Designing a Procedure for the Acquisition of Probability Constraints for Bayesian Networks

Eveline M. Helsper[1], Linda C. van der Gaag[1], and Floris Groenendaal[2]

[1] Institute of Information and Computing Sciences, Utrecht University,
P.O. Box 80.089, 3508 TB Utrecht, The Netherlands
{eveline,linda}@cs.uu.nl
[2] Department of Neonatology, University Medical Center Utrecht,
P.O. Box 85.090, 3508 AB Utrecht, The Netherlands
F.Groenendaal@azu.nl

Abstract. Among the various tasks involved in building a Bayesian network for a real-life application, the task of eliciting all probabilities required is generally considered the most daunting. We propose to simplify this task by first acquiring qualitative features of the probability distribution to be represented; these features can subsequently be taken as constraints on the precise probabilities to be obtained. We discuss the design of a procedure that guides the knowledge engineer in acquiring these qualitative features in an efficient way, based on an in-depth analysis of all viable combinations of features. In addition, we report on initial experiences with our procedure in the domain of neonatology.

1 Introduction

Bayesian networks are well established in artificial-intelligence research as intuitively appealing representations of knowledge, tailored to domains in which uncertainty is predominant [1]. 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 of the network encodes the variables of importance in the domain being represented, along with their probabilistic interrelationships. The strengths of the relationships between the variables are quantified by conditional probability distributions. These distributions constitute the numerical part of the network.

Bayesian networks for real-life applications are often constructed with the help of a domain expert. Experience shows that, although it may require considerable effort, configuring the graphical part of the network is quite practicable. In fact, building the graphical part has parallels to designing a domain model for any knowledge-based system; well-known knowledge engineering techniques can therefore, to at least some extent, be employed for this purpose [2]. Obtaining all probabilities required for the numerical part of the network is generally considered a far harder task, however, especially if these probabilities have to be assessed by the domain expert [3]. The more interactions there are among the represented variables, moreover, the harder the task is.

Given a well-constructed graphical part, we propose to simplify the task of obtaining all probabilities required for a Bayesian network in the making by first acquiring the probabilistic interactions among the represented variables in qualitative terms. The acquired qualitative features of the distribution being captured can subsequently be taken as constraints on the precise probabilities to be obtained [4, 5]. For the acquisition of the qualitative features, we designed a procedure based upon the concepts of qualitative influence and qualitative synergy [6]. Our procedure more specifically derives from an analysis of the fundamental properties of these concepts.

We conducted an in-depth study of the qualitative features that may hold for the conditional probability distributions for a variable and its possible causes. Our study revealed that the viable combinations of features constitute four classes [7]. We exploited these classes to design an efficient procedure for acquiring combinations of qualitative features for a Bayesian network in the making from a domain expert. Our procedure begins with the elicitation of knowledge to establish, for each variable and its causes, the class of the features that hold for its conditional probability distributions. As within each class various features are fixed, it serves to indicate the additional knowledge that needs to be acquired to fully specify the combination of features under study. By building upon the four classes, therefore, the procedure guides the knowledge engineer, step by step, in focusing further acquisition efforts.

We conducted an initial study of the use of our procedure with an intensive-care neonatologist to acquire probabilistic information for a real-life Bayesian network in the making. The results indicate that our procedure requires relatively little effort on the part of the expert. From our initial experiences, in fact, we feel that deriving an acquisition procedure from an analysis of the fundamental properties of the concepts used, has resulted in an efficient, dedicated procedure for one of the harder tasks in building Bayesian networks.

The paper is organised as follows. Sect. 2 briefly reviews Bayesian networks and the concepts of qualitative influence and qualitative synergy. Sect. 3 introduces the four classes of combinations of qualitative features. In Sect. 4, we present our procedure, and associated techniques, for acquiring qualitative features of conditional probability distributions from a domain expert. We report on an initial study of the use of our procedure in Sect. 5. The paper ends with our concluding observations in Sect. 6.

2 Bayesian networks

A *Bayesian network* is a model of a joint probability distribution over a set of statistical variables [1]. The model includes a graphical structure in which each node represents a variable. For ease of exposition, we assume all statistical variables to be binary, taking one of the values *true* and *false*; we further assume that the two values are ordered by *true* $>$ *false*. For abbreviation, we use a to denote that the variable A adopts the value *true* and $\bar{a}$ to denote $A = \mathit{false}$. The arcs in the graphical structure represent the presence of probabilistic influences

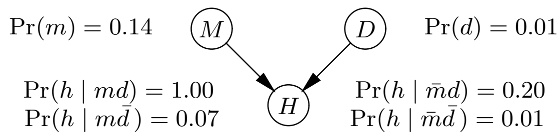

Fig. 1. A fragment of a Bayesian network in neonatology

between the variables. An arc $B \rightarrow C$ between the two variables B and C indicates that there is a direct influence between them; B then is generally referred to as the *cause* of the *effect* C. The variable C with all its possible causes is termed a *causal mechanism*. We note that the term causal mechanism is used to denote *any* variable and its parents in the graphical structure, also if the influences involved are not strictly causal. Associated with the graphical part of the Bayesian network are numerical quantities from the modelled probability distribution: with each variable C are specified *conditional probability distributions* $\Pr(C \mid \pi(C))$, that describe the joint effect of values for the causes $\pi(C)$ of C, on the probabilities of C's values.

Fig. 1 depicts a fragment of a Bayesian network in the domain of neonatology. The fragment pertains to the possible causes of a hypoplastic, or underdeveloped, lung in premature newborns. The presence or absence of a hypoplastic lung is modelled by the variable H. One of its possible causes is a spontaneous, prolonged rupture of the foetal membranes and the subsequent lack of amniotic fluid; the variable M captures whether or not there has been such a rupture. The other possible cause of a hypoplastic lung is a diaphragmatic hernia. Such a hernia is a defect in the diaphragm through which the abdominal contents herniate and thereby reduce the space for the lung to develop; the presence or absence of the defect is modelled by the variable D. The conditional probabilities specified for the variable H indicate the likelihood of a hypoplastic lung occurring in a premature newborn. These probabilities show, for example, that a hypoplastic lung is very unlikely to result from another cause than the two mentioned above: the associated probability is just 0.01. The conditional probabilities further reveal that the presence of a diaphragmatic hernia is more likely to result in a hypoplastic lung in a newborn than a prolonged rupture of the foetal membranes; the associated probabilities equal 0.20 and 0.07, respectively. While the presence of either one of the two causes is not very likely to result in a hypoplastic lung, the presence of both causes is certain to give rise to the condition: the influences of the two causes on the likelihood of a hypoplastic lung serve to strengthen one another. We note that, while the graphical structure of a Bayesian network models the presence of probabilistic influences between the variables, the interactions among the influences are modelled in the associated probabilities only.

To provide for capturing probabilistic influences and the interactions among them in a qualitative way, we build upon the concepts of qualitative influence and synergy [6]. A *qualitative influence* between two statistical variables expresses how observing a value for the one variable affects the probability distribution for the other variable. For example, a *positive qualitative influence* of a variable B on a variable C expresses that observing the higher value for B makes the higher value for C more likely, regardless of any other influences on C, that is,

$$\Pr(c \mid bx) \geq \Pr(c \mid \bar{b}x)$$

for any combination of values x for the other causes of C than B. A *negative influence* and a *zero influence* are defined analogously, replacing the inequality $\geq$ in the above formula by $\leq$ and $=$, respectively. If the influence of B on C is positive for one combination of values x and negative for another combination, then it is said to be *ambiguous*. From the probabilities specified for the variable H in the network from Fig. 1, for example, we observe that $\Pr(h \mid md) \geq \Pr(h \mid m\bar{d})$ and $\Pr(h \mid \bar{m}d) \geq \Pr(h \mid \bar{m}\bar{d})$. We conclude that D exerts a positive qualitative influence on H, that is, the presence of a diaphragmatic hernia makes the occurrence of a hypoplastic lung more likely.

An *additive synergy* expresses how the values of two variables combine to yield a joint effect on the third variable. A *positive additive synergy* of the variables A and B on their common effect C, for example, expresses that the joint influence of A and B on C is greater than the sum of their separate influences, regardless of any other influences on C, that is,

$$\Pr(c \mid abx) + \Pr(c \mid \bar{a}\bar{b}x) \geq \Pr(c \mid a\bar{b}x) + \Pr(c \mid \bar{a}bx)$$

for any combination of values x for the other causes of C than A and B. Negative, zero and ambiguous additive synergies are defined analogously. From the probabilities specified in Fig. 1, for example, we have that $\Pr(h \mid md) + \Pr(h \mid \bar{m}\bar{d}) = 1.01$ and $\Pr(h \mid m\bar{d}) + \Pr(h \mid \bar{m}d) = 0.27$. We conclude that M and D show a positive additive synergy on D, that is, the joint influence of a prolonged rupture of the foetal membranes and a diaphragmatic hernia on the occurrence of a hypoplastic lung is larger than the sum of their separate influences.

Product synergies express how observing a value for a variable affects the probability distribution for another variable in view of a value for a third variable [8]. A *negative product synergy* of a variable A on a variable B (and vice versa) given the value c for their common effect C, for example, expresses that, given c, observing the higher value for A renders the higher value for B less likely; formally, the negative product synergy given c is defined as

$$\Pr(c \mid abx) \cdot \Pr(c \mid \bar{a}\bar{b}x) \leq \Pr(c \mid a\bar{b}x) \cdot \Pr(c \mid \bar{a}bx)$$

for any combination of values x for the other causes of C than A and B. *Positive*, *zero* and *ambiguous product synergies* again are defined analogously. From the probabilities specified for the variable H in Fig. 1, for example, we have that $\Pr(h \mid md) \cdot \Pr(h \mid \bar{m}\bar{d}) = 0.01$ and $\Pr(h \mid m\bar{d}) \cdot \Pr(h \mid \bar{m}d) = 0.014$, from which we conclude that M and D show a negative product synergy given h. A product synergy of the variable A on the variable B in essence describes the influence that is induced between them by the observation of a value for their common effect. From a negative product synergy of A on B given c, we then have that

$$\Pr(b \mid acy) \leq \Pr(b \mid \bar{a}cy)$$

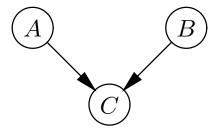

Fig. 2. A basic causal mechanism

for any combination of values y for all causes of B and C except A and B itself. From the negative product synergy that we found from Fig. 1, we thus have that, given the occurrence of a hypoplastic lung, the absence of a prolonged rupture of the foetal membranes makes the presence of a diaphragmatic hernia more likely.

3 Classes of qualitative features

In a Bayesian network, each causal mechanism has associated a *pattern of interaction* composed of qualitative influences and synergies. The pattern has a qualitative influence of each cause on the common effect. It further has an additive synergy for each pair of causes. For each such pair, moreover, it specifies two opposite product synergies, one for each value of the common effect. Not all combinations of qualitative features constitute viable patterns of interaction, however. To identify the viable patterns, we investigated the mechanism from Fig. 2 and studied the conditional probabilities for the presence of the effect C in detail [7]. Each ordering of these four probabilities gives rise to a specific pattern of interaction. By studying all possible orderings, we identified four classes of patterns, each resulting in a different type of combination of influences.

Class I includes all orderings of the four conditional probabilities under study in which the common effect is less likely to occur in the cases in which just one of the two causes is present than in the cases where both causes are either present or absent, or vice versa. These orderings give rise to an interaction pattern that specifies ambiguous influences of each cause separately; the additive synergy and the product synergy given the common effect both are positive, or negative alternatively. As an example, we consider

$$\Pr(c \mid ab) \geq \Pr(c \mid \bar{a}\bar{b}) \geq \Pr(c \mid \bar{a}b) \geq \Pr(c \mid a\bar{b})$$

We have that the influence of the cause A on the effect C is positive in the presence of B and negative in the absence of B. We conclude that the overall influence of A on C is ambiguous. A similar observation holds for the influence of B on C. We further find that the two causes exhibit a positive additive synergy on their common effect as well as a positive product synergy given c.

Class II includes all orderings of the four conditional probabilities under study in which the common effect is more likely to occur in the cases in which just one of the causes is present than in the case in which both causes are absent yet less likely to occur than in the case in which both causes are present, or vice versa. These orderings give rise to an interaction pattern that specifies positive influences of both causes on the common effect, or negative influences alternatively. The additive and product synergies are dependent upon the precise numbers for the four probabilities. As an example, we consider

$$\Pr(c \mid ab) \geq \Pr(c \mid \bar{a}b) \geq \Pr(c \mid a\bar{b}) \geq \Pr(c \mid \bar{a}\bar{b})$$

From the ordering, we find that both influences on C are positive.

Class *III* includes all orderings of the four conditional probabilities under study in which the common effect is more likely to occur in the two cases in which both causes are either present or absent than in the case in which just the one cause is present yet less likely than in the case in which just the other cause is present, or vice versa. These orderings give rise to an interaction pattern that specifies opposite non-ambiguous influences of the two causes separately. The orderings may again give rise to different additive synergies and different product synergies given the common effect. As an example, we consider

$$\Pr(c \mid \bar{a}b) \geq \Pr(c \mid ab) \geq \Pr(c \mid \bar{a}\bar{b}) \geq \Pr(c \mid a\bar{b})$$

From the ordering, we find that the variable A exerts a negative influence on C while the variable B exerts a positive influence on C.

Class *IV*, to conclude, comprises all orderings that give rise to an interaction pattern that specifies an ambiguous influence of one of the causes and a non-ambiguous influence of the other cause. The additive synergy and the product synergy given the common effect are both positive, or both negative alternatively. An example ordering of the four probabilities from this class is

$$\Pr(c \mid ab) \geq \Pr(c \mid \bar{a}b) \geq \Pr(c \mid \bar{a}\bar{b}) \geq \Pr(c \mid a\bar{b})$$

which gives rise to an ambiguous influence of A and a positive influence of B on the common effect C.

The four classes of orderings are defined for causal mechanisms composed of an effect and two possible causes. A mechanism that comprises three possible causes in essence is composed of three partial mechanisms, each consisting of the effect and two of its causes. The patterns of interaction for these partial mechanisms then combine into the overall pattern.

4 The acquisition of interaction patterns

Obtaining all probabilities required for a Bayesian network is a demanding task. To simplify the task, we propose to first acquire qualitative patterns of interaction for the network's causal mechanisms. These patterns can subsequently be used as constraints on the precise probabilities to be obtained. For the acquisition of interaction patterns from a domain expert, we designed a dedicated procedure by building upon the classes of probability orderings that we identified before. As these classes include viable patterns of interaction only, we are guaranteed not to obtain any impossible ones. By not using the concepts of qualitative influence and synergy for the elicitation directly, moreover, we circumvent any misinterpretation by the domain expert of these concepts. Furthermore, intermediate elicitation results serve to restrict the possible patterns for a mechanism and, hence, the additional knowledge that needs to be acquired to fully specify the pattern under study. The procedure thus minimises the amount of knowledge to be elicited from the expert, and guides the knowledge engineer in the acquisition process. In this section, we present our basic acquisition procedure and propose elicitation techniques to be used with it.

4.1 A procedure for acquiring interaction patterns

Our basic procedure for acquiring interaction patterns for mechanisms involving two causes, is shown schematically in Fig. 3; the steps in which knowledge is elicited from a domain expert are distinguished from the various other steps performed by the knowledge engineer by a bold outline. Before discussing the procedure in some detail, we would like to note that it uses two clusters of classes to guide the acquisition. We observe that, for the two classes *II* and *III*, the qualitative influences involved follow directly from the largest and smallest probabilities of the effect being present. For the classes *I* and *IV*, a similar observation holds for the additive and product synergies. The procedure therefore distinguishes between the two clusters of classes $\{I, IV\}$ and $\{II, III\}$.

The first phase of our acquisition procedure serves to identify the cluster of classes that applies to the causal mechanism under study. To this end, knowledge is elicited about the cases in which the presence of the effect is the most likely and the least likely to occur, respectively. Based upon the obtained partial ordering of the probabilities involved, the knowledge engineer associates the appropriate cluster of classes with the mechanism. Now, if the associated cluster equals $\{I, IV\}$, then the additive and product synergies are uniquely defined and, hence, readily looked up by the knowledge engineer; no further elicitation efforts are required for this purpose. For example, if for the mechanism of Fig. 2, the probabilities $\Pr(c \mid ab)$ and $\Pr(c \mid a\bar{b})$ are indicated by the expert as being the largest and the smallest, respectively, among the four conditional probabilities involved, then the knowledge engineer associates the cluster of classes $\{I, IV\}$ with the mechanism. From the characterisations of the two classes and the order of the selected probabilities, it now follows directly that both the additive synergy and the product synergy given c, are positive. For establishing the qualitative influences, however, a total ordering of the probabilities is required. In the next phase of our procedure, therefore, such an ordering is elicited from the domain expert. Based upon the ordering obtained, the knowledge engineer associates the appropriate class with the mechanism under study and decides upon the qualitative influences. For example, if the expert gives the ordering $\Pr(c \mid ab) \geq \Pr(c \mid \bar{a}\bar{b}) \geq \Pr(c \mid \bar{a}b) \geq \Pr(c \mid a\bar{b})$, then class *I* is associated with the mechanism. Since all orderings from class *I* give rise to ambiguous influences, the knowledge engineer now knows that the qualitative influences of the mechanism under study are ambiguous. Similar observations apply if class *IV* would have been associated with the mechanism.

In the first phase of our acquisition procedure, the cluster of classes $\{II, III\}$ is readily distinguished from the cluster $\{I, IV\}$ by the knowledge engineer

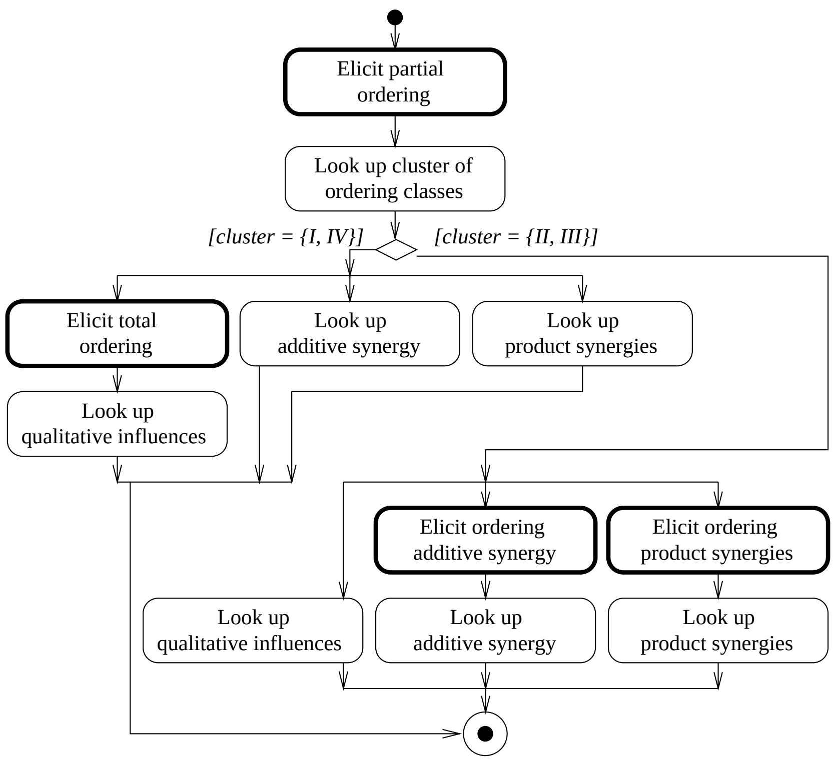

Fig. 3. The basic procedure for acquiring interaction patterns for causal mechanisms

through the largest and smallest probabilities indicated by the domain expert. The partial ordering obtained, moreover, allows the knowledge engineer to distinguish between the two classes *II* and *III*. For both classes, the qualitative influences of the mechanism under study are uniquely defined. In fact, there is no need to elicit a total ordering of the probabilities involved. The second phase of our procedure therefore focuses on the elicitation of additional knowledge that serves to identify the additive and product synergies of the mechanism. We will return to this elicitation task in further detail in Sect. 4.2. In summary, as soon as a cluster of classes is associated with a causal mechanism under study, some of the qualitative features of its pattern of interaction are fixed. The classes thus provide a means for readily distinguishing between features that can be looked up and features for which additional knowledge has to be elicited from the expert. We would like to note that the situations in which two or more probabilities appear to be equal are also covered by the procedure, albeit that one or more steps in the procedure may become redundant; the procedure can be easily extended to address these situations explicitly, however.

To conclude, we briefly comment on the acquisition of patterns of interaction for mechanisms involving three causes. As we argued in Sect. 3, such a mechanism

A group of 100 patients with:

Prolonged rupture of the foetal membranes: **yes**
Diaphragmatic hernia: **yes**

Fig. 4. A case card for the domain of neonatology

is composed of three partial mechanisms, each comprising two of the causes. For the acquisition, we build upon this fundamental observation and subdivide the mechanism into its partial mechanisms. For each partial mechanism, an appropriate pattern of interaction is established, basically by using the procedure outlined above. The three patterns of interaction obtained then are combined to yield the overall pattern for the mechanism under study.

4.2 Techniques for the elicitation of probability orderings

Our procedure for acquiring patterns of interaction includes various steps in which the domain expert is asked to indicate an ordering for a number of conditional probabilities. Since reasoning in terms of frequencies is generally perceived as less demanding than reasoning in terms of probabilities [9], we use the frequency format for presenting the probabilities to be ordered to the expert. For the mechanism of Fig. 2, for example, the probabilities $\Pr(c \mid ab)$ and $\Pr(c \mid \bar{a}\bar{b})$ equal

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

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

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. Indicating an ordering for the probabilities $\Pr(c \mid ab)$ and $\Pr(c \mid \bar{a}\bar{b})$ 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. To support the domain expert in this task, the various groups of cases are presented on *case cards*. For convenience of comparison, the information on each card is represented in an easily surveyable manner. To make the ordering task more concrete, moreover, we randomly chose a fixed size for the various groups. An example case card for the domain of neonatology is shown in Fig. 4.

In the first phase of our acquisition procedure, a partial ordering of probabilities is to be obtained. For Fig. 2, for example, a partial ordering is required of the probabilities $\Pr(c \mid ab)$, $\Pr(c \mid a\bar{b})$, $\Pr(c \mid \bar{a}b)$ and $\Pr(c \mid \bar{a}\bar{b})$. A case card is constructed for each of the four groups of cases as described above; the resulting four case cards are presented to the expert, who is asked to select the group

We only consider patients with
Sepsis: **yes**

Fig. 5. A context bar for the domain of neonatology

containing the largest number of cases in which c holds and the group having the smallest number of such cases.

After the expert has given the partial ordering, either a total ordering of the probabilities is required, or knowledge for establishing the additive and product synergies involved. For the first of these tasks, the same case cards are used as before. The domain expert is asked to assign an order to the four groups of cases, rather than to select specific groups. For establishing the additive synergy for Fig. 2, for example, information is to be obtained as to whether or not

$$\Pr(c \mid ab) + \Pr(c \mid \bar{a}\bar{b}) \geq \Pr(c \mid a\bar{b}) + \Pr(c \mid \bar{a}b)$$

In essence, the task involved for the domain expert again is an ordering task, in which the two sums of probabilities $\Pr(c \mid ab) + \Pr(c \mid \bar{a}\bar{b})$ and $\Pr(c \mid a\bar{b}) + \Pr(c \mid \bar{a}b)$ are to be compared. We visualise each sum by joining the two cards with the two groups of cases under study, to yield a compound group. The expert is then asked to assign an order to the two compound groups thus constructed.

For establishing the product synergies, we do not build upon their definition in terms of products of probabilities directly, since comparing products of fractions at a qualitative level is known to be very hard. Instead, we build upon the influence between two causes that is induced by the observation of a value for their common effect. For establishing for Fig. 2 the product synergy of A on B given c, for example, information is to be obtained as to whether or not

$$\Pr(b \mid acy) \geq \Pr(b \mid \bar{a}cy)$$

for any combination of values y for all causes of B other than A; note that since we consider mechanisms with two causes only, the effect C does not have any other causes to be taken into account. We refer to the combination of values y as the *context* of the influence. The domain expert in essence now has to perform a number of ordering tasks, one for each possible context y. We view the context as background knowledge and ask the expert to indicate an ordering for the probabilities $\Pr(b \mid ac)$ and $\Pr(b \mid \bar{a}c)$ for each such context, using case cards as before. For each separate ordering task, the context is visualised by a card, termed the *context bar*; this card is put over the two case cards to be ordered. An example context bar for the domain of neonatology is shown in Fig. 5. After obtaining the product synergy given c for each context, the overall synergy is established as defined in Sect. 2.

For the acquisition of interaction patterns for mechanisms with three causes, we subdivide such a mechanism into its partial mechanisms as described in Sect.

4.1. For each partial mechanism, an interaction pattern is established given each value of the third cause. The two patterns thus obtained are then combined to yield the overall pattern for the partial mechanism. The patterns of interaction for the three partial mechanisms are subsequently combined to yield the pattern for the entire mechanism. We would like to note that the probability orderings that are elicited from the domain expert for the different partial mechanisms may be mutually inconsistent. If such an inconsistency arises, we propose to ask the expert to assign a total ordering to eight case cards, describing the eight conditional probabilities for the common effect of the mechanism under study.

5 Experiences with our procedure

We used our procedure for acquiring the patterns of interaction for different parts of a Bayesian network in the making for the domain of neonatology. The network currently consists of some fifty causal mechanisms, a handful of which have been considered in the evaluation of our procedure. The graphical structure of the network has been developed with the help of an intensive-care neonatologist, who is the third author of the present paper. We acquired from the same expert, the knowledge for establishing the various qualitative influences and synergies. The knowledge engineers, who are the first two authors of the paper, divided the engineering tasks involved in conducting the knowledge acquisition session. The session was held in an office with a large empty table that provided sufficient space for the expert to rearrange the various cards used with our procedure.

In reviewing the experiences with the application of our acquisition procedure, we focus on the results obtained for the mechanism that pertains to a hypoplastic lung and its two causes, shown in Fig. 1. We would like to note that the four conditional probabilities indicated in the figure for the presence of a hypoplastic lung, were estimated from patient data available from the domain expert's home institute. For the elicitation of the pattern of interaction for the mechanism under study, we constructed four case cards, one for each combination of values for the two causes of a hypoplastic lung. In case knowledge had to be acquired for establishing the product synergies, two additional case cards were constructed, for the presence and for the absence of a diaphragmatic hernia, respectively; since the causes of a hypoplastic lung do not have any causes themselves in the network, no context bars were required.

We presented the expert with the first four case cards and asked him to select, from among the four groups, the group with the largest number of newborns with a hypoplastic lung and the group with the smallest such number. He indicated the group of newborns in whom both causes are present as the group with the highest occurrence of a hypoplastic lung; he further indicated the group of patients in whom both causes are absent as the group in which the fewest cases of an underdeveloped lung occur. Based on the partial ordering of the underlying probabilities, we associated class *II* with the mechanism. From the partial ordering, moreover, we established positive qualitative influences of the two causes separately. As prescribed by our procedure, further knowledge elicitation was focused on the two types of synergy. For the additive synergy, we constructed compound groups of patients, as described in the previous section. The expert initially indicated that the two compound groups of patients had equally many newborns with a hypoplastic lung; after some further consideration, however, he expressed that the compound group with the patients in whom just one of the causes is present, was likely to include more patients with a hypoplastic lung. The acquisition results further indicated the product synergy given the presence of a hypoplastic lung to be negative. We would like to note that from the expert a negative additive synergy was elicited, while the data indicated a positive one. The number of patients in the available data set in whom either one or both causes are present, is very small, however, as a consequence of which the probability estimates obtained may be strongly biased.

After the acquisition session, the domain expert mentioned that he found the task of ordering groups of patients to be quite easy in general. His observation was supported by the finding that he selected the groups with the highest and the lowest number of patients with a condition under study, within a second. The task of indicating an ordering for two compound groups of patients clearly was more demanding; the expert indicated that he tended to associate concrete numbers with the cards, in order to be able to compare the two groups involved. Still, the task was generally performed within 30 seconds.

To conclude, we would like to mention that some of the causal mechanisms under study included a medical intervention as one of its causes. During the acquisition session, the expert appeared to interpret the intervention as stating that *the intervention is called for* rather than as stating that *the intervention has been performed*. Afterwards the expert indicated that he found the co-occurrence of truly statistical events and interventions in a single mechanism to be confusing and counterintuitive.

6 Conclusions and future work

The task of obtaining all probabilities required for a Bayesian network, can be simplified by first establishing the qualitative pattern of interaction for each causal mechanism; these patterns can subsequently be used as constraints on the precise probabilities to be obtained. In this paper, we elaborated on this idea and designed a dedicated procedure for acquiring interaction patterns; in addition, we proposed various elicitation techniques to be used for interaction with a domain expert. Initial experiences with our procedure in the domain of neonatology indicated that it requires relatively little effort on the part of the expert. In the future, we intend to further develop our procedure, for example by explicitly providing for equal probabilities. Moreover, we plan to investigate methods for dealing with variables that capture interventions and to focus on the task of obtaining precise probabilities that meet the established constraints.

In designing our acquisition procedure, we built to a large extent on a profound study of the fundamental properties of the concepts of qualitative influence and synergy. We feel that designing an acquisition procedure from the results of such a study has proved beneficial and has in fact resulted in an efficient dedicated procedure for one of the harder tasks in building Bayesian networks.

References

1. F.V. Jensen. *Bayesian Networks and Decision Graphs*. Springer-Verlag, New York, 2001.
2. E.M. Helsper, L.C. van der Gaag. Building Bayesian networks through ontologies. In: F. van Harmelen, editor, *Proceedings of the 15th European Conference on Artificial Intelligence*. IOS Press, Amsterdam, 2002, pp. 680 – 684.
3. M.J. Druzdzel, L.C. van der Gaag. Building probabilistic networks: "Where do the numbers come from?" Guest editors' introduction. *IEEE Transactions on Knowledge and Data Engineering*, 2000, **12**, pp. 481 – 486.
4. M.J. Druzdzel, L.C. van der Gaag. Elicitation of probabilities for belief networks: combining qualitative and quantitative information. In: Ph. Besnard, S. Hank, editors, *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers, San Francisco, 1995, pp. 141 – 148.
5. S. Renooij, L.C. van der Gaag. From qualitative to quantitative probabilistic networks. In: A. Darwiche, N. Friedman, editors, *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers, San Francisco, 2002, pp. 422 – 429.
6. M.P. Wellman. Fundamental concepts of qualitative probabilistic networks. *Artificial Intelligence*, 1990, **44**, pp. 257 – 303.
7. L.C. van der Gaag, E.M. Helsper. Defining classes of influences for the acquisition of probability constraints for Bayesian networks. In: *Proceedings of the 16th Conference on Artificial Intelligence*, 2004, to appear.
8. M. Henrion, M.J. Druzdzel. Qualitative propagation and scenario-based approaches to explanation in probabilistic reasoning. In: P.P. Bonissone, M. Henrion, L.N. Kanal, J.F. Lemmer, editors, *Uncertainty in Artificial Intelligence 6*. Elsevier, North-Holland, 1991, pp. 17 – 32.
9. G. Gigerenzer, U. Hoffrage. How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review*, 1995, **102**, pp. 684 – 704.