

# Sensitivity Analysis for Threshold Decision Making with Bayesian Belief Networks<sup>\*</sup>

Linda C. van der Gaag<sup>1</sup> and Veerle M.H. Coupé<sup>2</sup>

<sup>1</sup> Department of Computer Science, Utrecht University, P.O. Box 80.089, 3508 TB Utrecht, The Netherlands; [linda@cs.uu.nl](mailto:linda@cs.uu.nl)

<sup>2</sup> Center for Clinical Decision Sciences, Erasmus University Rotterdam, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands; [coupe@mgz.fgg.eur.nl](mailto:coupe@mgz.fgg.eur.nl)

**Abstract.** The probability assessments of a Bayesian belief network generally involve inaccuracies. These inaccuracies influence the reliability of the network's output. An integral part of investigating the reliability of output is to study robustness. Robustness pertains to the extent to which varying the probability assessments of a Bayesian belief network influences its output. Output robustness is studied by subjecting the network to a sensitivity analysis. In this paper, we address the issue of robustness of a belief network's output in view of the threshold model for decision making. We present a method for sensitivity analysis that provides for the computation of bounds between which a network's assessments can be varied without inducing a change in recommended decision.

## 1 Introduction

Bayesian belief networks are widely accepted in artificial intelligence as intuitively appealing representations of domain knowledge. A Bayesian belief network basically is a concise representation of a joint probability distribution. It encodes, in a qualitative, graphical part, the variables of importance in the domain that is being represented, along with their probabilistic interrelationships; the strengths of these relationships are quantified by conditional probabilities, that with each other constitute the network's quantitative part. The increasing number of knowledge-based systems that build upon the framework of Bayesian belief networks for knowledge representation and inference, clearly demonstrate its usefulness for addressing complex real-life problem domains in which uncertainty is predominant. Most notably, applications are being realised in the medical domain, for diagnosis, prognostic assessment, and treatment planning.

Bayesian belief networks are generally constructed with the help of experts from the domain of application. Experience shows that, although it may require considerable effort, building the qualitative part of a belief network is quite practicable. In fact, as it has parallels to designing a domain model for a more traditional knowledge-based system, well-known knowledge-engineering techniques

---

<sup>\*</sup> The investigations were (partly) supported by the Netherlands Computer Science Research Foundation with financial support from the Netherlands Organization for Scientific Research (NWO).

can be employed. Constructing a belief network's quantitative part is generally considered a far harder task, not in the least because it tends to be much more time-consuming. It amounts to assessing various conditional probabilities for the variables represented in the network's qualitative part. Although, for most application domains, probabilistic information is available from literature or from statistical data, it often turns out that this information does not provide for estimating all probabilities required [DV95]. For most domains, therefore, many probabilities remain to be assessed by domain experts. Upon eliciting judgemental probabilities from experts, various problems of bias and poor calibration are typically encountered [KST82]. The probability assessments obtained for a belief network as a consequence tend to be inaccurate.

The inaccuracies in the probability assessments of a Bayesian belief network influence the reliability of its output. An integral part of investigating the reliability of a network's output is to study its robustness. Robustness pertains to the extent to which varying the probability assessments of the network influences its output. For gaining detailed insight in output robustness, a Bayesian belief network can be subjected to a sensitivity analysis [CV98]. Sensitivity analysis is a general technique for investigating the effects of the inaccuracies in the parameters of a mathematical model on the model's output [MH90]. The basic idea of performing a sensitivity analysis of a belief network is to systematically vary the assessments for the network's conditional probabilities and study the effects on the output. Upon such an analysis, some conditional probabilities will show a considerable impact, while others will hardly reveal any influence.

In this paper, we address the issue of output robustness of a Bayesian belief network in view of applications in which the output is used for decision making. To this end, we focus on the threshold model for decision making. Although generally applicable, this model is used most notably for patient management in medical applications [PK80]. With the threshold model, an attending physician decides whether or not to gather additional information from diagnostic tests and whether or not to give treatment based upon the probability of disease for a patient under consideration. The robustness of the output of a belief network now pertains not just to the probability of disease computed from the network, but also to the decision for patient management based upon it. For some conditional probabilities, varying their assessment may have a considerable effect on the probability of disease and yet not induce a change in patient management; for other probabilities, variation may have little effect on the probability of disease and nonetheless result in a different management decision. Studying the effects of varying the assessments for the network's conditional probabilities on the probability of disease therefore no longer suffices for establishing robustness: the effects on the recommended decision need also be taken into consideration.

To provide for studying output robustness of a Bayesian belief network in view of the threshold model for decision making, we enhance the basic method of sensitivity analysis with the computation of upper and lower bounds between which a belief network's assessments can be varied without inducing a change in recommended decision. Informally speaking, the more a belief network's proba-

bility assessments can be varied, the more robust the decision based upon the network is.

The paper is organised as follows. In Section 2, we briefly review the formalism of Bayesian belief networks. In Section 3, we outline the threshold model for decision making. In Section 4, we detail the basic method of sensitivity analysis and its enhancement for threshold decision making. The paper ends with some conclusions and directions for further research in Section 5.

## 2 Bayesian Belief Networks


A *Bayesian belief network* basically is a representation of a joint probability distribution on a set of statistical variables [Pea88]. It consists of a qualitative part and an associated quantitative part. The network's qualitative part takes the form of an acyclic directed graph, or digraph, for short. Each node in this digraph represents a statistical variable that takes its value from a finite set of discrete values. In this paper we will restrict the discussion to binary variables, taking one of the values *true* and *false*. If a variable  $V$  has the value *true*, we will write  $v$ ; the notation  $\neg v$  is used to indicate that  $V = \textit{false}$ . The arcs in the digraph represent the influential relationships among the represented variables. The tail of an arc indicates the cause of the effect at the head of the arc. Absence of an arc between two variables means that these variables do not influence each other directly and, hence, are conditionally independent.

For our running example we consider the following fragment of (fictitious and incomplete) medical information, adapted from [Coo84]:

Consider a primary tumour with an uncertain prognosis in an arbitrary patient. The cancer can metastasize to the brain and to other sites. We are interested in the course of the cancer within the next few years, especially with regard to the development of a brain tumour and its associated problems. Metastatic cancer (denoted  $MC$ ) may be detected by an increased level of serum calcium ( $ISC$ ). The presence of a brain tumour ( $B$ ) may be established from a CT scan ( $CT$ ). Severe headaches ( $SH$ ) are indicative of the presence of a brain tumour. Both a brain tumour and an increased level of serum calcium are likely to cause a patient to fall into a coma ( $C$ ) in due course.

In this fragment of information, six statistical variables are identified. The influential relationships among these variables are encoded in the digraph depicted in Figure 1. The digraph for example reflects, by means of the arc  $B \rightarrow SH$ , that the presence of a brain tumour is a possible cause of severe headaches.

The relationships among the variables that are represented in the qualitative part of a Bayesian belief network basically are probabilistic dependences. The strengths of these dependences are described by conditional probabilities: for each variable, the probabilities of its values are specified conditional on the various possible combinations of values for its parents in the digraph. For our



**Fig. 1.** The digraph of an example belief network; it expresses information concerning the presence of a brain tumour and its associated problems in an arbitrary patient.

running example, we assume the following probabilities:

$$\begin{array}{lll}
 p(b \mid mc) = 0.20 & p(mc) = 0.20 & p(ct \mid b) = 0.95 \\
 p(b \mid \neg mc) = 0.05 & & p(ct \mid \neg b) = 0.10 \\
 & p(c \mid b, isc) = 0.80 & \\
 p(isc \mid mc) = 0.80 & p(c \mid \neg b, isc) = 0.80 & p(sh \mid b) = 0.80 \\
 p(isc \mid \neg mc) = 0.20 & p(c \mid b, \neg isc) = 0.80 & p(sh \mid \neg b) = 0.60 \\
 & p(c \mid \neg b, \neg isc) = 0.05 & 
 \end{array}$$

The probabilities specified for the variable *ISC*, for example, express that knowing whether or not metastatic cancer is present has a considerable influence on the probability of an increased level of serum calcium in an arbitrary patient. The relationship between metastatic cancer and increased total serum calcium therefore is a strong dependence. On the other hand, severe headaches are expressed as quite common in both patients with and without a brain tumour. Severe headaches thus have a low predictive value for a brain tumour. The probabilities with each other constitute the network's quantitative part.

The qualitative and quantitative parts of a Bayesian belief network with each other uniquely define a joint probability distribution. A belief network therefore allows for the computation of any probability pertaining to its variables. For this purpose, various algorithms are available, that provide for computing probabilities of interest and for processing evidence, that is, for entering observations into the network and subsequently computing the revised probability distribution given these observations [Pea88,LS88]. The details of these algorithms are not relevant for the present paper.

### 3 Threshold Decision Making

In the medical domain, Bayesian belief networks are often used for diagnostic purposes. A diagnostic belief network typically comprises one or more variables modeling the presence or absence of disease, various variables modeling findings and results from diagnostic tests, and a number of intermediate variables


modeling unobservable pathophysiological states. In our example network, for instance, the variable  $B$  models the disease of interest, being the presence or absence of a brain tumour; the variable  $MC$  models an unobservable state and the remaining variables capture findings and test results. A medical diagnostic belief network is used for computing a most likely diagnosis for a patient given his or her presentation findings and test results.

The most likely diagnosis for a patient, along with its uncertainty, is generally taken by an attending physician to decide upon management of the patient. The physician may decide, for example, to start treatment rightaway. For our running example, the physician may decide to perform neurosurgery if a brain tumour is indicated. Alternatively, the physician may defer the decision whether or not to treat the patient until additional diagnostic information has become available, for example from a CT scan. Or, the physician may decide to withhold treatment altogether. To support choosing among these decision alternatives, the threshold model for patient management can be used.

The *threshold model for patient management*, or for decision making more in general, builds upon various threshold probabilities of disease [PK80]. The *treatment threshold probability* of disease, written  $P^*(d)$  for disease  $d$ , is the probability at which an attending physician is indifferent between giving treatment and withholding treatment. If, for a specific patient, the probability of disease  $\text{Pr}(d)$  exceeds the treatment threshold probability, that is, if  $\text{Pr}(d) > P^*(d)$ , then the physician will decide to treat the patient as if the disease were known to be present with certainty. Alternatively, if  $\text{Pr}(d) \leq P^*(d)$ , the physician will basically withhold treatment from the patient.

As a consequence of the uncertainty concerning the presence of disease in a patient, additional information from a diagnostic test may affect an attending physician's basic management decision. If the probability of disease exceeds the treatment threshold probability, then interpreting a negative test result may result in an updated probability of disease *below* the threshold probability. Alternatively, if the pretest probability of disease falls below the treatment threshold probability, a positive test result may raise the probability of disease to a value *above* the threshold probability. To reckon with such effects, the threshold model for patient management includes another two threshold probabilities. The *no treatment-test threshold probability* of disease, written  $P^-(d)$ , is the probability at which the physician is indifferent between the decision to withhold treatment and the decision to obtain additional diagnostic information. The *test-treatment threshold probability* of disease, written  $P^+(d)$ , is the probability at which the physician is indifferent between obtaining additional information and starting treatment rightaway.

Figure 2 summarises the basic idea of the threshold model for patient management. As long as the diagnostic test under consideration has not been performed, a physician has three decision alternatives at his or her disposal. If the probability of disease  $\text{Pr}(d)$  for a patient falls below the no treatment-test threshold probability, that is, if  $\text{Pr}(d) < P^-(d)$ , then the physician will withhold treatment from the patient without gathering additional diagnostic information.



**Fig. 2.** The threshold model for patient management, indicating three threshold probabilities and the various decision alternatives at a physician's disposal.

If the probability of disease exceeds the test-treatment threshold probability, that is, if  $\Pr(d) > P^+(d)$ , then the physician will start treatment rightaway. Otherwise, that is, if  $P^-(d) \leq \Pr(d) \leq P^+(d)$ , the physician will perform the diagnostic test. After testing, there are only two decision alternatives left. If the updated probability of disease for the patient exceeds the treatment threshold probability, the physician will start treatment; otherwise, treatment will be withheld from the patient.

The treatment threshold probability of disease  $P^*(d)$  used in the threshold model is typically established by a physician after carefully weighing the various utilities involved. These utilities pertain to the presence or absence of disease on the one hand and giving or withholding treatment on the other hand. From the expected utilities for giving and withholding treatment in view of the uncertainty concerning the presence of disease, the probability of disease at which the physician is indifferent between the two decision alternatives is readily determined. For our running example, the physician will typically take into consideration the life expectancy for a patient, with and without a brain tumour, and the patient's attitude towards impaired health states; we assume that the physician sets the treatment threshold probability of a brain tumour at 0.15. The two threshold probabilities  $P^-(d)$  and  $P^+(d)$  for deciding whether or not to perform a diagnostic test are established from the test's characteristics. For our example, the physician will typically weigh the discomfort of a CT scan for a patient against the additional diagnostic information yielded by the scan; we assume that the physician sets the no treatment-test threshold probability of a brain tumour at 0.045 and the test-treatment threshold probability at 0.56.

Although we have discussed the threshold model for decision making in a medical context, we would like to note that the model's use is not restricted to the medical domain but in fact is generally applicable.

## 4 Sensitivity Analysis for Threshold Decision Making


In our introduction, we have argued that the various probability assessments of a Bayesian belief network tend to be inaccurate. To gain insight into the effects of the inaccuracies involved, a belief network can be subjected to a sensitivity analysis. In Section 4.1, we outline sensitivity analysis of a Bayesian belief network with regard to a probability of interest. In Section 4.2, we address sensitivity analysis of a belief network in view of threshold decision making.

#### 4.1 Sensitivity analysis of a Bayesian belief network


For a Bayesian belief network, sensitivity analysis amounts to systematically varying the assessments for the network’s conditional probabilities and investigating the effects on a probability of interest [CV98]. In essence, for every conditional probability of the network, a number of deviations from the original assessment are investigated. For every investigated value, the probability of interest is computed from the network. The results thus obtained reflect the probability of interest as a function of the conditional probability under study.

We illustrate performing a sensitivity analysis of our example belief network, taking the prior probability of the presence of a brain tumour in an arbitrary patient for the probability of interest. The effects of varying the assessments for the probabilities  $p(mc)$  and  $p(b | \neg mc)$  on this probability of interest are shown in Figure 3. Figure 3(a) shows that systematically varying, from 0 to 1, the assessment for the probability  $p(mc)$  of the presence of metastatic cancer has a rather small effect on the probability of interest:  $\Pr(b)$  increases from 0.05 to 0.20. Figure 3(b) shows that varying the assessment for the conditional probability  $p(b | \neg mc)$  of the presence of a brain tumour in the absence of metastatic cancer has a much stronger effect:  $\Pr(b)$  now ranges from 0.04 to 0.84. As long as no further information is available about the inaccuracies in the assessments for the two probabilities under study, we conclude that the probability of interest is more robust with regard to inaccuracy in the assessment for the probability  $p(mc)$  than with regard to inaccuracy in the assessment for  $p(b | \neg mc)$ .

A sensitivity analysis of a Bayesian belief network with regard to a prior probability of interest allows for assessing the robustness of the network in its reflecting a prior probability distribution for the domain of application. In the presence of case-specific observations, however, a belief network may very well show different sensitivities. To reveal these, a sensitivity analysis can be performed with regard to a *posterior* probability of interest. Such an analysis



**Fig. 3.** A sensitivity analysis of the example belief network; the effects of varying the assessments for the probabilities  $p(mc)$ , (a), and  $p(b | \neg mc)$ , (b), on the prior probability of disease  $\Pr(b)$  are shown.



**Fig. 4.** A sensitivity analysis of the example belief network; the effects of varying the assessments for the conditional probabilities  $p(b | \neg mc)$ , (a), and  $p(sh | \neg b)$ , (b), on the posterior probability of disease  $\Pr(b | sh)$  are shown.

allows for investigating the robustness of the network’s output for specific cases or for case profiles.

We once again perform a sensitivity analysis of our example belief network, this time taking for the probability of interest the *posterior* probability  $\Pr(b | sh)$  of the presence of a brain tumour in a patient who is *known* to suffer from severe headaches. By doing so, we assess the robustness of the diagnosis of a brain tumour for an arbitrary patient with severe headaches. The effects of varying the assessments for the conditional probabilities  $p(b | \neg mc)$  and  $p(sh | \neg b)$  on the posterior probability of interest are shown in Figure 4. Figure 4(a) reveals that the observation that a patient suffers from severe headaches has little impact on the robustness of the probability of disease with regard to inaccuracy in the assessment for  $p(b | \neg mc)$ . Figure 4(b) shows that varying the assessment for the conditional probability  $p(sh | \neg b)$  can have a considerable effect on the posterior probability of disease; small deviations from the original assessment 0.60 have little effect, however.

In a sensitivity analysis of a Bayesian belief network, the relation between a probability of interest and a conditional probability under study is fully determined by the network. More specifically, the independences reflected by the network’s qualitative part constrain the relation to a simple mathematical function. In general, the probability of interest relates to a conditional probability under study as a quotient of two linear functions. For the posterior probability of interest  $\Pr(d | o)$ , given some observations  $o$ , and a conditional probability  $x$ , we have that

$$\Pr(d | o) = \frac{a \cdot x + b}{e \cdot x + f}$$

where  $a, b, e$ , and  $f$  are constants. For an example, we reconsider Figure 4(b) showing for our belief network the posterior probability of interest  $\Pr(b | sh)$  as



a function of the conditional probability  $x = p(sh | \neg b)$ ; the function equals

$$\Pr(b | sh) = \frac{0.06957}{x + 0.06957}$$

If the conditional probability under study pertains to a variable without any observed descendants, that is an ancestor of the variable of interest in the network's qualitative part, the mathematical function reduces to a linear function. For the probability of interest and a conditional probability  $x$  as indicated, we then have that

$$\Pr(d | o) = a \cdot x + b$$

where  $a$  and  $b$  are constants. In particular, a prior probability of interest relates linearly to any conditional probability from the network. For an example, we reconsider Figure 3(b) showing for our belief network the prior probability of interest  $\Pr(b)$  as a function of the conditional probability  $x = p(b | \neg mc)$ ; the function equals

$$\Pr(b) = 0.8 \cdot x + 0.04$$

The constants in the mathematical functions involved in a sensitivity analysis of a Bayesian belief network are readily determined by computing the probability of interest from the network for a small number of values for the conditional probability under study and solving the resulting system of linear equations. For further technical details, we refer the reader to [CV98].

#### 4.2 Sensitivity analysis in view of threshold decision making

Sensitivity analysis of a Bayesian belief network with regard to a probability of interest yields a functional relation between this probability of interest and every single conditional probability from the network. The relation indicates how the probability of interest will shift upon varying the assessment for the conditional probability under study. For a probability of interest that is used in the threshold model for decision making, not every shift is significant. In fact, only a shift that results in a different decision recommendation is of interest. In sensitivity analysis in view of threshold decision making, therefore, we have to take the various threshold probabilities employed into consideration. To this end, we enhance sensitivity analysis of a Bayesian belief network with the computation of upper and lower bounds between which the network's assessments can be varied without inducing a change in decision.

The computation of bounds on variation of a belief network's probability assessments builds upon the mathematical functions that we have detailed before, relating a probability of interest to the network's conditional probabilities. Once again focusing on patient management, we begin by considering a probability of disease  $\Pr(d | o)$  and a conditional probability  $x$  to which it is linearly related, that is, we have

$$\Pr(d | o) = a \cdot x + b$$

for some constants  $a$  and  $b$ . For ease of exposition, we assume that  $\Pr(d | o)$  increases with increasing values for  $x$ ; we will return to this assumption presently. If, in view of the threshold model, the probability of disease indicates withholding treatment, that is, if  $\Pr(d | o) < P^-(d)$ , then the decision will remain unaltered as long as the value of the conditional probability  $x$  is smaller than the value  $x^-$  that is computed from

$$a \cdot x^- + b = P^-(d)$$

More precisely, the decision to withhold treatment remains unaltered for any value of the conditional probability  $x$  within the interval  $(-\infty, x^-) \cap [0, 1]$ . If the probability of disease on the other hand indicates starting treatment rightaway, that is, if  $\Pr(d | o) > P^+(d)$ , then the decision will remain unaltered as long as the value of the conditional probability  $x$  is greater than the value  $x^+$  that is computed from

$$a \cdot x^+ + b = P^+(d)$$

More precisely, the decision to start treatment rightaway remains unaltered for any value of the conditional probability  $x$  within the interval  $(x^+, +\infty) \cap [0, 1]$ . If the probability of disease indicates testing, that is, if  $P^-(d) \leq \Pr(d | o) \leq P^+(d)$ , then this decision will be the same for any value of the conditional probability  $x$  within the interval  $[x^-, x^+] \cap [0, 1]$ .

So far, we have addressed the computation of bounds on the variation of a conditional probability that is related linearly to the probability of disease. For a conditional probability that is related to the probability of disease by a quotient of two linear functions, bounds on variation are computed in a similar fashion. We have further assumed so far that the probability of disease increases with increasing values of the conditional probability  $x$  under study. With this assumption, we have implicitly assumed that  $x^- \leq x^+$ . For a conditional probability  $x$  of which increasing values serve to decrease the probability of disease, we have that  $x^- \geq x^+$ . Using this observation, the bounds derived above are readily adjusted.

We illustrate the computation of bounds on variation for our example belief network; we recall from Section 3 that the three threshold probabilities of disease have been set at  $P^*(b) = 0.15$ ,  $P^-(b) = 0.045$ , and  $P^+(b) = 0.56$ . We begin by addressing the robustness of the decision for management of an arbitrary patient. From our belief network, the prior probability of disease is computed to be  $\Pr(b) = 0.08$ . For this probability, we have that  $P^-(b) \leq \Pr(b) \leq P^+(b)$ . Using the threshold model for patient management, therefore, the physician will decide to gather additional information from a CT scan. We investigate the robustness of this decision by computing an upper and lower bound on variation of the assessment for the conditional probability  $x = p(b | \neg mc)$ . The lower bound  $x^-$  on variation is computed from

$$\Pr(b) = 0.8 \cdot x^- + 0.04 = 0.045$$

yielding  $x^- = 0.00625$ ; the upper bound  $x^+$  on variation is computed from

$$\Pr(b) = 0.8 \cdot x^+ + 0.04 = 0.56$$

yielding  $x^+ = 0.65$ . For any value of the conditional probability  $p(b \mid \neg mc)$  within the interval  $[0.00625, 0.65]$ , therefore, the decision to gather additional diagnostic information will remain unaltered. Since the conditional probability under study has been assessed at 0.05, we conclude that the decision is fairly robust with regard to this assessment; variation of the assessment to smaller values, however, may change the decision to the recommendation to withhold treatment without gathering additional diagnostic information.

To conclude, we address the robustness of the management decision for a patient with a primary tumour who is known to suffer from severe headaches. From our belief network, the posterior probability of disease is computed to be  $\Pr(b \mid sh) = 0.1039$ . For this probability, we observe that  $P^-(b) \leq \Pr(b \mid sh) \leq P^+(b)$ . The physician will therefore order a CT scan for the patient. We investigate the robustness of this decision by computing the upper and lower bound on variation of the assessment of the conditional probability  $x = p(sh \mid \neg b)$ . Note that the probability of disease decreases with increasing values for this conditional probability. The lower bound  $x^+$  on variation is computed from

$$\Pr(b \mid sh) = \frac{0.06957}{x^+ + 0.06957} = 0.56$$

yielding  $x^+ = 0.1938$ . Upon computing the upper bound  $x^-$  on variation, we find a value greater than one. For any value of the conditional probability  $p(sh \mid \neg b)$  within the interval  $[0.1938, 1]$ , therefore, the decision to gather additional diagnostic information for the patient will remain unaltered. Since the conditional probability under study has been assessed at 0.60, we conclude that the decision is quite robust with regard to this assessment.

## 5 Conclusions

The probability assessments of a Bayesian belief network tend to be inaccurate. The belief network as a consequence will yield inaccurate output. If the network's output is used for decision making, its inaccuracy influences the reliability of a decision that is based upon it. An integral part of investigating reliability is to study output robustness. To investigate the robustness of a belief network's output in view of threshold decision making, we have presented an enhanced method for sensitivity analysis that provides for the computation of upper and lower bounds between which a network's assessments can be varied without inducing a change in decision.

We have addressed the issue of robustness in view of a simplified threshold model for decision making, involving binary variables and a single diagnostic test. The more general threshold model addresses variables that have multiple discrete values and provides for selecting among multiple tests. Our method of

sensitivity analysis will be further elaborated upon for use with this more general model. Although often used in practice, the threshold model is a simple model for decision making. With a Bayesian belief network, more complex models can be used. More specifically, a belief network can be extended to an influence diagram to provide for addressing more elaborate trade-offs in decision making. The results put forward in this paper hold unabatedly for influence diagrams. In the near future, we hope to extend our method of sensitivity analysis for decision making to apply to influence diagrams.

## References

- [Coo84] G.F. Cooper (1984). *NESTOR: a Computer-based Medical Diagnostic Aid that Integrates Causal and Probabilistic Knowledge*, Report HPP-84-48, Stanford University.
- [CV98] V.M.H. Coupé, L.C. van der Gaag (1998). Practicable sensitivity analysis of Bayesian belief networks. In: M. Hušková, P. Lachout, J.A. Víšek (Editors). *Prague Stochastics '98*, Union of Czech Mathematicians and Physicists, Prague, pp. 81 – 86; also available as Report UU-CS-1998-10, Utrecht University.
- [DV95] M.J. Druzdzel, L.C. van der Gaag (1995). Elicitation of probabilities for belief networks: combining qualitative and quantitative information. In: P. Besnard, S. Hanks (Editors). *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann, Palo Alto, pp. 141 – 148.
- [LS88] S.L. Lauritzen, D.J. Spiegelhalter (1988). Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society, Series B*, vol. 50, pp. 157 – 224.
- [MH90] M.G. Morgan, M. Henrion (1990). *Uncertainty, a Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*, Cambridge University Press, Cambridge.
- [Pea88] J. Pearl (1988). *Probabilistic Reasoning in Intelligent Systems. Networks of Plausible Inference*, Morgan Kaufmann, Palo Alto.
- [PK80] S.G. Pauker, J.P. Kassirer (1980). The threshold approach to clinical decision making. *New England Journal of Medicine*, vol. 302, pp. 1109 – 1117.
- [KST82] D. Kahneman, P. Slovic, A. Tversky (1982). *Judgement under Uncertainty: Heuristics and Biases*, Cambridge University Press, Cambridge.