Bringing Order into Bayesian-Network Construction

E.M. Helsper[1], L.C. van der Gaag[1], A.J. Feelders[1], W.L.A. Loeffen[2], P.L. Geenen[1], A.R.W. Elbers[2]

1. Dept. of Information and Computing Sciences, Utrecht University, Utrecht, The Netherlands
{eveline,linda,ad,petrag}@cs.uu.nl
2. Dept. of Virology, Central Institute for Animal Disease Control, Wageningen UR, Lelystad, The Netherlands
{willie.loeffen,armin.elbers}@wur.nl

ABSTRACT

Among the tasks involved in building a Bayesian network, obtaining the required probabilities is generally considered the most daunting. Available data collections are often too small to allow for estimating reliable probabilities. Most domain experts, on the other hand, consider assessing the numbers to be quite demanding. Qualitative probabilistic knowledge, however, is provided more easily by experts. We propose a method for obtaining probabilities, that uses qualitative expert knowledge to constrain the probabilities learned from a small data collection. A dedicated elicitation technique is designed to support the acquisition of the qualitative knowledge required for this purpose. We demonstrate the application of our method by quantifying part of a network in the field of classical swine fever.

Categories and Subject Descriptors: I.2.6 [Artificial Intelligence]: Learning – *Knowledge acquisition, parameter learning*

General Terms: Design

Keywords: Bayesian networks, interview techniques, probability estimation, uncertainty

1. INTRODUCTION

The formalism of Bayesian networks is used more and more often for representing knowledge in domains of application in which uncertainty is predominant [7]. A Bayesian network is a concise representation of a joint probability distribution, consisting of a graphical part and an associated numerical part. The graphical part represents the relevant variables in the domain, along with their interrelations. The numerical part models the strengths of these relations, by means of conditional probability distributions.

The construction of Bayesian networks is usually done with the help of domain experts. The task of eliciting the variables of importance as well as their interrelations from experts is comparable, to at least some extent, to knowledge engineering for other artificial-intelligence representations [5]. Although it may require significant effort, it is generally considered doable. The task of obtaining all required probabilities, on the other hand, is highly specific for probabilistic models and is commonly acknowledged to be more daunting.

Obtaining the probabilities required for a Bayesian network is relatively easy in data-rich application domains, where large data collections are available. These collections retrospectively document every-day problem solving and thereby contain highly valuable information about the relations between the variables in the domain. A comprehensive data collection is expected to embed all intricacies in the domain and to faithfully reflect the probability distribution over the variables. If such a data collection is available, therefore, reliable probabilities for the network under construction can be extracted from the data.

In many application domains, unfortunately, no comprehensive data collections are available. Often, at least some data have been collected, but these data are few. These data may moreover be biased by the data-collection strategies used or by the purpose for which they have been collected. With an insufficiently large data collection, the various subsets from which probabilities are estimated can be empty or too small to allow for meaningful numbers. Since the probabilities obtained from a relatively small data collection tend to be insufficiently reliable, common practice in building Bayesian networks is to not use the data at all and to exploit the knowledge and experience of experts in the domain of application as the only source of probabilistic informa-

K-CAP'05, October 2–5, 2005, Banff, Alberta, Canada.

tion. The problems encountered when directly eliciting probabilities from human experts, however, are widely known [8, 9]. In addition to the task being hard on the part of the expert and highly time consuming, the assessments obtained tend not to be properly calibrated.

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 [2, 11]. 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 [9].

Building upon the above observations, we designed a new method for obtaining the probabilities required for a Bayesian network. The basic idea of the method is to combine qualitative information about probabilities, obtained from domain experts, with the numerical information that is available from a (relatively small) data collection. Informally speaking, the elicited qualitative knowledge, which in essence is a probability order, serves as a framework that is filled in quantitatively by numbers estimated from the data. By thus combining the two types of information, we circumvent the need of directly eliciting numbers from experts: they just have to provide qualitative information about the probabilities. By exploiting the elicited order to constrain the estimates obtained from the data, the resulting probabilities match the experts' intuitions. The constructed network, as a consequence, is less likely to exhibit reasoning behaviour that is counterintuitive to the experts, leading to a higher acceptance.

In this paper, we describe our method for combining expert knowledge with data. Our method consists of two phases. First the qualitative knowledge, which in essence is a probability order, is elicited from a domain expert. The elicited order is then used, in the second phase, for constraining the probabilities to be estimated from a data collection. To support the knowledge engineer in the acquisition of probability orders, we designed a dedicated elicitation technique. We demonstrate the application of our method by quantifying part of a network in the field of classical swine fever.

The remainder of the paper is structured as follows. Section 2 briefly introduces Bayesian networks and the domain of classical swine fever. Section 3 presents our technique for eliciting qualitative knowledge. Section 4 shows how the elicited knowledge is used to constrain the estimated probabilities. Our experiences with the method in the field of classical swine fever are discussed in Section 5. The paper ends with our conclusions.

2. PRELIMINARIES

We briefly review the formalism of Bayesian networks and introduce the network that we are currently developing for the early detection of classical swine fever. Classical swine fever is an infectious viral disease of pigs, which may have serious socio-economic consequences upon an outbreak in pig herds. The network is being built in close cooperation with three domain experts, who are the last three authors of the present paper. Part of the network pertains to the consistency of faeces in pigs. The faeces of pigs may have abnormal consistency due to three causes: the presence of enteritis, which is an infection of the intestinal mucous membrane of the pig, a high body temperature, and unbalanced feed. Figure 1 shows the domain knowledge involved, represented in the formalism of Bayesian networks.

A Bayesian network is a model of a joint probability distribution over a set of stochastic variables [7]. It consists of a graphical structure and a set of associated probabilities. The nodes in the structure represent stochastic variables; in this paper, we assume all variables to be binary. The structure from Figure 1, for example, contains the variable C, representing the consistency of faeces in pigs; we use c to indicate abnormal faeces and $\bar{c}$ for normal faeces. The variable E models the presence or absence of enteritis, indicated by e and $\bar{e}$, respectively. The body temperature of a pig is modelled by the variable B; b and $\bar{b}$ represent a high and a normal body temperature, respectively. The composition of the pig's feed is represented by the variable F, where f indicates unbalanced feed and $\bar{f}$ denotes balanced feed. The arcs between the variables in the structure represent the presence of probabilistic influences between these variables. The arc $E \rightarrow C$ in Figure 1, for example, indicates that there is a direct influence between the two variables; E is referred to as the *parent* of C in the network. A combination of values for the parents of a variable is referred to as a *parent configuration*.

Associated with the graphical part of a network are numerical quantities from the modelled probability distribution: with each variable are specified the conditional probabilities of its values, given its possible parent con-

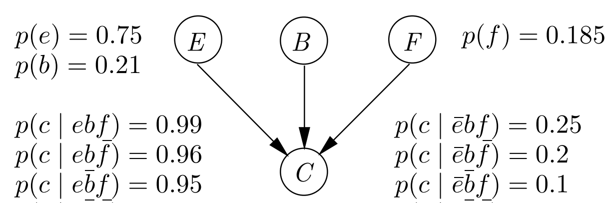

Figure 1: Part of a Bayesian network in the field of classical swine fever

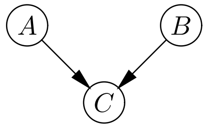

Figure 2: The variable C and its parents

figurations. For the variable C from our example, the probabilities $p(C \mid EBF)$ describe the joint effect of values for E, B and F, on the probabilities of C's values. More informally, they represent the probability of abnormal faeces in pigs, for all combinations of the three causes. From the probabilities in Figure 1 we learn, for example, that abnormal faeces are unlikely to result from another cause than the three causes mentioned above: the associated probability $p(c \mid \bar{e}\bar{b}\bar{f})$ equals 0.01. Furthermore, in the presence of enteritis, abnormal faeces is somewhat more likely to occur in pigs with a high body temperature that took balanced feed than in pigs with a normal body temperature that took unbalanced feed; the associated probabilities are 0.96 and 0.95, respectively. Our Bayesian network under construction currently contains 43 variables; the numerical part of the network includes some 2400 probabilities.

3. ELICITING PROBABILITY ORDERS

The first phase of our method is to elicit qualitative knowledge about probabilities from an expert in the domain of application; the elicited knowledge will be used in the second phase for constraining the probability estimates from a data collection. The knowledge to be acquired in the first phase concerns an order of the conditional probabilities required for the network under construction. We consider, as an example, the variable C and its parents A and B, as depicted in Figure 2. In the example, we denote the two values of A by a and $\bar{a}$, respectively; the values of the parents B and C are defined analogously. We now aim to elicit a (possibly partial) order of the conditional probabilities $p(c \mid AB)$; an example of such an order is $p(c \mid ab) \geq p(c \mid \bar{a}b) \geq p(c \mid a\bar{b}) \geq p(c \mid \bar{a}\bar{b})$.

For acquiring probability orders from an expert, we designed a dedicated elicitation technique, that is directed at minimising the cognitive effort required from the expert. Since reasoning in terms of frequencies is generally perceived as less demanding than reasoning in terms of probabilities [4], we use the frequency format for presenting the probabilities to be ordered to the expert. As an example, we consider again the structure from Figure 2. The probabilities $p(c \mid ab)$ and $p(c \mid \bar{a}\bar{b})$ equal

$$p(c \mid ab) = \frac{\text{number of cases in which } cab \text{ holds}}{\text{number of cases in which } ab \text{ holds}}$$

$$p(c \mid \bar{a}\bar{b}) = \frac{\text{number of cases in which } c\bar{a}\bar{b} \text{ holds}}{\text{number of cases in which } \bar{a}\bar{b} \text{ holds}}$$

A group of 100 pigs with:

Enteritis: **yes**
Body temperature: **normal**
Feed: **balanced**

Figure 3: A case card in the domain of classical swine fever

The conditional probability $p(c \mid ab)$ thus is the relative frequency of cases with c among the cases in which ab holds. Since the precise numbers are not relevant for the ordering task, we take the number of cases in which ab holds to be equal to the number of cases in which $\bar{a}\bar{b}$ holds. By doing so, we render the task of comparing the relative frequencies practicable for the expert: assigning an order to the probabilities $p(c \mid ab)$ and $p(c \mid \bar{a}\bar{b})$ now amounts to establishing which of the two groups of cases, in which ab and $\bar{a}\bar{b}$ holds, respectively, contains the largest number of cases in which also c holds. The ordering task is made more concrete for the expert by assigning the groups a fixed size, based on the characteristics of the domain.

To further reduce the cognitive effort that is asked of the expert in the ordering task, the various groups of cases are presented on *case cards*, for visual reference. For convenience of comparison, the information on each card is represented in an easily surveyable manner. An example case card in the domain of classical swine fever is shown in Figure 3. Now, for the structure from Figure 2, an order is to be obtained of the probabilities $p(c \mid ab)$, $p(c \mid a\bar{b})$, $p(c \mid \bar{a}b)$ and $p(c \mid \bar{a}\bar{b})$. A case card is constructed for each of the four groups of cases described above; the resulting four case cards are presented to the expert, who is asked to assign an order to the groups according to the number of cases in which c holds. This question is verbally posed, yet is also presented on a *question card*, again for visual reference.

Now, in performing the ordering task, the domain expert visualises (provisional) orders by arranging the case cards on a table. The visualisation enables the expert to maintain an overview of the order relations assigned so far. The possibility to shove the cards around allows for flexibility in changing the order relations. The case cards in addition facilitate an easy division of the ordering task in subtasks, for example for pairwise comparisons. Most experts find the method of pairwise comparison easy to use and tend to provide reliable information using the method [9]. We would like to note that the number of properties used for describing a group on a case card should be restricted, since the number of items that people are capable of processing simultane-

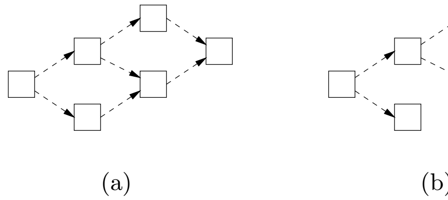

Figure 4: Example partial orders

ously, is limited [10]. We observe that the number of properties is equal to the number of parents of the variable under consideration. Now, in a real-life Bayesian network, this number is restricted anyhow, to guarantee a feasible running time of probabilistic inference.

An expert may not always be able to assign a total order to the probabilities under study. We therefore also allow for partial orders, thus creating the possibility for the expert to indicate that he does not know the order relations between two or more cards. The expert therefore is never forced to assign order relations that he feels uncomfortable with. Unlike total orders, partial orders cannot be visualised by simply arranging the cards as a sequence. Arranging the cards in a two-dimensional area, however, may introduce ambiguity about the exact order meant by the expert. We consider, for example, the arrangement from Figures 4(a) and (b), where the squares represent the cards. This arrangement can be interpreted in various ways, as illustrated by the arcs in the figures, where an arc from one card to another indicates that the frequency on the former card immediately precedes that on the latter card in the order. We therefore provide for visualising the order relations explicitly by cards depicting arrows; we feel that arrows are more easy to perceive than text, and more easy to interpret than mathematical symbols. Note that the *arrow cards* are needed only in case of partial orders.

With our elicitation method, we ask a domain expert to provide a probability order that is not intended to be strict. We feel that enforcing strict orders would render the task more demanding on the expert and, hence, would result in additional “don’t know’s”, or partial orders. If the expert is convinced of the equality of two frequencies, this relation can easily be made explicit by piling the associated case cards.

We would like to note that the idea of using case cards that describe frequencies rather than probabilities, has been introduced before [6]. The two elicitation techniques differ considerably, however, since they are aimed at obtaining quite different types of (qualitative) knowledge. We note, for example, that there exist total orders that cannot be elicited using the technique described before [6]; these orders were not relevant for the goal for which that particular technique was designed. All possible orders can be acquired using the technique described in the present paper, however. We further note that the former technique built on the assumption that a domain expert is always able to indicate a total order. In contrast, our present technique allows for partial orders and for explicit equalities.

4. LEARNING PROBABILITIES WITH ORDER CONSTRAINTS

In the second phase of our method, the order relations specified by the domain expert are used as constraints on the probability estimates that are extracted from the available data collection. The general idea is to first estimate the probabilities from the data without using the constraints. If the estimates happen to satisfy the order relations, then no adjustments are made. Otherwise, the estimates are adjusted so as to satisfy the elicited relations. We argue that, since the adjusted estimates satisfy the expert’s qualitative knowledge, the resulting network is less likely to exhibit counterintuitive reasoning behaviour. Furthermore, empirical evidence indicates that the adjusted estimates give a better fit of the true (unknown) probability distribution, if the elicited order is consistent with this distribution [3].

We consider again the graphical structure from Figure 2, representing the variable C and its parents A and B. For this structure, we need to obtain the conditional probabilities $p(c \mid ab)$, $p(c \mid \bar{a}b)$, $p(c \mid a\bar{b})$, and $p(c \mid \bar{a}\bar{b})$. The *unrestricted maximum-likelihood estimates* of these probabilities are learned from the data collection; henceforth, we refer to these unrestricted estimates as the *basic estimates*. The basic estimate for the probability $p(c \mid ab)$, for example, is given by

$$\hat{p}(c \mid ab) = \frac{\text{number of records in which } cab \text{ holds}}{\text{number of records in which } ab \text{ holds}}$$

The basic estimates for the remaining three probabilities are defined analogously. The basic estimate of the probability of c given a specific parent configuration thus is the relative frequency of records with c among the records with the given parent configuration. Note that the basic estimates are indeed estimates for the conditional probabilities stated in Section 3.

The basic estimates obtained from the data collection may not satisfy the order that was elicited from the expert. Especially for smaller data collections, violations of the elicited relations are likely to occur. We now adjust the basic estimates to meet the expert’s intuition, in a mathematically justified manner. For this purpose, we use an algorithm by Brunk [1] (see also [12], pp. 24-25), called the minimum lower sets (MLS) algorithm. This algorithm produces maximum-likelihood estimates from the data, subject to the constraint that these must satisfy the expert’s order relations. In other words, of all estimates that satisfy these order relations, the ones

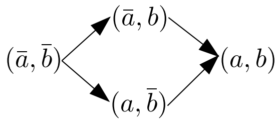

Figure 5: Order on parent configurations

produced by the algorithm give the highest probability of the observed data. We henceforth refer to these estimates as the *order-constrained estimates*.

We illustrate the MLS algorithm by means of our example from Figure 2. Suppose that the expert has specified the order given in Figure 5. The tuples in the figure indicate the four parent configurations for the variable C. An arc from one configuration to another indicates that, within the order, the probability of c given the former configuration immediately precedes that given the latter configuration. The constraints on the probabilities of c given the various parent configurations that are derived from this order thus are

(1) $p(c \mid \bar{a}\bar{b}) \leq p(c \mid \bar{a}b)$
(2) $p(c \mid \bar{a}\bar{b}) \leq p(c \mid a\bar{b})$
(3) $p(c \mid \bar{a}b) \leq p(c \mid ab)$
(4) $p(c \mid a\bar{b}) \leq p(c \mid ab)$

Now suppose that in addition we have a data collection containing 54 observations on the variables C, A and B. In the left part of Table 1, containing a cell for each parent configuration, each cell specifies the relative frequency of observations in the data collection with that parent configuration, that also have c. Note that the basic estimates obtained from the data are equal to these relative frequencies. We observe that the basic estimates violate the order constraints (1) and (2), since $\hat{p}(c \mid \bar{a}\bar{b}) = 2/4 = 0.5 > \hat{p}(c \mid \bar{a}b) = 4/16 = 0.25$, and $\hat{p}(c \mid \bar{a}\bar{b}) = 0.5 > \hat{p}(c \mid a\bar{b}) \approx 0.33$.

The MLS algorithm now resolves the identified violations among the basic estimates by taking the weighted average of the estimates over as few conflicting cells as possible, and subsequently assigning this weighted average as the order-constrained estimate to each of those cells. Note that setting two conflicting basic estimates to the same value always serves to resolve the conflict, since the inequalities to be satisfied are not strict. In our example, the violations of the order constraints (1)

Table 1: Basic (left) and order-constrained (right) estimates for the first example

	$\bar{b}$	b
$\bar{a}$	2/4	4/16
a	8/24	8/10

	$\bar{b}$	b
$\bar{a}$	6/20	6/20
a	8/24	8/10

Table 2: Basic (left) and order-constrained (right) estimates for the second example

	$\bar{b}$	b
$\bar{a}$	2/4	4/16
a	6/24	8/10

	$\bar{b}$	b
$\bar{a}$	12/44	12/44
a	12/44	8/10

and (2) are resolved by averaging the basic estimates $\hat{p}(c \mid \bar{a}\bar{b})$ and $\hat{p}(c \mid \bar{a}b)$. Note that the weighted average is computed by adding the numerators and denominators, respectively, of the relative frequencies in the cells involved. We then obtain $p^*(c \mid \bar{a}\bar{b}) = p^*(c \mid \bar{a}b) = 6/20 = 0.3$, where p^* is used to denote the order-constrained estimates produced by the algorithm. Note that in computing the average, the basic estimate $\hat{p}(c \mid \bar{a}b)$ is assigned a higher weight than $\hat{p}(c \mid \bar{a}\bar{b})$, because it is based on more observations and therefore has a higher precision. Most likely the identified violations are due to an overestimate of $p(c \mid \bar{a}\bar{b})$ based on relatively few observations. By the adjustment of the estimates $\hat{p}(c \mid \bar{a}\bar{b})$ and $\hat{p}(c \mid \bar{a}b)$, all order constraints have been satisfied. No further adjustments are required, therefore. The algorithm thus sets $p^*(c \mid a\bar{b}) = \hat{p}(c \mid a\bar{b}) = 0.33$ and $p^*(c \mid ab) = \hat{p}(c \mid ab) = 0.8$. The resulting order-constrained estimates are summarised in the right part of Table 1.

As a second example, we consider the basic estimates in the left part of Table 2. The basic estimates again violate the order constraints (1) and (2), but this time they cannot be resolved simultaneously by averaging $\hat{p}$ over just one of the sets of cells $\{(\bar{a}, \bar{b}), (\bar{a}, b)\}$ and $\{(\bar{a}, \bar{b}), (a, \bar{b})\}$. The violations now are resolved by taking the weighted average of the basic estimates from the cells $(\bar{a}, \bar{b})$, $(\bar{a}, b)$, and $(a, \bar{b})$; the algorithm sets $p^*(c \mid \bar{a}\bar{b}) = p^*(c \mid \bar{a}b) = p^*(c \mid a\bar{b}) = 12/44 \approx 0.27$, and $p^*(c \mid ab) = \hat{p}(c \mid ab) = 0.8$. For a more detailed description of the MLS algorithm and its application to learning probabilities in Bayesian networks, we would like to refer to [3].

5. EXPERIENCES WITH OUR METHOD

We have applied our method for quantifying various parts of our Bayesian network in the domain of classical swine fever. We first elicited the probability orders for various parts of the network. The elicited order relations were then used to constrain the probabilities that were extracted from a small data collection.

5.1 Eliciting probability orders

In a single knowledge-acquisition session, we elicited the probability orders for various parts of our network. Two of the domain experts involved in building the network were present; the orders were elicited from one of them. Among the knowledge engineers, one conducted the interview; the other two were observers.

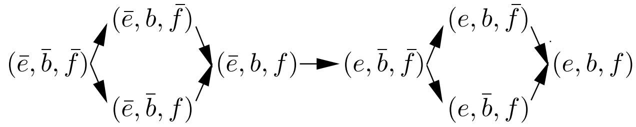

Figure 6: Order on parent configurations in the domain of classical swine fever

The expert was asked to provide probability orders for three parts of the network, consisting of a variable and two or three parents. The ordering tasks were preceded by a verbal instruction, that illustrated the basic idea of what was expected of the expert by an example. The example was taken from a domain other than classical swine fever, to prevent any bias. For the same reason, the purpose of the task was not explained to the domain expert beforehand.

We now focus on the elicitation of a probability order for the graphical structure depicted in Figure 1; we recall that this part of the network models the consistency of a pig's faeces being influenced by the presence or absence of enteritis, the pig's body temperature, and the composition of its feed. The probabilities in the figure were not shown to the domain expert during the elicitation; we will comment on this issue shortly. For this part of the network, eight probabilities had to be ordered; one of the cards that we prepared for this task, was shown in Figure 3. The cards were presented to the expert as a set instead of one-by-one, to prevent any suggestion of a pre-existing order. The expert was asked to order the groups on the case cards according to their number of pigs with abnormal faeces. The expert was further handed a question card; apart from the ordering question, the card contained reminders of what to do in case of equal numbers and "don't know's". Arrow cards of various sizes were also made available to the expert. Figure 6 shows the order that was provided by the expert, again in terms of parent configurations.

During the session, the expert ordered sets of four to twelve case cards. He performed each of the ordering tasks in 40 seconds up to about 7.5 minutes. The order in Figure 6 was assigned in four minutes, for example. Ordering the largest set, containing twelve case cards, took 7.5 minutes. In ordering these twelve cards, he explicitly divided the overall task in smaller subtasks. First he organised the cards into three columns, according to the three possible values of one of the properties described on the cards. Next, the cards from each column separately were ordered, after which they were merged into an overall order. Although for several parts of the network a partial order was elicited, the domain expert decided not to use the arrow cards to visualise the order relations. The arrangements of the cards on the table appeared to be sufficient.

At the end of the interview, the domain expert indicated that he had perceived the ordering tasks as "neutral", that is, as not particularly difficult or easy. The cards were considered to serve well as a visual reference during the ordering tasks. We would like to note that our experiment pertained to parts of the network for which the same expert had provided numerical probability assessments. The expert could therefore possibly have been biased in assigning the order relations. We would like to note, however, that the probabilities had been assessed six to eight months prior to the current interview, and that the probabilities had not been made available to the expert during the ordering tasks. Moreover, at the end of the interview he indicated that he had found the tasks of assessing the probabilities and assigning the orders incomparable. In addition, he remarked that he had ordered the cards without having exact numbers in mind. We conclude that our elicitation technique enabled him to focus on qualitative rather than quantitative knowledge about probabilities. We observe furthermore that the probabilities from Figure 1, which were assessed by the expert, are consistent with the elicited order from Figure 6; this property does not hold for all elicited orders, however.

5.2 Learning probabilities with constraints

The order that we elicited for the graphical structure from Figure 1 is now used for constraining the probability estimates obtained from data. We generated a small data sample of $n = 50$ observations from the network fragment, to illustrate the drawbacks of the basic probability estimates, and to show how these are forestalled by the order-constrained estimates. For generating the data, we used the probability assessments that the expert had provided before. We recall that the order relations specified by the expert are consistent with these assessments. The data sample that was generated using these probabilities, however, is small and may well produce basic estimates that are not consistent with the elicited relations.

The generated data is summarised in Table 3. The table displays the number of observations $n(EBF)$ for

Table 3: Data for the experiment

	E	B	F	$n(EBF)$	$n(c, EBF)$
1	$\bar{e}$	$\bar{b}$	$\bar{f}$	8	0
2	e	$\bar{b}$	$\bar{f}$	21	20
3	$\bar{e}$	b	$\bar{f}$	2	0
4	e	b	$\bar{f}$	7	7
5	$\bar{e}$	$\bar{b}$	f	2	1
6	e	$\bar{b}$	f	8	7
7	$\bar{e}$	b	f	1	0
8	e	b	f	1	1

Table 4: Basic estimates of $p(c|EBF)$

$\bar{f}$	$\bar{b}$	b	f	$\bar{b}$	b
$\bar{e}$	0	0	$\bar{e}$	0.5	0
e	0.95	1	e	0.88	1

each parent configuration, and the number $n(c, EBF)$ among these with c. Table 4 gives the basic estimates of $p(c \mid EBF)$ that were extracted from the data, for each parent configuration. From the table, we observe that these estimates violate the order constraints that are derived from the elicited order, depicted in Figure 6. The basic estimates, for example, imply that the probability of abnormal faeces for pigs with enteritis (e) and a normal body temperature ($\bar{b}$) is 95% if the pig takes balanced feed ($\bar{f}$) but drops to 88% if the pig takes unbalanced feed (f).

The violations

(1) $\hat{p}(c \mid \bar{e}\bar{b}f) > \hat{p}(c \mid \bar{e}bf)$
(2) $\hat{p}(c \mid e\bar{b}\bar{f}) > \hat{p}(c \mid e\bar{b}f)$

now are resolved by averaging the basic estimates over the conflicting cells. The algorithm sets $p^*(c \mid \bar{e}\bar{b}f) = p^*(c \mid \bar{e}bf) = 1/3 \approx 0.33$, and $p^*(c \mid e\bar{b}\bar{f}) = p^*(c \mid e\bar{b}f) = 27/29 \approx 0.93$. Note that the order-constrained estimate 0.93 is closer to the basic estimate of $p(c \mid e\bar{b}\bar{f})$ than to the basic estimate of $p(c \mid e\bar{b}f)$. Since we have only eight observations in the cell $(e, \bar{b}, f)$ and 21 observations in the cell $(e, \bar{b}, \bar{f})$, the violation is more likely due to an underestimate of the former than to an overestimate of the latter probability. The order-constrained estimates are given in Table 5.

The violations by the basic estimates of the order relations provided by the domain expert, have now been resolved. The order-constrained estimates now imply that the probability of abnormal faeces for pigs with enteritis and a normal body temperature is 93% if the pig takes balanced feed and is also 93% if the pig takes unbalanced feed. These numbers are an improvement over the basic estimates, but still may not be entirely satisfactory. The equality of the two order-constrained estimates is nevertheless a consequence of the qualitative knowledge that was specified by the expert. If we would want to enforce a strict increase of the probability of abnormal faeces when the feed changes from balanced to unbalanced, the expert would have to indicate some

Table 5: Order-constrained estimates of $p(c|EBF)$

$\bar{f}$	$\bar{b}$	b	f	$\bar{b}$	b
$\bar{e}$	0	0	$\bar{e}$	0.33	0.33
e	0.93	1	e	0.93	1

Table 6: Order-constrained estimates of $p(c|EBF)$ with minimum differences

$\bar{f}$	$\bar{b}$	b	f	$\bar{b}$	b
$\bar{e}$	0	0	$\bar{e}$	0.31	0.46
e	0.90	1	e	0.95	1

minimum numeric difference between the two probabilities. As an illustration, we suppose that the expert has specified the following minimum differences:

(1) $p(c \mid \bar{e}bf) - p(c \mid \bar{e}\bar{b}f) \geq 0.15$
(2) $p(c \mid e\bar{b}f) - p(c \mid e\bar{b}\bar{f}) \geq 0.05$

The first constraint, for example, states that for pigs that do not have enteritis and that take unbalanced feed, the percentage of pigs with abnormal faeces is at least 15 percentage points higher for pigs with a high body temperature than it is for pigs with a normal body temperature. The order-constrained estimates given this additional knowledge are given in Table 6. Note that the differences between the estimates for the pairs of violating cells $\{(\bar{e}, \bar{b}, f), (\bar{e}, b, f)\}$ and $\{(e, \bar{b}, \bar{f}), (e, \bar{b}, f)\}$ are exactly equal to the specified minimum differences.

The disadvantage of the approach above is that the expert is required to specify numeric information in addition to the order relations. This disadvantage can be alleviated to at least some extent by eliciting the additional information only if necessary, that is, when averaging leads to unwanted equalities between probability estimates. Another option is to enforce some standard minimum difference between two order-constrained estimates whenever the order indicates a direct relation between the two probabilities involved. It is questionable however, whether such a procedure would yield results that are acceptable to the expert. Another disadvantage of the approach above is that the MLS algorithm has not been designed to handle numeric differences and therefore can not longer be applied to compute the constrained estimates. A numerical optimisation algorithm could be used instead; such an algorithm is much more demanding from a computational point of view, however. We feel that the issue of enforcing minimum numeric differences deserves further research.

6. CONCLUSIONS

In this paper, we proposed a new method for quantifying a Bayesian network. Our method exploits probability orders elicited from domain experts to constrain the probability estimates that are extracted from a data collection. To support the acquisition of the probability orders, we designed a dedicated elicitation technique. We applied our method and the associated elicitation technique for quantifying parts of a Bayesian network in the field of classical swine fever.

From our experiences, we conclude that our elicitation technique enabled the domain expert to focus on qualitative rather than quantitative knowledge about the required probabilities. Moreover, constructing the requested orders proved to be practicable for our expert. The use of our method further showed that counterintuitive reasoning behaviour of the resulting network is likely to be prevented by using the elicited order relations as constraints on the probabilities estimated from a small data collection. We expect that in domains where only a small data collection is available, our method facilitates a knowledge engineer to build a Bayesian network that has a higher acceptance among domain experts, while requiring less cognitive effort from the experts during elicitation.

7. ACKNOWLEDGEMENTS

The research presented in this paper was partly supported by the Netherlands Organisation for Scientific Research (NWO).

8. REFERENCES

[1] Brunk H.D. Maximum likelihood estimates of monotone parameters. *Annals of Mathematical Statistics*, 1955, 26:607–616.

[2] Druzdzel M.J., van der Gaag L.C. Elicitation of probabilities for belief networks: combining qualitative and quantitative information. In: Besnard P., Hanks S. (Eds.). *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers, San Francisco, California, 1995, 141–148.

[3] Feelders A., van der Gaag L.C. Learning Bayesian network parameters under order constraints. In: Lucas P. (Ed.). *Proceedings of the Second European Workshop on Probabilistic Graphical Models (PGM'04)*. Leiden, 2004, 73–80.

[4] Gigerenzer G., Hoffrage U. How to improve Bayesian reasoning without instruction: frequency formats. *Psychological Review*, 1995, 102:684–704.

[5] Helsper E.M., van der Gaag L.C. Building Bayesian networks through ontologies. In: van Harmelen F. (Ed.). *Proceedings of the 15th European Conference on Artificial Intelligence*. IOS Press, Amsterdam, 2002, 680–684.

[6] Helsper E.M., van der Gaag L.C., Groenendaal F. Designing a procedure for the acquisition of probability constraints for Bayesian networks. In: Motta E., Shadbolt N.R., Stutt A., Gibbins N. (Eds.). *Engineering Knowledge in the Age of the Semantic Web: 14th International Conference, EKAW 2004*. Springer-Verlag, Heidelberg, 2004, 280–292.

[7] Jensen F.V. *Bayesian Networks and Decision Graphs*. Springer-Verlag, New York, 2001.

[8] Kahneman D., Slovic P., Tversky A. *Judgment under Uncertainty: Heuristics and Biases*. Cambridge University Press, Cambridge, 1982.

[9] Meyer M.A., Booker J.M. *Eliciting and Analyzing Expert Judgment: A Practical Guide*. ASA-SIAM Series on Statistics and Applied Probability, 2001.

[10] Miller G.A. The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review*, 1956, 63:81–97.

[11] Morgan M.G., Henrion M. *Uncertainty, a Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. Cambridge University Press, Cambridge, 1990.

[12] Robertson T., Wright F., Dykstra R.L. *Order Restricted Statistical Inference*. Wiley, New York, 1988.