element $L(x)$ of the lattice is associated the conditional probability $\Pr(c | x, e^-)$. The lattice now provides for studying monotonicity for the set X *in the context of* e^- . Both the subset X and the assignment e^- for which the properties of monotonicity are to be verified, are dependent upon the domain under study and should be chosen in close consultation with the experts.

4.2 ELICITING MONOTONICITY PROPERTIES

The monotonicity properties in a domain under study can in essence be elicited directly from experts by exploiting the structure of the assignment lattice introduced above. We will briefly describe how vignettes can be used for this purpose. We will then argue that the tasks involved in the elicitation, although not demanding from a cognitive perspective, are quite time consuming. To reduce the time that is required from the experts, we propose to not elicit monotonicity properties directly, but to verify instead only the violations of these properties that are identified from the assignment lattice.

We have argued above that, if properties of monotonicity are commonly acknowledged in a domain, experts will naturally produce statements suggesting monotonicity during knowledge acquisition, even so without explicit prompting from the knowledge engineer. Since the experts making these statements may have a different conception of the idea of monotonicity than is captured by the mathematical concept reviewed above, these statements have to be carefully verified before they can be further exploited. To study the mathematical properties of monotonicity directly with the experts, we begin by observing that the knowledge to be acquired concerns orderings of conditional probabilities. We consider again a set X of observable variables and an output variable C . For any two joint value assignments x and x' with $x \preceq x'$, the expert more specifically will have to indicate which of the two probabilities $\Pr(c | x)$ and $\Pr(c | x')$ is the largest. By building upon the associated assignment lattice in fact, the expert will only have to perform this task for pairs of value assignments x and x' of which the set $L(x)$ directly precedes the set $L(x')$ in the lattice.

To reduce the cognitive effort that is asked of the experts in the comparison task, we propose to present the various probabilities in a way that is easily accessible by the experts. For each of the two probabilities to be compared, a description of a concrete case is constructed. For the probability $\Pr(c | x, e^-)$, where e^- is the context assign-

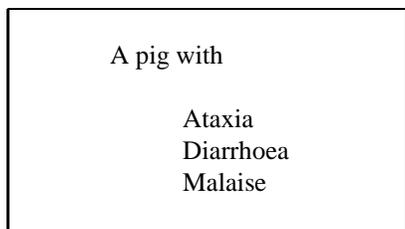


Figure 3: A vignette for the domain of classical swine fever.

ment introduced above, more specifically, a case is constructed with the evidence that is described by x, e^- . For the two constructed cases, separate *case cards* or *vignettes* are created for presentation to the experts. Figure 3 shows, as an example, a vignette describing a value assignment to the five observable variables *Diarrhoea*, *Ataxia*, *Fever*, *Malaise*, and *Skin haemorrhages* from our Bayesian network for classical swine fever. The description on the vignette mentions just the variables for which the value *true* has been observed; the other variables are implicitly taken to be *false*, which fits in with the implicit assumptions humans make in observational reasoning. For each pair of vignettes, the experts are asked to indicate in which of the two described cases the output c is more likely to occur. We would like to mention that in domains in which experts have little experience with probabilities, the constructed cases can be presented using the frequency format [3, 4].

While human experts tend to feel uncomfortable expressing their knowledge and experience in terms of probabilities and are known to provide imperfectly calibrated assessments, they typically are able to state probabilistic information of a semi-numerical or qualitative nature with relative conviction and clarity, and with less cognitive effort [5]. Experts, for example, can often easily indicate which of two probabilities is smallest. In addition to requiring less cognitive effort, such relative judgements tend to be more reliable than direct numerical assessments [6].

Even though the task of comparing two probabilities is not very demanding on the part of the experts, direct elicitation of all properties of monotonicity as outlined above requires a large number of such comparisons. By building upon the assignment lattice for n observable variables in fact,

$$\sum_{i=0}^n \binom{n}{i} \cdot (n - i) > 2^n$$

pairs of probabilities have to be compared. For five observable variables, for example, the experts will have to perform as many as 80 comparisons. Since statements implying monotonicity tend to arise naturally during the elicitation interviews with the domain experts while constructing a Bayesian network, we feel that eliciting all suggested properties of monotonicity explicitly is too time consuming. The elicitation will moreover give the impression of

unnecessary duplication and, as a result, will generate irritation. We therefore propose to apply the method described above only for the pairs of conditional probabilities that are identified from the assignment lattice as violating the monotonicity properties that were implied by the experts in their statements.

5 APPLICATION OF THE VERIFICATION METHOD

During the elicitation interviews, our veterinary experts had made various statements that suggested properties of monotonicity for the observable variables of our Bayesian network for classical swine fever. To study whether or not our network adhered to the mathematical properties implied, we used the verification method described in the previous section for various sets of observable variables. In this section, we review the results that we obtained for one such set.

We consider the five observable variables *Diarrhoea*, *Ataxia*, *Fever*, *Malaise*, and *Skin haemorrhages* from our network and take for our output variable the variable *CSF Viraemia*. From the five observable variables, we constructed an assignment lattice. The lattice, which was depicted in Figure 2, includes $2^5 = 32$ elements to capture all possible joint value assignments to the five variables under study; it further includes 80 direct set-inclusion statements.

Before the lattice could be enhanced with conditional probabilities of the presence of a viraemia of classical swine fever, we had to decide upon the context in which the properties of monotonicity would be verified. For this context, we decided to take the value assignment in which all other observable variables of the network had adopted the value *false*. We chose this assignment since the various clinical signs have a rather small probability of occurrence and moreover it is highly unlikely to find a large number of signs in a single live pig. Our choice of assignment further had the advantage of fitting in with the mental model of humans which presumes signs that are not mentioned explicitly to be absent. Note that, if we would have chosen a different context, we would have had to adapt the vignettes to include the clinical signs that were presumed to be present in the context. Given the chosen context, we computed the various conditional probabilities to be associated with the elements of the assignment lattice.

For each pair of directly related elements from the assignment lattice, we compared the computed conditional probabilities of a viraemia of classical swine fever. We found four violations of the monotonicity properties that had been implied by the experts during the elicitation interviews. We presented the pairs of violating assignments to two veterinarians using vignettes as described in the previous section and asked them to indicate the pig that would be more

likely to have a viraemia of classical swine fever. We thus asked the veterinarians to perform four comparisons rather than the 80 comparisons that would have been necessary for eliciting the properties of monotonicity among the five variables involved directly.

The four violations of monotonicity in distribution that were identified from the assignment lattice all pertained to adding the clinical sign of diarrhoea to the combination of findings of ataxia and malaise. Being confronted with the associated pairs of vignettes, the two veterinarians independently and with conviction indicated that the probability of a viraemia of classical swine fever should increase upon finding the additional sign of diarrhoea. Both veterinarians mentioned that the combination of ataxia and diarrhoea especially pointed to classical swine fever; within the scope of our Bayesian network, they could not think of another disease that would be more likely to give this combination of signs. Through their orderings, both veterinarians indicated that the network should indeed have been monotone in the mathematical sense for the five variables under study in the absence of any other signs. The four identified violations thus were indicative of modelling inadequacies.

During the interviews, the veterinarians suggested that diagnostic reasoning patterns in the domain of infectious animal diseases are not monotone in general. Both could rather easily generate, from their accumulated knowledge and experience, examples in which the output would be neither isotone nor antitone in the various clinical signs observed. As a side remark, one of the veterinarians moreover suggested that for studying monotonicity it would not be necessary to include more than six observable variables, since with more clinical signs a pig would be dead.

6 CONCLUDING OBSERVATIONS

In this paper, we have presented a method for verifying monotonicity of Bayesian networks with the help of domain experts. The method focuses on a subset of the observable variables of a network and builds upon a lattice of possible joint value assignments to these variables. The lattice is enhanced with probabilistic information about the effects of these assignments on the probability distribution over the network's main variable of interest. The enhanced lattice then is used for identifying any violations of monotonicity. The experts subsequently are presented with these violations by means of pairs of vignettes stated in the domain's terminology, for their careful consideration. The method has been designed specifically so as to ask little time as well as little cognitive effort from the experts in the verification of their statements of monotonicity.

The results from applying our method for verifying monotonicity to a real network in veterinary science indicate that it presents a practicable method for studying reasoning patterns in Bayesian networks. We feel in fact that through the availability of our method, it has become worthwhile to devote additional attention to any statements made by experts during knowledge elicitation that appear to imply properties of monotonicity.

To conclude, with our method, we are able to verify either isotonicity or antitonicity for a set of observable variables. Also, our method applies to binary variables only. We currently are extending our method to apply to sets of observable variables of mixed monotonicity and to non-binary variables. In the near future, we hope to be able to further extend our method to include techniques for identifying the parts of a Bayesian network that have to be modified to ensure the required properties of monotonicity.

References

- [1] L.C. van der Gaag, H.L. Bodlaender, A. Feelders. Monotonicity in Bayesian networks. In: M. Chickering, J. Halpern (editors). *Proceedings of the Twentieth Conference on Uncertainty in Artificial Intelligence*, AUAI, Arlington, pp. 569 – 576, 2004.
- [2] L.C. van der Gaag, S. Renooij, C.L.M. Witteman, B. Aleman, B.G. Taal. How to elicit many probabilities. In: K.B. Laskey, H. Prade (editors). *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann Publishers, San Francisco, California, pp. 647 – 654, 1999.
- [3] G. Gigerenzer, U. Hoffrage. How to improve Bayesian reasoning without instruction: frequency formats. *Psychological Review*, pp. 684 – 704, 1995
- [4] E.M. Helsen, L.C. van der Gaag, A.J. Feelders, W.L.A. Loeffen, P.L. Geenen, A.R.W. Elbers. Bringing order into Bayesian-network construction. *Proceedings of the Third International Conference on Knowledge Capture*, ACM Press, New York, pp. 121 – 128, 2005.
- [5] M.A. Meyer, J.M. Booker. *Eliciting and Analyzing Expert Judgment: A Practical Guide*. ASA-SIAM Series on Statistics and Applied Probability, 2001.
- [6] M.G. Morgan, M. Henrion. *Uncertainty, a Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. Cambridge University Press, Cambridge, 1990.