# BUILDING KNOWLEDGE-BASED SYSTEMS BY CREDAL NETWORKS: A TUTORIAL

***Alberto Piatti*[1], *Alessandro Antonucci*[2] *and Marco Zaffalon*[2*]**

[1] Mathematical and Physical Sciences Unit (SMF),
University of Applied Sciences and Arts of Southern Switzerland (SUPSI),
Galleria 2, CH-6928 Manno (Lugano), Switzerland

[2] Istituto Dalle Molle di Studi sull'Intelligenza Artificiale (IDSIA),
Galleria 2, CH-6928 Manno (Lugano), Switzerland

**Abstract**

*Knowledge-based systems* are computer programs achieving expert-level competence in solving problems for specific task areas. This chapter is a tutorial on the implementation of this kind of systems in the framework of *credal networks*. Credal networks are a generalization of Bayesian networks where *credal sets*, i.e., closed convex sets of probability measures, are used instead of precise probabilities. This allows for a more flexible model of the knowledge, which can represent ambiguity, contrast and contradiction in a natural and realistic way. The discussion guides the reader through the different steps involved in the specification of a system, from the evocation and elicitation of the knowledge to the interaction with the system by adequate inference algorithms. Our approach is characterized by a sharp distinction between the *domain knowledge* and the process linking this knowledge to the perceived evidence, which we call the *observational process*. This distinction leads to a very flexible representation of both domain knowledge and knowledge about the way the information is collected, together with a technique to aggregate information coming from different sources. The overall procedure is illustrated throughout the chapter by a simple knowledge-based system for the prediction of the result of a football match.

## 1 Introduction

*Knowledge-based systems* (KBSs, [40]) are computer programs that achieve expert-level competence in solving problems on specific task areas. KBS are based on a coded body of *human knowledge* represented through a mathematical model. The task of developing a KBS is generally known as *knowledge engineering*, while the specific task of collecting

*E-mail addresses: alberto.piatti@supsi.ch, alessandro@idsia.ch, zaffalon@idsia.ch.

human knowledge and representing it through a mathematical model is known as *elicitation*. The kind of human knowledge we consider in this chapter is meant to correspond to the beliefs of a single expert or to the shared beliefs of a pool of experts about a specific domain. We refer to the single expert or to the pool of experts queried during the development of the system by using the term *expert* and the male pronoun *he*, without any implicit reference to the gender of the expert(s). Furthermore, we refer to his beliefs using the term *expert knowledge*.

When developing a KBS, one should select a mathematical formalism for the representation of the expert knowledge. Research in elicitation has mainly focused on the quantification of expert knowledge through subjective probabilities (like in Bayesian approaches).[1] In particular, the combination of knowledge representation through directed acyclic graphs and conditional probabilities has led to the development of *Bayesian networks* (BN), which are nowadays one of the most important mathematical formalisms for KBSs.[2] A BN is based on a collection of variables and on a directed acyclic graph, whose nodes are in one-to-one correspondence with the given variables. The graph is used to represent probabilistic independence relations between the variables according to the *Markov condition*: every variable is independent of its non-descendant non-parents conditional on its parents.[3] The Markov condition is clarified in Figure 1 by means of three examples, illustrating the meaning of the typical patterns that can be observed in a directed graph. More general graphs can be understood on the basis of these three patterns. The graph associated with a BN can be regarded as the qualitative part of the model, while the quantitative part is specified by means of conditional probabilities: for each variable in the network, a conditional probability mass function for each possible combination of values of the parent variables should be specified. Depending on the topology of the graph associated to the BN, the amount of conditional probabilities to be specified can be very large and consequently the specification of the quantitative part can become problematic. Besides this potentially huge number of probabilities to be quantified, another issue is that expert knowledge is mostly qualitative, while the approaches to BNs quantification usually require the expert, implicitly or explicitly, to model his beliefs by “precise” (i.e., single-valued) probabilities.[4]

We agree with Walley [58] in saying that Bayesian precise probabilities are not the most suited mathematical formalism for representing uncertain expert knowledge in a KBS. Walley provides many motivations for that; among them the fact that a proper measure should be able to model partial or complete ignorance, limited or conflicting information, qualitative judgments: all these features can be hardly represented in the Bayesian framework without making some extra assumptions. Consider for instance an expert claiming that *it is more probable that a football team will win a match playing on its home ground than away*: how can we formalize this sentence in the Bayesian framework? This comparative judgment defines a constraint between two (conditional) probabilities, which can be satisfied by infinitely many assignments of their values. Thus, if we want to meet the requirement of *precision* of Bayesian probabilities (i.e., that a probability is specified by a single number),

[1] See O’Hagan et al. [46] or Jenkinson [37] for an overview on elicitation of subjective probabilities.

[2] See Jensen and Nielsen [38] for an overview on BNs and Pourret et al. [48] for a list of applications.

[3] Given an arrow from node $A$ to node $B$, we call $A$ *parent* of $B$, and $B$ *child* of $A$. $C$ is *descendant* of $A$, if there is a directed path connecting $A$ and $C$.

[4] See for example Van Der Gaag et al. [31, 32, 56], Druzdzel et al. [25] and Renooij et al. [49, 50].

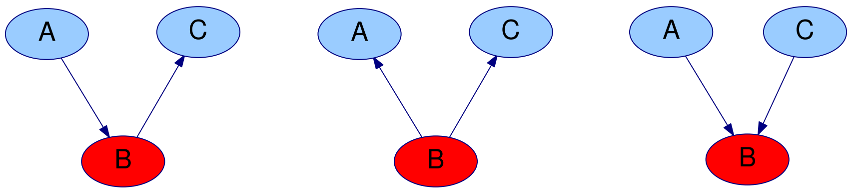


Figure 1: Three types of connections to interpret the conditional (in)dependence relations in a directed graph with three variables according to the Markov condition. In the first case, on the left-hand side, $C$ is conditionally independent of $A$ given $B$ and vice versa. In other words, if we know the value of $B$, then the variables $A$ and $C$ are independent. Also in the second case, $A$ and $C$ are conditionally independent given $B$. The remaining case illustrates a situation of conditional dependence between the variables $A$ and $C$ given $B$. In other words, if $B$ is known, then $A$ and $C$ are dependent, otherwise they are independent.

we need some additional assumption in order to select a particular assignment among the feasible ones. Similar assumptions cannot always be allowed or justified, as the arbitrariness they involve could unpredictably bias the conclusions inferred by the system. This is a common problem encountered in the elicitation through precise-probability models, which is related to the qualitative (and hence imprecise) nature of the expert knowledge.

In order to cope with problems of this kind, a number of approaches relaxing precision in probabilistic assessments has been proposed in the past, among them Choquet capacities [14], belief functions [30], possibility measures [26], and Walley's theory of coherent lower previsions [57]. Under the subjective interpretation of probability, these theories can be all regarded as modelling beliefs by *credal sets* (CSs) [42], i.e., closed convex sets of precise measures; in particular, Walley's theory includes all the others as special cases [23, 24, 58]. The combination of knowledge representation through directed acyclic graphs and knowledge quantification by CSs has led to the development of *credal networks* (CNs) [16, 17], which can be regarded as a generalization to CSs of BNs. Also CNs are represented through a qualitative and a quantitative part.

The qualitative part of a CN consists of a directed acyclic graph representing independence relations between a collection of discrete variables through the Markov condition. The only difference with respect to BNs is in the notion of independence. The standard notion of *stochastic* independence[5] employed for BNs can be extended to the case of CNs in many ways. Among the different notions of independence considered in the literature,[6] two of them seem to be particularly suited for the modelling of real problems: *epistemic independence* and *strong independence*. The former is appealing because of its intuitive behavioral interpretation, while the interpretation of the latter is more related to a sensitivity analysis approach. Loosely speaking: two variables are *epistemically independent* if the observation of one of the two variables does not affect our beliefs about the other variable, while they are *strongly independent* if their joint probability according to the *extreme points*[7] of their (joint) CS obey the standard notion of stochastic independence. Strong inde-

[5]Two variables are stochastically independent if their joint probability factorizes.

[6]See Couso et al. [15], Moral and Cano [45] and Cozman [17] for a discussion and further references.

[7]An extreme point of a CS is an element that cannot be expressed as convex combination of other points in the set.

pendence is a particularly suitable assumption if the variables are assumed to be described by a probability measure whose numerical values are not precisely known [45]. It can be shown (e.g., [45]) that strong independence implies epistemic independence: the latter is therefore a less committal assumption than the former. Despite its less intuitive interpretation and its strength, strong independence is the most used notion of independence for CNs. In fact, several exact and approximate inference algorithms have been developed for CNs under strong independence, while algorithms for CNs under epistemic independence have been developed only for particular cases.[8] Accordingly, in this chapter, we only consider CNs with strong independence.

The quantitative part of a CN consists in a collection of conditional CSs, i.e., convex sets of conditional probability mass functions, one for each variable in the network and state of its parents. BNs are very useful tools for the modelling of KBSs, in particular thanks to the graphical representation of the variables (and their independence relations), which is considered as particularly intuitive by the experts. CNs maintain this desirable feature, allowing at the same time for a more suited and cautious approach to the elicitation task.[9] Accordingly, we regard CNs as an effective and flexible mathematical tool for KBSs construction.

Here we present a procedure for building KBSs based on the formalism of CNs, which has been proved to be effective in practice, like for example in airspace surveillance [1] and environmental monitoring [3]. The approach we propose satisfies several conditions that, according to our practical experience, are important for an effective representation of the expert knowledge in a KBS. The conditions are the following.

- **Imprecision**: it should be possible to model qualitative expert knowledge without introducing unrealistic assumptions.
- **Aggregation**: if knowledge comes from many experts, it should be possible to fuse the different pieces of information into a single representation, again without forcing unrealistic assumptions.
- **Modularity**: an easy identification and modification of specific parts of the representation should be possible, without requiring the modification of the rest.
- **Transferability**: the final representation of knowledge should be easily represented by a computer.
- **Observation**: it should be possible to provide a realistic description of the way the evidence is collected. For example, this should hold also for sensors fusion, even if they are reporting contrasting or missing values.
- **Separation**: the expert knowledge can be divided in two parts: a part describing the theoretical knowledge, which we call the *domain knowledge*, and the knowledge about the process linking this knowledge to the perceived evidence, which we call the *observational process*. Often, in practice, the same domain knowledge is used in situations characterized by different observational processes. For this reason, it is

[8]See for example de Cooman et al. [21], where an exact algorithm has been developed for dealing with epistemic irrelevance in tree-shaped CNs.

[9]The counterpart of such a higher expressive power is an increased computational complexity in the inferences [17]. Yet, as shown in Section 4, this does not prevent the implementation of KBSs based on CNs.

convenient to separate the specification and representation of domain knowledge and observational process. In this way, it is possible to adapt a KBS to practical situations characterized by different observational processes, without the need of redefining the domain knowledge.

- **Reasoning**: the representation should offer a valid support to the expert reasoning, by inducing a further elaboration and organization of the knowledge.
- **Completeness**: the representation should allow for a complete specification of the relevant knowledge in the problem under consideration.
- **Ontology**: the representation should be accessible and understandable without particular technical skills.

The **imprecision**, **aggregation**, **modularity** and **transferability** conditions are automatically satisfied by CNs, while the other conditions are satisfied because of the particular procedure proposed here. As discussed in Walley [58], CSs are particularly suitable for representing qualitative knowledge, and therefore the imprecision condition is satisfied. Aggregation can be considered as a special case of imprecision [10] and is satisfied as well. The modularity condition is satisfied by probabilistic graphical models in general, and in particular by CNs, because these models represent knowledge separately for each variable in the network given the parents: the quantification of a variable can be therefore changed without the need of modifying the rest of the network. Finally, CNs satisfy transferability for the simple reason that graphs and CSs can be easily stored and elaborated by a computer. The **observation** and the **separation** conditions are satisfied by our particular approach, thanks to the sharp distinction we adopt between the *domain knowledge* and the *observational process*. The **reasoning**, **completeness** and **ontology** conditions are satisfied by the proposed elicitation procedure, because we focus not only on the translation of beliefs into CSs, but also on a preliminary evocation, organization and representation of them in natural language. In fact we have introduced an intermediate step between the specification of the qualitative part of the CN and the quantification of the CSs, in which the beliefs of the expert are listed in an organized way, using arguments as a semi-informal language [54]. As highlighted by Helsper and Van der Gaag [35], during the specification of a probabilistic network the expert knowledge is often elicited by directly using mathematical formalisms, without constructing an (intermediate) *explicit knowledge model*, i.e., an *ontology* [54]. The task to explicitly reconstruct the knowledge embedded in a probabilistic network can be therefore difficult and time consuming, especially for people without technical skills. Accordingly, the knowledge embedded in a large and complex network may often become inaccessible to people which have not been involved in its development from the beginning, with consequent problems in its documentation, maintenance and adoption by new users. The combination of the qualitative part of the network and the list of arguments produced in the proposed elicitation procedure can be regarded as an ontology providing an explicit model for the knowledge embedded in the network. This ontology allows non-experts to access the expert knowledge in the KBS without the need of accessing and interpreting the CSs. At the same time, the construction of this ontology helps the expert to recall in mind and focus on all the aspects that are relevant for the problem under consideration, through reasoning and judgement.

The outline of the chapter is as follows. In Section 2 we detail and motivate the distinction between domain knowledge and observational process. In Section 3 we describe a procedure for the elicitation of a CN satisfying all the conditions stated above. The different elicitation steps are illustrated through the construction of a CN for predicting the result of a football match. The developed model is then used in Section 4 to illustrate how CNs are used in practice. Concluding remarks are finally reported in Section 5.

## 2 Domain Knowledge and Observational Process

In this section, we introduce and motivate the concepts of *domain knowledge* and *observational process* by means of the following example.

**Example 2.1 (Cold or flu?)** *Suppose that a kid shows the symptoms of a chill and his mother wants to verify whether these symptoms are due to a simple cold or to a flu (other possible causes are excluded for sake of simplicity). The mother knows that the presence of a flu is typically associated with a high body temperature (fever), while a simple cold usually is not. To verify the presence of a high body temperature, the mother can simply put her hand on her kid's forehead, or measure his body temperature through a thermometer. If the perceived (or measured) temperature is high, usually the mother concludes that the kid is having a flu, while if the perceived (or measured) temperature is low, the mother excludes the presence of a flu.*

This example can be described by the following four variables.

- ***Cold or flu*.** This variable describes the presence of a cold or of a flu, with possible values *cold* and *flu*.
- ***Fever*.** This variable describes the presence of fever, with possible values *yes* and *no*.
- ***Temperature (hand)*.** This variable describes the result of the measurement of the temperature by hand, with possible values *yes*, if the perceived temperature is high, *no* if the perceived temperature is not high, and *missing* if the temperature has not been measured by hand or it has been measured but the result is not available.
- ***Temperature (thermometer)*.** This variable describes the result of the measurement of the temperature by a thermometer, with as possible values *yes*, if the measured temperature exceeds 37 Celsius degrees, *no* if the measured temperature is below 37 Celsius degrees, and *missing* if the temperature has not been measured by thermometer or it has been measured but the result is not available.

The theoretical knowledge relevant for this example correspond to the assessed relationships between the first two variables. These variables cannot be directly observed: they are *latent* variables [51] whose actual values are directly inaccessible. The values of the latter two variables describe how information about the latent theoretical variables is gathered. The value of these variables is always known: they are therefore *manifest* variables [51]. More specifically, the two manifest variables are observed in order to gather information about the latent variable *Fever*, that in turn is related to the latent variable of interest *Cold or flu*.

Now suppose you want to build a CN describing the situation in our example. Zaffalon and Miranda [60] have proposed an approach to the construction of CNs based on the distinction of two information levels: the *latent* and the *manifest level*. The latent level consists of (theoretical) latent variables that are relevant for the problem but that cannot be observed directly. The manifest level consists of manifest variables that are observed in practice to gather information about the variables of the latent level. In our case, the variables *Cold or flu* and *Fever* belong to the latent level, while the other two variables belong to the manifest level. The combination of the variables of the latent level and their conditional (in)dependence relations is called *domain knowledge*. In general, the domain knowledge is used to represent the theoretical knowledge that is relevant for the problem under consideration. A CN describing the domain knowledge for our example is in Figure 2(a). Suppose now that the temperature is measured only by the mother's hand: we can extend the network in Figure 2(a) in order to include a description of how evidence about the presence of fever is gathered in practice. A possible extension is illustrated in Figure 2(b). In general, the combination of the variables of the manifest level and the conditional (in)dependence relations between latent variables and manifest variables used to gather information about them is called *observational process*. In general, the observational process is used to describe the relationship between theoretical knowledge and perceived reality. In this case, the observational process consists in the variable *Temperature (hand)* and in the link connecting it with the latent variable *Fever*.

In general, the construction of a CN based on the distinction between domain knowledge and observational process consists of two steps. In the first step the dependence relations between the variables of the latent level are represented through a directed acyclic graph and CSs are quantified for each variable. The result of the first step is a CN describing the domain knowledge. In the second step, the graph is augmented by the variables of the manifest level and with arrows connecting each variable of the latent level with the variables in the manifest level that are used to gather information about it. The corresponding CSs are then quantified for each variable of the manifest level.[10] The result is a CN describing both the domain knowledge and the observational process. This approach to the construction of CNs has proven to be useful and convenient to deal with missing or incomplete data [5, 22, 59, 60] and with multiple observations and sensor fusion [1] and satisfies consequently the **observation** condition. Furthermore, this approach satisfies the **separation** condition: the same domain knowledge can be adapted to different practical situations simply redefining the observational process, and leaving the domain knowledge unchanged. For example, in Figures 2(b)–2(e), several observational processes have been defined for the same domain knowledge.

Although the steps required for the specification of the domain knowledge and the observational process are the same, the experts taking care of them are usually different. For example, in the situation above, the *domain knowledge* consists of the knowledge relating illness (cold, flu) and body temperature. The expert responsible for the specification of this part could be, for example, a physician (or a pool of them). The *observational process* depends on how the evidence about the body temperature is collected. If the mother measures the temperature by using her hand, then the observational process consists in the relation

[10] For a general philosophical discussion about the modelling of the observational process see Bagozzi et al. [9] and [27] and Borsboom et al. [12].

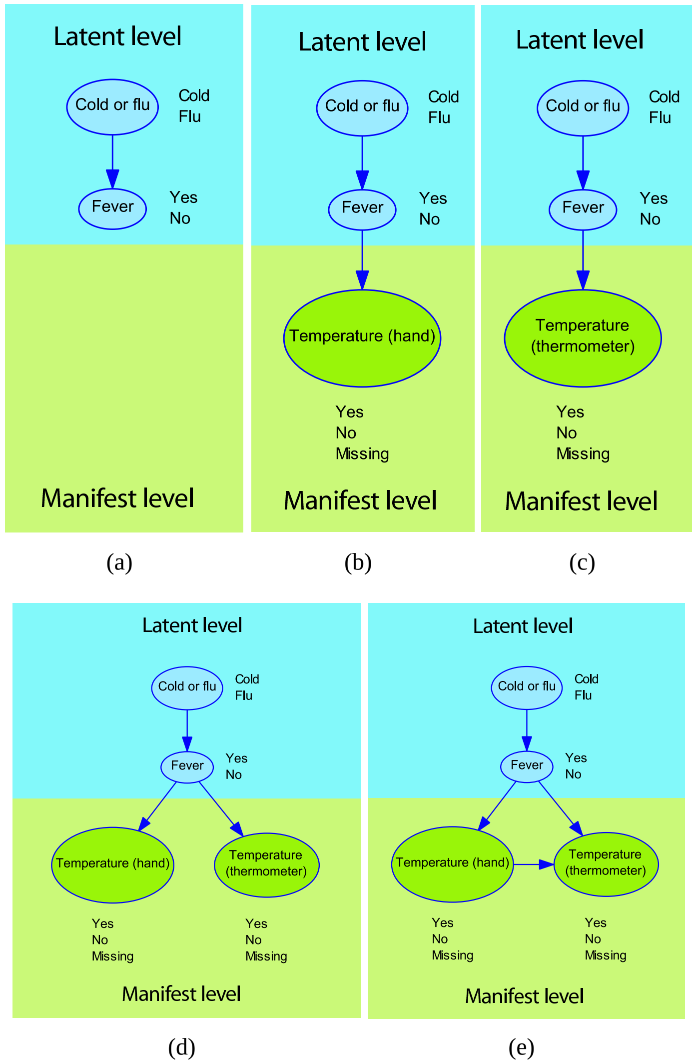


Figure 2: Several networks modelling the situation of Example 2.1. The domain knowledge, describing the relationship between the presence of flu or cold and the presence of fever, is the same for all the networks and it is displayed in Figure 2(a). The other networks represent the combination of the domain knowledge with different observational processes. Figure 2(b) represents an observation through temperature measurement by hand only, Figure 2(c) an observation through temperature measurement by thermometer only, Figure 2(d) an observation through temperature measurement by hand and by thermometer where the two observations are considered to be conditional independent given the (unobservable) value of the variable *Fever*, and finally Figure 2(e) represents an observation through temperature measurement by hand and by thermometer where the two observations are considered to be dependent. The latter structure could for example model a situation where the temperature is measured through thermometer only if the temperature measured by hand indicates fever. In this case, if the measurement by hand indicates *no* fever, then the temperature measured by thermometer is surely *missing* because it is not measured.

between true and perceived temperature. The expert to be queried during this part can be the mother herself, as she knows by experience how her kid's body temperature and her sensation are related. In other cases, the body temperature of the kid could be measured by a thermometer, and the observational process would be the relation between true and measured body temperatures. The expert to be queried for the specification of this part would be for instance the manufacturer of the thermometer, who knows the level of precision and the drawbacks of the device.

## 3 Elicitation

In a definition shared by many authors [37, 46], Walley [57, p. 167] defines elicitation as the *process by which beliefs are measured through explicit judgments and choices*. In our procedure, we intend elicitation as a more general process that, apart from measuring beliefs, is also concerned with their evocation and organization. We follow Benson et al. [11] in distinguishing two sequential steps for a complete elicitation process:

(i) the *belief assessment*, which is the process of inducing the expert to recall in his mind all the aspects that are relevant for the problem under consideration through reasoning and judgment;

(ii) the *response assessment*, which involves the measurement of the evoked expert knowledge by an appropriate mathematical model, after the necessary structuring of this information.

Most of the existing procedures for elicitation focus on the response assessment step, while the role of the belief assessment is often underestimated. According to our own experience, belief assessment is crucial in order to make elicitation effective and complete, and satisfy the **reasoning** and **completeness** conditions without biases in the specification of knowledge.[11] Furthermore, the belief assessment can be structured in order to produce an explicit knowledge model, thus satisfying the **ontology** condition.

Our elicitation procedure, which is graphically depicted in Figure 3, is divided in four steps. In the first two steps the domain knowledge is described through a complete CN. In the second two steps, the former network is augmented by the model of the observational process. The final result is a new CN embedding both the domain knowledge and the observational process, and an ontology describing these two parts in natural language. The four steps are the following.

1. **Belief assessment for the domain knowledge:** in this step expert knowledge is recalled, organized and finally represented through a directed acyclic graph and a list of arguments.
2. **Response assessment for the domain knowledge:** in this step, CSs consistent with all the arguments expressed in the previous step are constructed for each variable (node) in the network.
3. **Belief assessment for the observational process:** in this step, the way evidence is collected in practice is described, for each variable in the domain knowledge, through

[11] For a survey of the psychological aspects of elicitation of probabilities see Daneshkhah [18] and Kahneman et al. [39].

an extension of the previous graph with variables modelling the observations and a list of arguments describing them.

4. **Response assessment for the observational process:** finally, CSs consistent with the arguments expressed in the previous step are constructed for each variable of the observational process.

This procedure is particularly convenient if the domain expert and the observational process expert are different individuals. If the expert is the same, the whole procedure can be reduced to two steps: belief and response assessment for the whole network. Yet, according to our experience, reasoning at the same time about domain knowledge and observational process can be somewhat misleading and we therefore recommend to apply the four steps procedure anyway. In the rest of this section, we describe in detail each step of the elicitation procedure and we illustrate them with the help of the following example.

**Example 3.1 (Football example)** *We consider the Swiss Football Super League. The goal is to develop a KBS that predicts the result of a football match between Bellinzona and Basel, to be played next week, after some months from the beginning of the championship. The domain expert is supposed to be a person that knows the two teams and the Swiss football championship quite well. His knowledge is based on technical considerations and on his own experience. The observational process expert is supposed to be a supporter of the Bellinzona football team aiming to predict the result of the match. His knowledge consists in knowing where and how evidence about the domain knowledge can be gathered.*

The complete CN for this example is illustrated in Figure 12.

## 3.1 Belief Assessment for the Domain Knowledge

We follow Browne et al. [13] in adopting a belief assessment procedure based on two complementary evocation and organization techniques: *knowledge map representation* [28, 36, 44] and *reasoning-based directed questions* [29]. The *knowledge map* (KM) consists, in our case, of a directed acyclic graph in which the relevant variables are modelled through discrete variables and depicted as nodes of the network, while dependence relationships between the variables are represented through the arrows according to the Markov condition. This type of KMs corresponds roughly to the one described by Howard [36]. Concerning the *reasoning-based directed questions*, they are a structured list of arguments reporting the expert knowledge in an organized way using natural language. The result of the belief assessment procedure is a directed graph whose nodes are associated to discrete variables and, for each variable, a list of arguments summarizing the knowledge of the expert about that variable. The combination of the graph, the description of the variables and their possible values in natural language and the arguments can be considered as an *ontology*, providing an explicit model for the expert knowledge. While the graphical structure can be passed as it is to a computer, the knowledge represented through arguments cannot, and should therefore be translated into a suitable mathematical model. This is exactly the task of the *response assessment*, where the expert knowledge that is not explicitly contained in the KM, i.e., the knowledge expressed through arguments, should be represented through CSs as described in Section 3.2. Here, we begin by describing the belief assessment procedure followed for the specification of the domain knowledge. The main features and a construction method

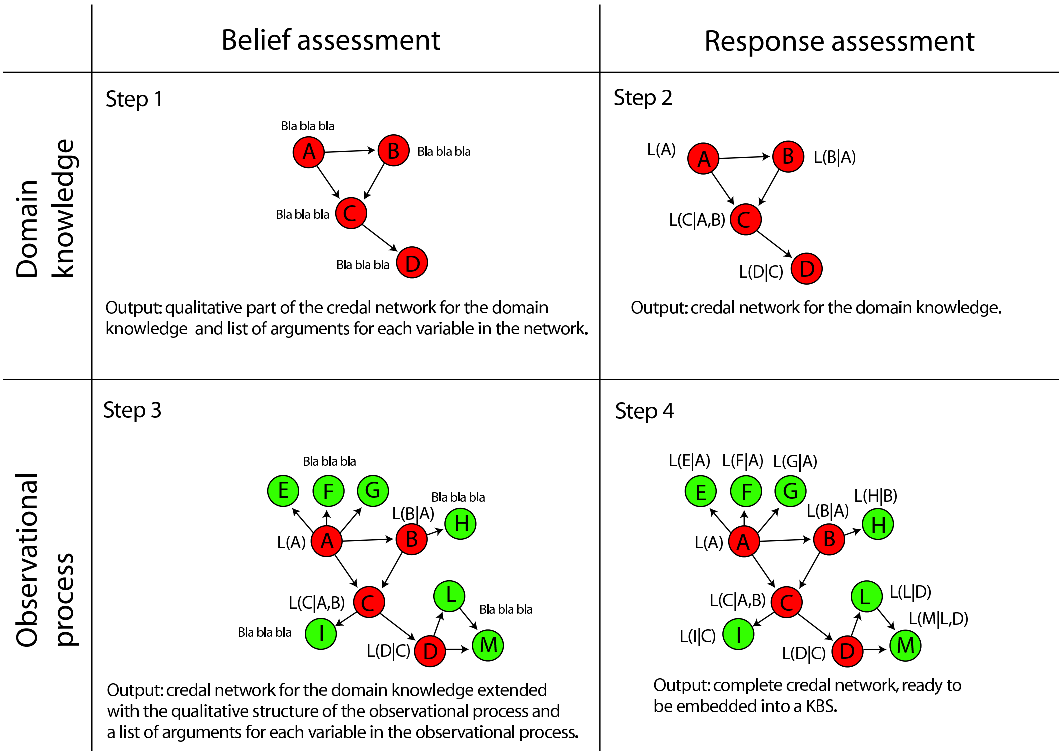


Figure 3: The elicitation procedure proposed in this chapter is divided in four steps. The final result is a CN embedding both the domain knowledge and the observational process. The symbol $L(\cdot\,|\,\cdot)$ is used to denote the set of conditional CSs associated with each variable and it is defined in Section 3.2.

for KMs are described in Section 3.1.1, while the specification of arguments is discussed in Section 3.1.2.

### 3.1.1 Knowledge Maps

Vail [55] gives the following definition of KM.

> *"A* knowledge map *is a visual display of captured information and relationships, which enables the efficient communication and learning of knowledge by observers with different backgrounds at multiple levels of detail. The individual items of knowledge included in such a map can be text, stories, graphics, model, or numbers. [...] Knowledge mapping is defined as the process of associating items of information or knowledge in such (preferably visual) a way that the mapping itself also creates additional knowledge."*

As illustrated by Eppler [28], KMs have been applied in a variety of contexts, with different purposes, contents and representation techniques. In this chapter, we consider the qualitative part of the domain knowledge as a KM, focusing thus on the particular definition of KM originally given by Howard [36] and further studied by Merkhofer [44] and by Browne et al. [13].

The first step in the construction of the KM in our case is to define exactly the *hypothesis variable*, i.e., the variable that should be predicted (or evaluated) by the KBS. In practical problems, we can have many hypothesis variables, but for the sake of simplicity we assume here a single hypothesis variable for the problem of interest. For instance, in the football example we consider, the variable that should be predicted is the result of the match of the next week between Bellinzona and Basel. We call it *Result* and as possible outcomes we assume *win*, *draw* and *loss*. An appropriate specification of the hypothesis variable is clearly crucial. In practice, one should feel comfortable with the chosen possible outcomes. On the one side, they should be detailed enough in order to allow one to specify, ideally, everything the domain expert knows about the problem under consideration (**completeness**). On the other side, they should be not too detailed, in order to avoid that the expert knowledge would be insufficient to construct a KBS producing informative results. An example of this latter case would be a definition of the possible outcomes of the match considering several possible results, like for example *win 2:1*, *win 5:3*, *loss 1:2*, etc..

The next step in the specification of the KM consists in the identification of the variables that the domain expert considers relevant for the problem, as they could have a (either direct or indirect) causal influence on the state of the hypothesis variable. The procedure is usually iterative: the variables having a direct influence on the hypothesis variable are first identified, then the variables having an influence on these variables, and so on. In this way the expert adds variables until he feels comfortable with the directed graph modelling these relations. The only aspect the domain expert should take care during this specification is to keep the graph acyclic. The procedure is illustrated by the following example.

**Example 3.2 (Variables affecting the result)** *In the football example, assume that, in his prediction of the result of the match, the domain expert considers the following variables:*

1. *the relative ranking of the two teams in the championship (*Relative ranking*);*
2. *the relative fitness of the two teams, defined as the comparison between the fitness state of the Bellinzona football team and the fitness state of the Basel football team (*Relative fitness*);*
3. *the level of support that can be expected from the supporters (*Support*).*

*These variables are assumed by the domain expert to have a direct causal influence to the state of the variable* Result*. These causal relations are depicted in Figure 4.*

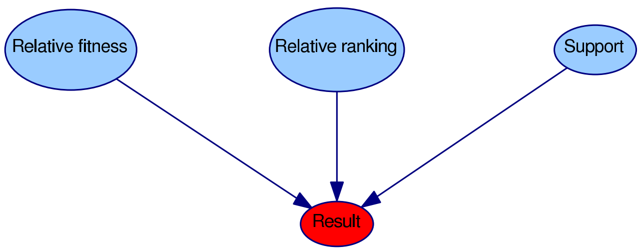


Figure 4: A simple network modelling the football example.

*Suppose now that the domain expert does not feel comfortable with this model because he needs other variables to represent his knowledge effectively. He has the possibility to add* *further variables, that can have a direct causal influence on the three variables specified in the first step, like for example the following.*

- *The* Relative ranking *of the two teams depends on the ranking of the Bellinzona football team (*Ranking*).*
- *The* Relative fitness *of the two teams depends on the* Fitness *of the Bellinzona football team and on the* Fitness of the opponent.
- *The quality of the* Support *depends on the enthusiasm of the supporters (*Supporters*) and on the place where the match is played (*Home match*).*

*Eventually, the domain expert can also add further variables that are relevant for those specified during the second step, like for example the following.*

- *The enthusiasm of the* Supporters *is influenced by the* Recent performance *of the team and by the* Ranking *of the team.*

*The graph in Figure 5 is obtained by considering these variables.*

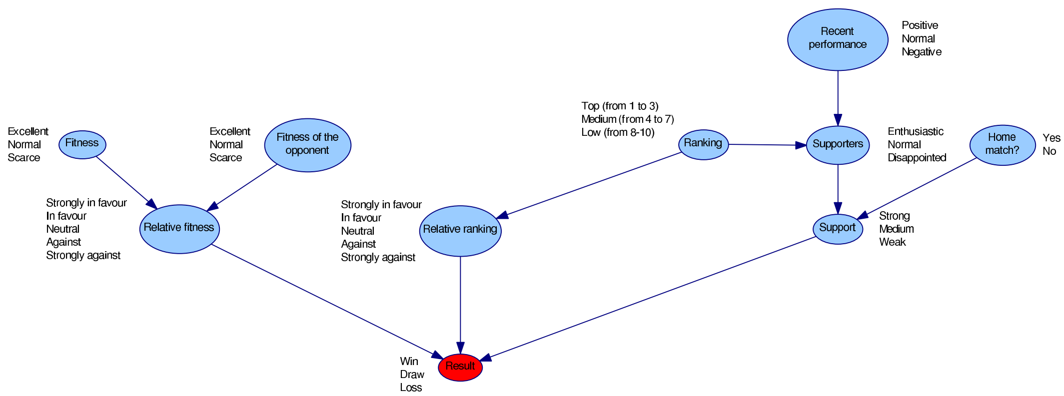


Figure 5: A network modelling the football example.

The above procedure returns a directed acyclic graph modelling the relevant variables considered in the domain knowledge and their relationships in terms of causal influence. The Markov condition allows for an interpretation of the graph in terms of conditional independence relations. In order to verify whether of not the relations associated to the graph are meaningful, the basic definitions and representations of conditional independence, depicted in Figure 1, can be used as a dictionary for the interpretation and validation. Let us describe this process by a simple example.

**Example 3.3 (Dependencies verification)** *Consider the variables* Relative Ranking, Result *and* Relative Fitness *in the graph of Figure 5. They are linked by arcs reproducing the pattern on the right-hand side of Figure 1. This corresponds to the assumption that, given the* Result, *the last two variables are dependent, and independent otherwise. Actually, a good* Relative ranking *is often uninformative about the current* Relative fitness *and vice versa.*

Once the verification of the dependencies is concluded, the last step for building the KM is an explicit specification of the possible values of the variables in the graph (i.e., those associated with the nodes of the graph). We model these quantities as *categorical variables*. This requires the specification, for each variable, of a finite set of *possible outcomes*, which are assumed to be *exclusive* and *exhaustive*.

**Example 3.4 (Football example)** *A possible specification of the possible outcomes of the variables considered in the domain knowledge is the following.*

- Result*.* Win, draw *and* loss*.*
- Fitness *and* Fitness of the opponent*. These two variables describe the fitness status of the two teams with three possible values:* excellent*, if the current fitness state of the team is extraordinary,* normal*, if the current fitness state reflects what could be expected from the team, and* bad*, if the current fitness state of the team does not reach what could be expected from the team.*
- Relative fitness*.* Strongly in favour*, if the* Fitness *is* excellent *and the* Fitness of the opponent *is* bad*,* in favour*, if the* Fitness *is* excellent *and the* Fitness of the opponent *is* normal *or if the* Fitness *of your team is* normal *and the* Fitness of the opponent *is* bad*,* neutral*, if the* Fitness *is the same as the* Fitness of the opponent*; the two further outcomes* against *and* strongly against *are defined in a similar way.*
- Ranking*.* Top*, from rank 1 to rank 3,* medium*, between rank 4 and rank 7,* low*, from rank 8 to rank 10.*[12]
- Relative ranking*.* Strongly in favour*, if the* Ranking *is* top *and the* Ranking of the opponent *is* low*,* in favour*, if the* Ranking *is* top *and the* Ranking of the opponent *is* medium *or if the* Ranking *is* medium *and the* Ranking of the opponent *is* low*,* neutral*, if the* Ranking *is the same as the* Ranking of the opponent*; the two further outcomes* against *and* strongly against *are defined in a similar way.*
- Support*.* Strong*, if the support is particularly motivating for the players,* medium*, if the support is as usual,* weak*, if the support is absent, depressing or stressing.*
- Supporters*.* Enthusiastic*,* normal*, or* disappointed*.*
- Home match*.* Yes*, if the match is played on the home football field,* no*, if the match is played on the home football field of the opponent.*[13]
- Recent performance*.* Positive*, if the recent performance of the team has been above the expectations,* medium*, if the recent performance of the team has reflected the expectations, and* negative*, if the recent performance has been unsatisfactory.*

The list of possible outcomes is usually accompanied by a description of the different outcomes, in order to be sure that all the persons involved in the specification of the domain knowledge (the domain expert) interpret the possible outcomes in (approximately) the same way (**aggregation**, **ontology**). The definition of the possible outcomes completes the specification of the KM. It is now interesting to illustrate what kind of knowledge the map contains. A KM refers to knowledge in three ways: (i) it contains knowledge expressed in an explicit way, (ii) it refers implicitly to tacit knowledge and (iii) it points to

[12] The Swiss Super League consists of ten teams.

[13] In the Swiss Super League matches are never played on neutral football fields.

explicit knowledge represented outside the map [28]. In our particular case, the three types of knowledge the network refers to are the following.

**Knowledge expressed explicitly in the KM:** our network describes explicitly what are the variables the domain expert considers for the hypothesis evaluation and their dependencies. The variables are defined using textual descriptions and through their possible outcomes, that are in turn also defined through textual descriptions. As suggested by Helsper and van der Gaag [35], a possible way to facilitate the understanding of the KM, is to prepare a *glossary*. The glossary should explicitly specify the meaning of all the (technical) terms used in the definitions of variables and possible outcomes in order to avoid misconceptions and ambiguities (**ontology**, **aggregation**).

**Tacit knowledge the KM implicitly refers to:** in the network the domain expert does not explain why he has specified the given dependencies. But usually these dependencies do not need to be explained, because they are self-evident. For example, in our network, the enthusiasm of the *Supporters* depends on the *Recent performance* of the team and on the *Ranking*. This is self-evident for every person who knows how a football match and a football championship work. This is what we mean by *tacit knowledge*: evidence that is readily available to every person looking at the KM (**ontology**).

**Explicit knowledge the KM refers to:** the network does not explain, given an arbitrary variable, what the domain expert expects about this variable given an arbitrary combination of outcomes of its parents. Consider the former example. The parents of the variable *Supporters* are the variables *Recent performance* and *Ranking*. Suppose that the domain expert knows that the *ranking* is *top* and the *Recent performance* of the team has been *positive*, what are his beliefs about the enthusiasm of the supporters? To specify his knowledge about the variable *Supporters*, he has to answer questions of this kind for each possible combination of outcomes of the two variables *Recent performance* and *Ranking* (there are actually 9 possible combinations). This task is performed through the arguments described in the next section. The result is a list of arguments expressing in a textual way the beliefs of the domain expert about the variable he is considering.

### 3.1.2 Arguments

As explained in the previous section, there is a kind of knowledge that is not explicitly contained in the KM, but which the map points to. In the CN, this part of the expert knowledge is coded through CSs and therefore is not easily accessible for persons without particular technical skills. On the other hand, it can be very difficult for the expert to specify his beliefs directly through CSs. Therefore, to satisfy the **reasoning** and the **ontology** conditions, we have introduced an intermediate state in the elicitation procedure. This intermediate state consists in specifying the expert knowledge that is not explicitly represented in the qualitative part of the network through a list of *arguments*. By argument we intend a statement of the domain expert that is based on practical reasoning (Benson et al. [11], Browne et al. [13]) or on logical deduction, in the light of his knowledge. In this section, we provide a theoretical description of the construction and the structure of these arguments. Let us introduce the discussion by the following example.

**Example 3.5 (Football Example)** *According to Figure 5, the variables* Recent performance *and* Ranking *have a direct influence on the enthusiasm of the* Supporters*. Nevertheless, the figure does not provide any further information about this influence. Therefore, we should still ask the domain expert about that, by means of questions like: what is expected about the enthusiasm of the* Supporter*, assuming a* top Ranking *and a* positive Recent performance*?*

This task can be performed by reporting a list of arguments expressing in a textual way the beliefs of the domain expert about the considered variable. Each variable of the KM points to a body of knowledge expressed in textual way through arguments. Note that, for each variable, *all* the possible scenarios (**reasoning**, **completeness**), corresponding to all the possible joint states of the variables that directly affect the variable under consideration, should be considered. E.g., in the above example the nine joint states of *Recent performance* and *Ranking* should be considered. Expert's answers to questions of this kind are usually based on arguments, whose truth cannot be demonstrated formally. Consider for instance the following argument, based on technical considerations and on historical cases.

> *The Bellinzona football team will probably win playing at home, because playing at home the support is stronger than playing elsewhere, since for the supporters it is easier and cheaper to go to the stadium. This unless the supporters are disappointed because of the recent bad performance of the team.*

According to Toulmin [52], an argument is a statement composed of at least three parts. The first is the *data*, that is, the evidence on which the argument is based. Then the *claim*, which means the conclusion that follows from the data. Finally, the *warrant*, which is the explanation of the relationship between data and claim. Formally, the basic structure of an argument is the following:

> Given *data*, since *warrant*, therefore *claim*.

In our example we have:

> Given that the team plays at home (*Data*), since support playing at home is usually stronger than elsewhere (*Warrant*), it will win (*Claim*).

There are also other components of an argument that can be used if needed [52]. Namely: (i) the *backing* is a further explanation of the warrant, in the form of a statement or of an entire argument. A backing is used when the audience should be convinced of the validity of the warrant, or if for some reason the warrant should be better detailed. (ii) The *qualifier* is a single word or an expression that is used to quantify the strength of the claim. A qualifier is typically used for example if the domain expert is uncertain about his argument. (iii) A *rebuttal* is used to take explicitly into consideration a special case or a scenario in which the given argument does not hold. In our example we have:

> The Bellinzona football team playing at home (*Data*), will $\overbrace{\text{probably}}^{\text{Qualifier}}$ win (*Claim*), because playing at home the support is stronger than playing elsewhere (*Warrant*), since for the supporters it is easier and cheaper to go to the stadium

(*Backing*). This unless the supporters are disappointed because of the recent bad performance of the team (*Rebuttal*).

Often, arguments consist only in data and claim. In these cases the warrant is considered to be tacit knowledge. Consider the following example:

If the recent performance of the team is positive (*Data*), supporters are enthusiastic (*Claim*).

In this case, it is self-evident that supporters are enthusiastic if the performance of the team is positive. Although the warrant is not explicitly stated in the statement, we consider this an argument as well. The domain expert should avoid stating arguments without warrant, if he is not aware of the knowledge he shares with other users or developers of the KBS. The knowledge engineer, who is usually aware of the skills of the different persons involved in the specification of the model, can help the domain expert to decide which argument requires an explicit warrant (**ontology**). In our specific problem, arguments should be specified for each variable in the KM. The data of the arguments are a combination of possible values of the parents of the variable under consideration. The claim expresses the (uncertain) beliefs of the domain expert, given the data, with respect to one or more possible outcomes of the variable. There are basically three possible types of claims about the possible outcomes, that we now describe in detail.

- Classificatory judgements: given a combination of values of the parents of the variable, the domain expert expresses his beliefs about a particular outcome of the variable. An example could be *if the team has a low ranking and has had a negative recent performance, the supporters will probably be disappointed*.
- Comparative judgements: given a combination of values of the parents of the variable, the domain expert compares two possible outcomes of the variable. An example could be *if the relative ranking of the two teams is strongly in favour of the opponent, then a loss is much more probable than a win.*
- Conditional comparative judgements: the domain expert compares his beliefs about a given outcome of the variable with respect to two different combinations of values of the parents of the variable. An example could be *if the support is strong, the probability of winning is higher than if the support is weak.*

The distinction between the three types of judgement is particularly important in the response assessment, as illustrated in Section 3.2. These judgements are particular types of argument. Throughout the chapter, we will refer to them interchangeably using the generic word *arguments* or using their technical name. The procedure that should be followed to specify all the required arguments consists in considering each variable of the network and in specifying, for each combination of its parents, the beliefs of the expert about a given combination using one or more arguments. The arguments can be more or less informative, depending on the level of expertise about the given combination and, in case of an expert consisting of many individuals, on the level of agreement of the different individuals. In extreme cases, if the expert has no beliefs about the given combination, or if the disagreement between different individuals is such that no shared argument can be specified, the corresponding arguments can be completely uninformative (*I/we don't know*).

After having specified all the beliefs of the experts about the variables in the network through textual arguments, the belief assessment procedure is finished. The network structure can be indeed passed directly to a computer (**transferability**), while the beliefs of the experts, expressed in textual way, should be translated or expressed through CSs. The latter task is achieved by the response assessment procedure described in Section 3.2.

## 3.2 Response Assessment for the Domain Knowledge

After the belief assessment, the knowledge of the domain expert is represented through a KM and a list of arguments. The response assessment represents this body of knowledge through a suitable mathematical formalism. The arguments specified by the expert are characterized by imprecision and uncertainty, we need therefore a mathematical formalism able to model these features. As already emphasized in the Introduction, we consider *credal sets* (CSs) [42] as a very suitable and general tool for the representation of uncertainty in KBSs. Actually, CSs have many features that are strongly desirable in order to develop KBSs:

- They allow to model imprecision, being therefore suited to model in a natural way ignorance, vagueness and contrast (Walley [57, 58], Cozman [17]).
- They allow to easily translate arguments into a mathematical model (Walley [57, 58], Cozman [17]).
- They are very general. Actually, the majority of the models used nowadays for the representation of uncertainty can be considered as special cases of CSs (Walley [58], Destercke et al. [23, 24]). Examples of such models are belief functions, possibility measures, probability intervals and Bayesian probabilities.

A CS is a closed convex set of probability mass functions. The idea of modelling subjective uncertain beliefs by convex sets of probability mass functions comes from Levi [42], who has formalized the notion of CS. For the sake of simplicity, we consider in this chapter only *finitely generated* CSs [16], i.e., obtained as the convex hull of a finite number of probability mass functions for a given variable. Geometrically, a finitely generated CS is a *polytope*. CSs of this kind contain infinitely many probability measures, but only a finite number of extreme mass functions, corresponding to the *vertices* of the corresponding polytope. In general, given an arbitrary variable $X$, we denote with $K(X)$ the CS associated to it and with $\mathrm{ext}[K(X)]$ the finite set of vertices of the CS. The following simple example, taken from Walley [57], explains how CSs can be obtained from arguments in practice.

**Example 3.6 (Walley's football example)** *Consider a football match. The possible results are* win *($W$),* draw *($D$) and* loss *($L$). Suppose that the domain expert has specified the following arguments:* *win* is improbable, *win* is at least as probable as *draw*, *draw* is at least as probable as *loss*, I would bet against *loss* only at odds that are no more than 4 to 1. *The CS corresponding to the arguments is the set of all the probability measures*

$$\mathbf{p} = \begin{pmatrix} P(W) \\ P(D) \\ P(L) \end{pmatrix}$$

*that are compatible with the arguments. The above arguments can be directly translated into (linear) inequalities concerning the components of* $\mathbf{p}$*. In particular we have,*

- Win *is improbable. Applying the conservative translation of the term* improbable *given by Walley [57] we have:* $P(W) \leq 0.5$.
- Win *is at least as probable as* Draw*:* $P(W) \geq P(D)$.
- Draw *is at least as probable as* Loss*:* $P(D) \geq P(L)$.
- *I would bet against* Loss *only at odds that are no more than 4 to 1:* $4 \cdot P(L) \geq P(W) + P(D)$*, and hence* $P(L) \geq 0.2$.

*Furthermore, the components of the vector* $\mathbf{p}$ *should satisfy the following consistency conditions,* $P(W)+P(D)+P(L) = 1$, $P(W) \geq 0$, $P(D) \geq 0$, $P(L) \geq 0$. *All the inequalities specified above can be collected in the following system:*

$$\begin{cases} P(W) \leq 0.5 \\ P(W) \geq P(D) \\ P(D) \geq P(L) \\ P(L) \geq 0.2 \\ P(W) + P(D) + P(L) = 1 \\ P(W) \geq 0 \\ P(D) \geq 0 \\ P(L) \geq 0. \end{cases}$$

*The above linear system of inequalities is already a representation of the CS associated to the arguments of the expert. Yet, another possible way to represent it is to solve the system and represent the solution. Note that the solution of such a system, if the system is not impossible, is a polytope defined by a finite number of vertices (i.e., a finitely generated CS). A finitely generated CS can be represented by enumerating explicitly its vertices. There exist algorithms that, given a linear system of inequalities, calculate the vertices of the corresponding polytope. In this chapter, we have used a Matlab implementation of the* Reverse Search Vertex Enumeration Algorithm *(Avis and Fukuda [8], Avis [7]), lrs in short, based on Kovačec and Ribeiro [41]. Solving the system above we find the following four vertices.*[14]

| $P(W)$ | $P(D)$ | $P(L)$ |
|---|---|---|
| *0.5* | *0.3* | *0.2* |
| *0.5* | *0.25* | *0.25* |
| *0.4* | *0.4* | *0.2* |
| $0.\overline{3}$ | $0.\overline{3}$ | $0.\overline{3}$ |

For a variable with three possible states, like for example the variable considered in the example above, a particularly effective geometric representation is based on the use of *barycentric coordinates* (Walley [57]). In the barycentric representation, each vertex of a triangle of height one is labeled with one possible outcome of the variable considered and a probability vector $(P(W), P(D), P(L))$ is represented through a point in the interior of the triangle, whose distance from each edge opposite to a vertex labeled with a particular outcome corresponds to the probability assigned to the given outcome, as shown in Figure 6(a). Barycentric coordinates are based on the following property of equilateral triangles:

[14] We denote with $0.\overline{3}$ the number $\frac{1}{3}$.

the sum of the distances of a point in the interior of the triangle from the edges of the triangle itself is equal to the height of the triangle. In the case of barycentric coordinates, the height of the triangle is 1 and the distances of a point in the interior of the triangle from the edges correspond to the probabilities $(P(W), P(D), P(L))$ assigned by a given mass function; thanks to the above property, we have $P(W) + P(D) + P(L) = 1$. The CS obtained in Example 3.6 is represented in Figure 6(b).

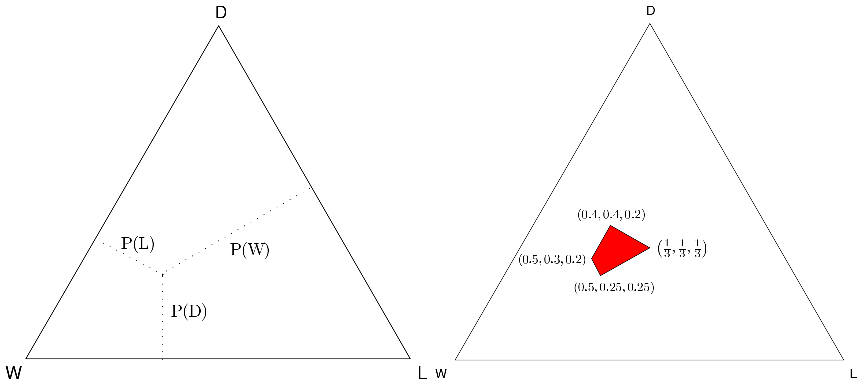


(a) (b)

Figure 6: In barycentric coordinates, a probability vector $(P(W), P(D), P(L))$ is represented as a point in the interior of an equilateral triangle of height 1, as shown in Figure 6(a). To represent a CS, it is sufficient to represent the polytope defined by the vertices of the CS. For example, The CS obtained in Example 3.6 is represented in Figure 6(b).

Very often in applications, given a variable $X$ and a CS $K(X)$, we are interested only in the *lower and upper probabilities* of an outcome. Denoting with $x$ an arbitrary outcome of the variable $X$, these are defined by

$$\underline{P}(x) = \min_{\mathbf{p} \in \mathcal{K}(\mathcal{X})} P(x), \quad \overline{P}(x) = \max_{\mathbf{p} \in \mathcal{K}(\mathcal{X})} P(x).$$

Because a CS is closed and convex, the lower and upper probabilities are attained at the vertices; we have

$$\underline{P}(x) = \min_{\mathbf{p} \in \mathrm{ext}[K(X)]} P(x), \quad \overline{P}(x) = \max_{\mathbf{p} \in \mathrm{ext}[K(X)]} P(x).$$

For example, the lower and upper probabilities in Example 3.6 are given by $\underline{P}(W) = 0.\overline{3}$, $\overline{P}(W) = 0.5$, $\underline{P}(D) = 0.25$, $\overline{P}(D) = 0.4$, $\underline{P}(L) = 0.2$ and $\overline{P}(L) = 0.\overline{3}$.

In order to apply the approach in Example 3.6 in the general case, we need a procedure to convert the domain expert's arguments into linear constraints on probabilities. Walley [57] has identified several types of arguments that can be expressed by a domain expert, in particular, he has distinguished between *classificatory probability judgements* and *comparative probability judgements*. The translation of the judgements into constraints depends on the type of judgement expressed. Some judgements can be directly translated into a constraint in a unique way, while other judgements require some probability scale for the translation and are therefore subject to arbitrariness. In the following, we illustrate the different types of probability judgements and their translation and representation through linear inequalities. In particular, in Section 3.2.1, we discuss classificatory probability judgements, while in Section 3.2.2 we discuss comparative probability judgements.

### 3.2.1 Classificatory Probability Judgements

Classificatory probability judgements refer to the (conditional) probability of a particular outcome of the variable. An example could be the statement used in the previous football example, *win is probable*. In this kind of judgements often the degree of belief in an outcome is quantified through a verbal qualifier like *probable*, *unlikely*, etc. Several authors,[15] like for example Walley [58] and Renooij [49, 50], have proposed verbal-numerical scales for the translation of such type of judgements into linear inequalities. The verbal-numerical scale proposed by Walley in [58] is represented in Figure 7. According to this scale, the above judgement should be translated into the inequality $P(W) \geq 0.5$. The verbal-numerical scale proposed by Renooij et al. in [49, 50] is illustrated in Figure 8. This scale is not a proper verbal-numerical scale, because the words are not directly translated into probabilities. An expert using this scale to specify precise probabilities should select a single probability using the words as guide, but not as anchor points. For example, according to this scale, the above judgement should be translated into the equality $P(W) = \alpha$, where $\alpha \in [0.75, 0.85]$. The scale can be used as a true verbal-numerical scale only when specifying imprecise probabilities. In this case, as suggested by Cozman [17], it is sufficient to assign to a given word the whole interval of possible values. For example, *win is probable* could be translated as $0.75 \leq P(W) \leq 0.85$. The translation of classificatory probability judgements is strictly dependent on the verbal-numerical scale the domain expert refers to and is therefore questionable.

There are also cases of classificatory probability judgements where the translation is unique. In particular, it is possible to assign explicitly a probability interval to a given outcome. For example, the judgement *the probability of win is between 0.3 and 0.5* is translated as $0.3 \leq P(W) \leq 0.5$. A special case of this type of judgements are precise probability assignments, like for example *the probability of win is 0.85* that can be translated as $P(W) = 0.85$. Precise assignments include also judgements of impossibility and certainty, like *win is impossible*, $P(W) = 0$, and *win is sure*, $P(W) = 1$.

**Example 3.7** *Consider the football match between Bellinzona and Basel, with possible results* win *($W$),* draw *($D$) and* loss *($L$). Suppose that the domain expert has specified the following arguments,* win is improbable, loss is probable. *To translate these judgements, we have to refer to a verbal-numerical probability scale. We consider the scale of Renooij [50] in Figure 8 and we follow Cozman's [17] approach.* Win is improbable, *is translated as $P(W) \leq 0.15$, while* loss is probable *is translated as $0.75 \leq P(L) \leq 0.85$. These inequalities, combined with the consistency conditions $P(W) + P(D) + P(L) = 1$, $P(W) \geq 0$, $P(D) \geq 0$, $P(L) \geq 0$, define a CS with the following vertices.*

[15] See Jenkinson [37] and O'Hagan et al. [46] for a review of elicitation methods for precise probability assessments.

| $P(W)$ | $P(D)$ | $P(L)$ |
|---|---|---|
| *0.15* | *0.10* | *0.75* |
| *0.15* | *0* | *0.85* |
| *0* | *0.15* | *0.85* |
| *0* | *0.25* | *0.75* |

*Calculating lower and upper probabilities for each outcome, we have that* $P(L) \in [0.75; 0.85]$, $P(D) \in [0; 0.25]$ *and* $P(W) \in [0; 0.15]$.

During several elicitation sessions with domain experts, we have noticed that it is not easy for them to express classificatory judgements about a variable. This happens because of several effects. First of all, we have experienced biases due to the problem of *availability* (Kahneman et al. [39]), i.e., the tendency of the domain expert to assign higher probability to an outcome that he can display clearly in mind, for example because it has occurred recently. Furthermore, we have observed that it is sometimes difficult for the domain expert to keep in mind all the possible outcomes of a variable and consequently to compare one particular outcome with all the other outcomes together. Actually, a classificatory probability judgement is the same as a comparative probability judgement between an outcome and its complement, i.e., all the other possible outcomes. This effect is illustrated in the above example, where the domain expert has specified classificatory judgements for the two outcomes *win* and *loss*, but has neglected the outcome *draw*. We argue that for the domain expert it is easier to specify comparative probability judgements on outcomes, than to specify classificatory probability judgements.

| Verbal | Numerical |
|---|---|
| $x$ is extremely probable | $P(x) \geq 0.98$ |
| $x$ has very high probability | $P(x) \geq 0.9$ |
| $x$ is highly probable | $P(x) \geq 0.85$ |
| $x$ is very probable | $P(x) \geq 0.75$ |
| $x$ has a very good chance | $P(x) \geq 0.65$ |
| $x$ is quite probable | $P(x) \geq 0.6$ |
| $x$ is probable (likely) | $P(x) \geq 0.5$ |
| $x$ has a good chance | $0.4 \leq P(x) \leq 0.85$ |
| The probability of $x$ is about $\alpha$ | $\alpha - 0.1 \leq P(x) \leq \alpha + 0.1$ |
| $x$ is improbable (unlikely) | $P(x) \leq 0.5$ |
| $x$ is somewhat unlikely | $P(x) \leq 0.4$ |
| $x$ is very unlikely | $P(x) \leq 0.25$ |
| $x$ has little chance | $P(x) \leq 0.2$ |
| $x$ is highly improbable | $P(x) \leq 0.15$ |
| $x$ is has very low probability | $P(x) \leq 0.1$ |
| $x$ is extremely unlikely | $P(x) \leq 0.02$ |

Figure 7: The verbal-numerical scale proposed by Walley in [58].

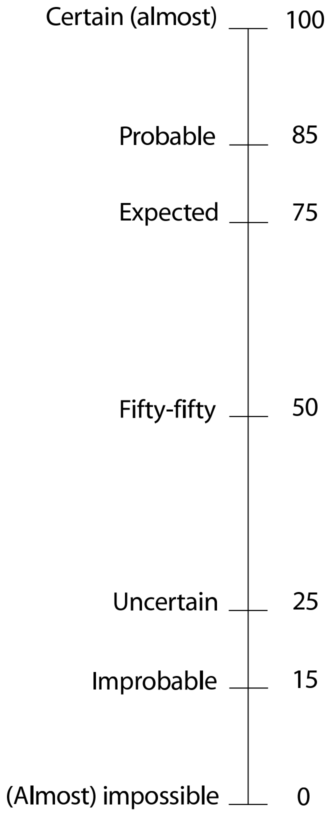


Figure 8: The verbal-numerical scale proposed by Renooij et al. in [50, 49].

#### 3.2.2 Comparative Probability Judgements

Comparative probability judgements are based on the comparison of the (conditional) probabilities for different outcomes of the variables. An example of such a judgement is *win is at least as probable as draw*. This type of judgement can be directly translated as $P(W) \geq P(D)$. We can also compare conditional probabilities. For example, we can compare the conditional probabilities of two different outcomes with the same conditioning events event like *playing on the home field, win is at least as probable as loss*. We can also compare the probabilities of the same outcome with different conditioning events, like for example *win playing at home is at least as probable as winning on other football fields*. A comparison is possible also between probabilities of different outcomes with different conditioning events, like *win playing at home is at least as probable as draw playing on other football fields*. On the basis of our experience, this latter type of judgement is not used very often by the expert.

Finally, there is a particular type of comparative probability judgement that occurs very rarely in practice: it consists of comparisons between outcomes of different variables, like for example *win playing at home is at least as probable as having a strong support when the supporters are enthusiastic playing at home*. Although there are particular special cases in which such judgements are meaningful, we do not consider them in the present discussion.

Comparative probability judgements should not necessarily compare probabilities of outcomes, they can also compare in general probabilities of events. Yet, in practice, domain experts often prefer to consider only single outcomes. This happens because considering events in general, the domain expert can incur the same error displayed above for classificatory probability judgements, i.e., neglecting some outcome. For example, if the expert says that *win is much more probable than not winning*, we cannot be sure that he consider the event *not winning* as the set containing draw and loss, maybe he is focusing on the outcome loss and he is neglecting the outcome draw. For this reason, and for the sake of simplicity, in the rest of the section we focus on the specification of inequalities representing comparative probability judgements on outcomes.

**Example 3.8 (Assessing a conditional CS)** *Consider the simplified football network displayed below.*

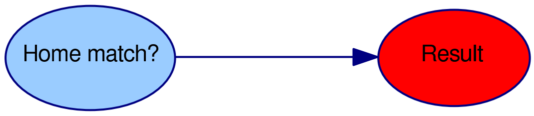


*The possible values of* Home match *are* $H$ *if the match is played on the home football field and* $\neg H$ *if not. The possible values of* Result *are* Win *(*$W$*),* Draw *(*$D$*) and* Loss *(*$L$*). We should specify the set of all the vectors*

$$\mathbf{p} = \begin{pmatrix} P(W \mid H) \\ P(D \mid H) \\ P(L \mid H) \\ P(W \mid \neg H) \\ P(D \mid \neg H) \\ P(L \mid \neg H) \end{pmatrix}$$

*that are consistent with the arguments reported by the domain expert. This set is denoted as* $L(\text{Result} \mid \text{Home match})$*. The set* $L(\text{Result} \mid \text{Home match})$ *is not itself a CS, because a CS is a convex set of probability measures, while a point in this set describes two conditional probability measures at the same time. Loosely speaking, the set* $L(\text{Result} \mid \text{Home match})$ *defines two CSs. Suppose that* $L(\text{Result} \mid \text{Home match})$ *should be constructed according to the following arguments specified by the domain expert:* If the match is played at home, win is more probable than loss, draw is more probable than loss and win is probable. If the match is not played at home, win is more probable than loss, draw is more probable than loss and draw is probable. *These arguments should be translated into a system of linear inequalities. We have two classificatory probability judgements and four comparative probability judgements. The comparative judgements can be translated directly, while the two classificatory judgements require some verbal-numerical probability scale for translating the term* probable *into a number. In this example, we interpret the term* probable *according to Walley [57]: win is probable if* $P(W) \geq 0.5$*. The above arguments can therefore be translated as follows:*

$$\begin{cases} P(W \mid H) \geq P(L \mid H) \\ P(D \mid H) \geq P(L \mid H) \\ P(W \mid H) \geq 0.5 \\ P(W \mid \neg H) \geq P(L \mid \neg H) \\ P(D \mid \neg H) \geq P(L \mid \neg H) \\ P(D \mid \neg H) \geq 0.5. \end{cases}$$

*To the above inequalities, we should add the following consistency constraints:*

$$\begin{cases} P(W \mid H) + P(D \mid H) + P(L \mid H) = 1 \\ P(W \mid H) \geq 0, \quad P(D \mid H) \geq 0, \quad P(L \mid H) \geq 0 \\ P(W \mid \neg H) + P(D \mid \neg H) + P(L \mid \neg H) = 1 \\ P(W \mid \neg H) \geq 0, \quad P(D \mid \neg H) \geq 0, \quad P(L \mid \neg H) \geq 0. \end{cases}$$

*We are now ready to solve the system of inequalities corresponding to the specified arguments. The set* $L(\text{Result} \mid \text{Home match})$ *is the convex set defined by the following nine vertices.*

| $P(W \mid H)$ | $P(D \mid H)$ | $P(L \mid H)$ | $P(W \mid \neg H)$ | $P(D \mid \neg H)$ | $P(L \mid \neg H)$ |
|---|---|---|---|---|---|
| 0.5 | 0.25 | 0.25 | 0.25 | 0.5 | 0.25 |
| 0.5 | 0.5 | 0 | 0.25 | 0.5 | 0.25 |
| 1 | 0 | 0 | 0.25 | 0.5 | 0.25 |
| 0.5 | 0.25 | 0.25 | 0.5 | 0.5 | 0 |
| 0.5 | 0.5 | 0 | 0.5 | 0.5 | 0 |
| 1 | 0 | 0 | 0.5 | 0.5 | 0 |
| 0.5 | 0.25 | 0.25 | 0 | 1 | 0 |
| 0.5 | 0.5 | 0 | 0 | 1 | 0 |
| 1 | 0 | 0 | 0 | 1 | 0 |

*In this case, there is also an easier way to obtain the vertices of* $L(\text{Result} \mid \text{Home match})$. *The system is the union of two disjoint systems. This is due to the absence, in the list of arguments of the expert, of judgements comparing probabilities with different conditioning events. The two systems can be solved separately, obtaining thus a CS for the probabilities conditioned on* $H$, *denote it by* $K(\text{Result} \mid H)$, *and a CS for the probabilities conditioned on* $\neg H$, *denote it by* $K(\text{Result} \mid \neg H)$. *The vertices of the set* $L(\text{Result} \mid \text{Home match})$ *can then be obtained combining the vertices of the two separate CSs in all the possible ways. The CS* $K(\text{Result} \mid H)$ *can be obtained solving the system consisting of the equations of the previous system containing the probabilities conditioned on* $H$, *i.e.,* $P(W \mid H)$, $P(D \mid H)$ *and* $P(L \mid H)$. *We obtain that* $K(\text{Result} \mid H)$ *is defined by the following three vertices.*

| $P(W \mid H)$ | $P(D \mid H)$ | $P(L \mid H)$ |
|---|---|---|
| 0.5 | 0.25 | 0.25 |
| 0.5 | 0.5 | 0 |
| 1 | 0 | 0 |

*Similarly,* $K(\text{Result} \mid \neg H)$ *is defined by the following three vertices.*

| $P(W \mid \neg H)$ | $P(D \mid \neg H)$ | $P(L \mid \neg H)$ |
|---|---|---|
| *0.25* | *0.5* | *0.25* |
| *0.5* | *0.5* | *0* |
| *0* | *1* | *0* |

More generally speaking, consider a variable $X$ and denote by $\mathrm{Pa}(X)$ its parents and with $\{x_1, \dots x_n\}$ the joint values of the parents. Denote by $K(X \mid Pa(X))$ the set defined by all the possible combination of vertices of the CSs

$$K(X \mid \mathrm{Pa}(X) = x_1), \dots, K(X \mid \mathrm{Pa}(X) = x_n).$$

If the vertices of $L(X \mid \mathrm{Pa}(X))$ coincide with $K(X \mid \mathrm{Pa}(X))$, we say that these CSs are *separately specified*. Otherwise, we say that the CSs defined jointly by $L(X \mid \mathrm{Pa}(X))$ are *extensively specified*.

If the expert specifies only classificatory probability judgements and comparative judgements comparing probabilities of different outcomes but with the same conditioning events, as is the case above, then the CSs corresponding to the arguments are automatically separately specified because the corresponding system of inequalities decomposes into subsystems, one for each possible combination of the values of the parents. If the expert specifies also comparative probability judgements for probabilities with different conditioning events, then the corresponding system of inequalities cannot be further decomposed. To obtain a CS that is not separately specified, an expert should specify arguments comparing probabilities with different conditioning events that are not already implied by the classificatory probability judgements and the comparative probability judgements comparing probabilities with same conditioning events, as illustrated in the following examples.

**Example 3.9 (Redundant constraint)** *Consider the situation of the previous example. Suppose that the domain expert adds the following argument to the arguments listed previously:* the probability of win if we play at home is at least as large as the probability of win if we do not play at home. *This argument compares two probabilities with different conditioning events. It corresponds to the inequality:*

$$P(W \mid H) \geq P(W \mid \neg H).$$

*To find the set* $L(\text{Result} \mid \text{Home match})$*, we consider the system used previously to compute* $L(\text{Result} \mid \text{Home match})$ *using in addition the above inequality. In this way, the system cannot be decomposed into two separate systems. But the solution we find solving it is exactly the same* $L(\text{Result} \mid \text{Home match})$ *with nine vertices as in the previous case. It follows that the CSs are however* ***separately specified****, although the domain expert has specified a comparative probability judgement comparing probabilities with different conditioning events. This happens because the additional condition stated above is already (implicitly) satisfied by the arguments specified in the previous example and could be omitted.*

**Example 3.10 (Extensively specified CSs)** *Consider again the situation of the previous example with the following additional argument of the domain expert:* the probability of winning if we play at home is at least three times as large as the probability of winning if we do not play at home. *This argument corresponds to the inequality*

$$P(W \mid H) \geq 3 \cdot P(W \mid \neg H).$$

*To find the set* $L(\text{Result}\,|\,\text{Home match})$*, we consider once more the system used previously to compute* $L(\text{Result}\,|\,\text{Home match})$ *using in addition the inequality above. In this case the system cannot be decomposed into two separate systems. The solution of this system is a polytope defined by twenty vertices, and therefore this set is not the same that we have obtained in the previous example. In this case, the additional argument is not redundant and therefore the CSs corresponding to those arguments are* ***extensively specified****.*

A CN whose quantitative part consists only of separately specified CSs is said a *separately specified* CN. If the quantitative part of a CN contains both separately and extensively specified CSs, then the CN is said extensively specified. Extensively specified networks are a special of *non-separately specified CNs*, which are networks where there are relationships between the specifications of the conditional CSs. The existing algorithms for the updating of CNs have been developed only for separately specified CNs. For non-separately specified CNs, there is a procedure that has been recently proposed by Antonucci and Zaffalon [6], that transforms an arbitrary non-separately specified CN into an equivalent separately specified CN allowing thus to make inference also in this case with the same algorithms employed for CNs with separately specified CSs (see [4] for an example of this approach). It follows that, when a CN is extensively specified, it should be transformed according to [6] in order to update it with available updating algorithms.

In the above examples, comparative probability judgements can be immediately expressed through inequalities, but this is not always the case. For example, we could say that *win is much more probable than draw*. In this case, it is necessary to refer to some verbal-numerical scale and also the translation of comparative judgements can be subject to arbitrariness. In particular, we have noticed, during our practical sessions with the domain expert queried for the development of a KBS for airspace surveillance [1], that the use of a verbal-probability scale is not ideal in a multilingual contest like for example the Swiss one. Actually, in Switzerland, it occurs very often to have meetings with people having three or more different mother tongues and the use of verbal scales is difficult. Also the use of a neutral language, like English, has proved not to be effective, because of the different interpretations assigned to the same words by different persons. These considerations have led us to conceive a procedure for specifying comparative probability judgements through inequalities without referring to a verbal-numerical scale. We explain our procedure through the following examples.

**Example 3.11 (Football example)** *Consider the variable* Support *in the KM in Figure 5. The variable has three possible values:* strong *($S$),* medium *($M$) and* weak *($K$). The parents of the variable* Support *are the variable* Supporters*, with as possible values* enthusiastic *($E$),* normal *($N$) and* disappointed *($I$), and the variable* Home match *with as possible values* yes *($H$) and* no *($\neg H$). The possible combinations of the parents are* $(H,E)$, $(H,N)$, $(H,I)$, $(\neg H,E)$, $(\neg H,N)$, $(\neg H,I)$. *Suppose that the domain expert has expressed the following arguments about this variable.*

> *Playing at home with supporters that are not disappointed, I expect the support to be strong, in particular if the supporters are enthusiastic. A weak support is extremely unlikely, although not impossible. If the support is not strong, I*

> *expect much more a normal support than a weak one. If the supporters are disappointed, the probability of having a strong support at home is low. I expect in this case a medium or weak support. Not playing at home, if the supporters are enthusiastic, I expect however a strong support,* but in this case the probability of having a medium and a weak support are higher than playing at home. *If the supporters are normal, then I expect a weak or medium support; a strong support is extremely unlikely.* The probability of having a strong support in this case is lower than playing at home with normal supporters. *If the supporters are disappointed and we are not playing at home, then I expect the support to be weak; if it is not weak then I expect it to be medium, but a strong support is almost impossible.* In general, assuming the same state for the other variables, the probability of having a strong support is higher playing at home than not playing at home. At home, the probability of having a strong support is higher with enthusiastic supporters than with normal ones, and it is higher with normal supporters than with disappointed ones. The same holds when not playing at home.

*Our aim is to construct* $L(\text{Support} \,|\, \text{Home match}, \text{Supporters})$ *according to the arguments of the domain expert. The arguments above contain several statements comparing the different outcomes of the variable given a particular combination of parents (in italic) and several statements comparing the same outcome for different combinations of parents (not in italic). We see immediately that arguments of the second type can be immediately translated as follows.*

- *Not playing at home, but with enthusiastic supporters,* the probability of having a medium and a weak support are higher than playing at home.

$$\begin{cases} P(M \,|\, \neg H, E) \geq P(M \,|\, H, E) \\ P(K \,|\, \neg H, E) \geq P(K \,|\, H, E). \end{cases}$$

- *Not playing at home with normal supporters,* the probability of having a strong support is surely lower than playing at home with normal supporters.

$$P(S \,|\, \neg H, N) \leq P(S \,|\, H, N).$$

- Other things being equal (same supporters), the probability of having a strong support is higher playing at home than not playing at home.

$$\begin{cases} P(S \,|\, H, E) \geq P(S \,|\, \neg H, E) \\ P(S \,|\, H, N) \geq P(S \,|\, \neg H, N) \\ P(S \,|\, H, I) \geq P(S \,|\, \neg H, I). \end{cases}$$

- Playing at home, the probability of having a strong support is higher with enthusiastic supporters than with normal ones, and it is higher with normal supporters than with disappointed ones. The same holds not playing at home.

$$\begin{cases} P(S \mid H, E) \geq P(S \mid H, N) \\ P(S \mid H, N) \geq P(S \mid H, I) \\ P(S \mid \neg H, E) \geq P(S \mid \neg H, N) \\ P(S \mid \neg H, N) \geq P(S \mid \neg H, I). \end{cases}$$

*The rest of the arguments of the domain expert cannot be translated directly, and also the selection of a verbal-numerical scale does not help in this case. A possible specification of the inequalities describing the remaining arguments is described in the following.*

We propose a procedure for the specification of inequalities describing comparative probability judgements based on the specification of lower and upper *probability ratios* (Walley [57, chapter 4.6.2]). I.e., given two possible outcomes $A$ and $B$, the expert has to specify two numbers: a lower ratio $l$ and an upper ratio $u$ such that

$$l \leq \frac{P(A)}{P(B)} \leq u.$$

In principle, for each combination of the outcomes, $l$ and $u$ can be arbitrary positive real numbers (provided that $l \leq u$). But, to simplify and accelerate the procedure, we advise the use of a scale with only finitely many possible ratios. The choice of a discrete scale is supported by Ofir and Reddy [47] who have shown, in the case of the specification of probabilities, that using a discrete scale with seven values instead of a continuous scale does not reduce significantly the quality of the specified values. We argue that this could be the case also specifying probability ratios, although, at the best of our knowledge, nothing has been published on this topic. A possible probability ratio scale is represented in Figure 9. In this scale, the ratios are graphically depicted through a visual proportion between the two probabilities and each possible ratio is associated with a symbol. The use of this scale allows the domain expert to specify his beliefs without the need of using words, avoiding thus the problem related to the use of verbal-numerical scales. It is important to remark that the scale used for this procedure, i.e., the symbols, the number of ratios and their values can be chosen arbitrarily by the domain expert. A practical application of this scale is illustrated in the following example.

**Example 3.12 (Football example)** *Consider once more the variable* Support *and suppose that the domain expert has specified the following argument,* Playing at home with supporters that are not disappointed, I expect the support to be strong, in particular if the supporters are enthusiastic. *The argument is not a comparative judgement, and therefore cannot be translated directly. This argument is a sort of summary of the knowledge of the domain expert for the given combination of parents and should be expanded into more explicit arguments about the probabilities of the different outcomes, in order to express it through inequalities. To do this, the domain expert can use the scale in Figure 9 to express his knowledge more precisely through comparative probability judgements. Given a particular combination of parents, he should fill some cells of the first table of Figure 10, using symbols of the probability ratio scale, where the symbol specified on the left of a particular cell corresponds to the lower probability ratio between the outcome in the row (numerator of*

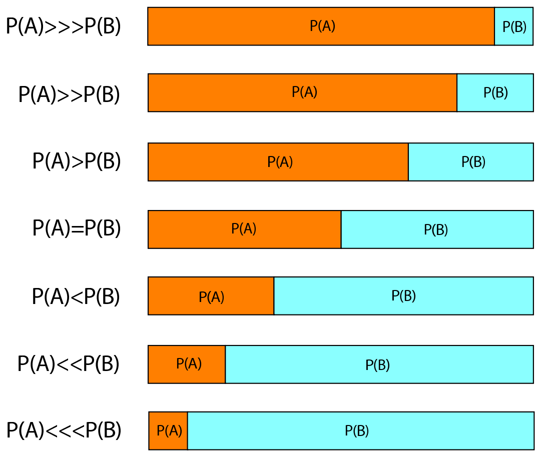


Figure 9: A possible symbolic-visual-numerical scale for the specification of probability ratios between the different outcomes of a variable. The ratios between probabilities are represented through visual proportions. the first row represents a proportion between $P(A)$ and $P(B)$ of 9 to 1 (9:1). The other proportions are 4:1, 2:1, 1:1, 1:2, 1:4 and 1:9.

*the ratio) and the outcome in the column (denominator of the ratio), the symbol specified on the right of a particular cell corresponds to the upper probability ratio between the outcome on the row (numerator of the ratio) and the outcome on the column (denominator of the ratio). The cells on the diagonal of the table are excluded because because they would be comparisons between equal probabilities. Note that, specifying ratios for an ordered combination of outcomes, ratios are (implicitly) specified also for the same combination in reverse order. Consider for example the pair* strong *and* weak *and suppose that we know that*

$$l \leq \frac{P(\text{strong})}{P(\text{weak})} \leq u,$$

*then we know automatically that*

$$\frac{1}{u} \leq \frac{P(\text{weak})}{P(\text{strong})} \leq \frac{1}{l}.$$

*Suppose that the ratios specified in the first table of Figure 10 correspond to ratios specified by the domain expert considering a situation where the supporters are enthusiastic and the match is played at home. Each filled cell in the table can be translated into a linear inequality taking into consideration the ratios in the probability ratio scale corresponding to the symbols used by the domain expert. We obtain the following inequalities,*

$$\begin{cases} P(S\,|\,H,E) \geq 4 \cdot P(M\,|\,H,E) \\ P(S\,|\,H,E) \leq 9 \cdot P(M\,|\,H,E) \\ P(S\,|\,H,E) \geq 9 \cdot P(K\,|\,H,E) \\ P(M\,|\,H,E) \geq 4 \cdot P(K\,|\,H,E) \\ P(M\,|\,H,E) \leq 9 \cdot P(K\,|\,H,E). \end{cases}$$

*Solving the above system of inequalities with the usual consistency conditions, we obtain the following vertices that define the CS* $K(\text{Support}\,|\,H,E)$.

| $P(S \mid H, E)$ | $P(M \mid H, E)$ | $P(K \mid H, E)$ |
|---|---|---|
| *0.76* | *0.19* | *0.05* |
| *0.89* | *0.10* | *0.01* |
| *0.78* | *0.20* | *0.02* |
| *0.88* | *0.10* | *0.02* |

*In the same way, the domain expert can also specify inequalities comparing probabilities of the same outcome with different conditioning events or probabilities of different outcomes with different conditioning events. Suppose for example that the domain expert needs to specify inequalities describing his beliefs about the difference between playing at home or away if the supporters are enthusiastic. He can fill a table like the second table of Figure 10. In this case the cells outside the diagonal are excluded because the domain expert is comparing probabilities of the same outcome with different conditioning events. The second table of Figure 10, combined with the probability ratio scale, defines the following inequalities, that have been already specified in Example 3.11:*

$$\begin{cases} P(S \mid H, E) \geq P(S \mid \neg H, E) \\ P(M \mid H, E) \leq P(M \mid \neg H, E) \\ P(K \mid H, E) \leq P(K \mid \neg H, E). \end{cases}$$

| | Strong | | Medium | | Weak | |
|---|---|---|---|---|---|---|
| Strong | | | >> | >>> | >>> | |
| Medium | | | | | >> | >>> |
| Weak | | | | | | |

Home match? Yes

| Home match? No | Strong | | Medium | | Weak | |
|---|---|---|---|---|---|---|
| Strong | = | | | | | |
| Medium | | | | = | | |
| Weak | | | | | | = |

Figure 10: The tables of Example 3.12.

So far, we have illustrated the construction of the CSs associated with a single variable of the domain knowledge. Consider now a whole KM constructed during the belief assessment procedure for the domain knowledge. Because we exclude judgements comparing probabilities of different variables, the specification of the CSs associated with a variable is independent of the specification of the CSs associated with the other variables of the network. This procedure for the specification of the conditional CSs for a single variable should then be simply iterated for all the variables in the KM. This ends the work of the domain expert.

## 3.3 Belief Assessment for the Observational Process

The domain knowledge describes the (theoretical) knowledge of the domain expert, but does not describe how information about the (theoretical) variables considered by the domain expert is gathered in practice. Roughly speaking, the collection of domain knowledge could be completely unrelated to the process in which the evidence is collected. This is shown in the following example.

**Example 3.13 (Sources of information)** *In Section 3.1 and 3.2 we have reported the specification of the domain knowledge by a hypothetical domain expert, regarding the football match of the next week between Bellinzona and Basel. See the network in Figure 5. Suppose now that a supporter of the Bellinzona football team aims at using this domain knowledge to predict the result of the match in the light of the evidence that he has collected about the variables in the domain knowledge. Suppose that he describes his sources of information about the different variables as follows.*

- Fitness*: the source of information are three persons that follow the training of the Bellinzona football team. Their opinion about the fitness state of the team is usually reliable. The three persons do not know one another.*
- Fitness of the opponent*: the source of information is a newspaper supporting, in a moderate manner, the Basel football team.*
- Relative ranking, Home match *and* Ranking*: the source of information is a newspaper reporting the championship ranking and the matches of the next week.*
- Supporters, Support *and* Relative fitness*: no source of information about these variables.*
- Recent performance*: the source of information is a person that is neither supporter of Bellinzona, nor a supporter of Basel. He is a person that is well informed about the Swiss football championship in general. His opinion about the recent performance of the different football teams is usually reliable and objective.*

*How can we link these sources of information with the existing domain knowledge, in order to be able to issue a prediction?*

In general, we call *source of information* about a variable a person, sensor, or any entity producing an evidence that is related in some probabilistic way to the underlying (unknown) value of the variable. The main aim of the specification of the observational process is exactly to describe how the available *sources of information* produce information about the latent variables of the domain knowledge. The observational process should be specified by an expert, called *observational process expert*. Often, in practice, the observational process expert is different from the domain expert. The same domain knowledge can be used under different practical circumstances specifying different observational processes. To specify the observational process, the procedure that should be followed by the observational process expert is the same used by the domain expert for the specification of the domain knowledge. In the *belief assessment*, he should complete the KM defined by the domain expert with discrete variables modelling the sources of information for each variable in the domain knowledge. Furthermore, he should describe, through arguments, the link between the conceptual variables in the domain knowledge and their sources of information. In the *response assessment*, the arguments describing the sources of information should be translated into inequalities and used to construct CSs for each variable modelling a source of information.

The first task to be performed by the observational process expert in the belief assessment is to complete the KM assessed by the domain expert with categorical variables modelling the sources of information and to connect them with the variables of the domain knowledge. In this tutorial, we model the observational process through a *reflective model* [12], i.e., a model in which the arrows go from the latent to the manifest level. See Borsboom et al. [12] for further details about the (philosophical) assumptions underlying this particular way of modelling. This is done variable by variable, for each variable in the domain knowledge.

**Example 3.14 (Single source of information)** *Consider the (latent) variable* Fitness of the opponent *in the football example in Figure 5. From Example 3.13, we know that the source of information about this variable consists of the fitness state reported by a newspaper supporting, in a moderate manner, the Basel football team. This situation can be modelled by the network in Figure 11(a), where the values reported near the nodes are the possible outcomes of the variables. The possible outcomes of the (manifest) variable modelling the source of information are the same of the underlying (latent) variable with in addition the value* missing*. The outcome* missing *models the situation in which the source of information does not provide any evidence about the variable of interest, for example because the newspaper does not report any information about the fitness of the Basel football team, or even because the newspaper has not been received by the user of the network, that in this case is the observational process expert himself. If the possible outcomes of a manifest variable are the same of an underlying latent variable with in addition the value* missing*, then the former variable is called an* observation *or* measurement *of the latter.*[16]

In Example 3.14 we have modelled a situation in which only one source of information is available. Often, in practice, there are several sources of information for the same variable. In practice, the evidence provided by the different sources of information can be contrasting or even contradictory. The problem of assessing the evidence coming from many sources of information is called *information fusion*. Our formalism allows to address the problem of information fusion in a flexible way, allowing to model contrast and contradiction through imprecise probabilities. This issue is further discussed in Section 4.

**Example 3.15 (Independent sources of information)** *Consider the variable* Fitness *in Figure 5. According to the observational process expert in Example 3.13, the sources of information about this variable consists the opinions of three well informed persons. The three persons do not know each other, and therefore, given the underlying fitness of the team, their opinions are independent. This assumption is very frequent in measurement models with latent variables and it is called* local independence assumption*, see Skrondal and Rabe-Hesketh [51] for details. This situation can be modelled by the network in Figure 11(b), which is a straightforward generalization of the network presented in Example 3.14 and illustrated in Figure 11(a).*

The case in which the local independence assumption holds is very easy to be modelled, but, as shown by the following example, the sources of information are not always independent given the underlying latent variable.

[16] See Skrondal and Rabe-Hesketh [51, Sect. 1.2] for a discussion about measurements model with latent variables.

**Example 3.16 (Dependent sources of information)** *Suppose that the sources of information about the variable* Fitness *consist of two persons, and that the opinion expressed by the second person is not only influenced by the underlying state of the latent variable, but also by the opinion expressed by the first person. Then the local independence assumption does not hold. It follows that the influence of the opinion of the first person on the opinion of the second person should be modelled explicitly as represented in Figure 11(c).*

In the examples above, the variables modelling the sources of information about a variable were always characterized by the same outcomes of the latent variable with in addition the value *missing*. But this is not always the case, as shown by the following example.

**Example 3.17 (Cold or flu?)** *Consider the situation described in Example 2.1. The situation can be modelled, depending on the way temperature is measured, by one of the networks in Figure 2. In this application, the only source of information are the measurements of the temperature. The variable* Cold or flu *has no direct source of information, and therefore the only way to gather information about it is to collect evidence about another latent variable (in this case the* body temperature*), that is related to* Cold of flu *in some (probabilistic) way.*

Finally, we can now complete the KM of the football example according to the instructions of the observational process expert reported in Example 3.13. The complete network modelling the domain knowledge and the observational process of the football example is displayed in Figure 12. After having completed the KM, each source of information should be described by the observational process expert through arguments. Some example of arguments are reported, together with their translation into linear inequalities, in the next section.

## 3.4 Response Assessment for the Observational Process

In the response assessment step, the observational process expert should translate the arguments expressed for each source of information into inequalities and obtain in this way the CSs required for the quantification of the network. The procedures are exactly like those in Section 3.2.

In the literature, KBSs are often developed by taking into consideration only the domain knowledge, without any explicit modelling of the observational process. Yet, as shown in [2], these approaches are implicitly assuming that: (i) the observed variables whose observations report non-missing outcomes are observed by *perfectly reliable* sources of information; (ii) the non-observed variables, and the observed ones whose observations report missing outcomes, are *missing at random* (MAR) (see Little and Rubin [43] and Gill et al. [33]). Let us first describe these two particular cases by means of examples and explain why they do not require an explicit response assessment of the observational process.

**Example 3.18 (Perfectly reliable source of information)** *Assume that the observation of the variable* Home match *is achieved as follows: the expert goes to the stadium the day of the match and, if the players and the supporters are there, he reports the value* yes *for his observation, and* no *otherwise. Assume that this procedure is particularly safe and it*

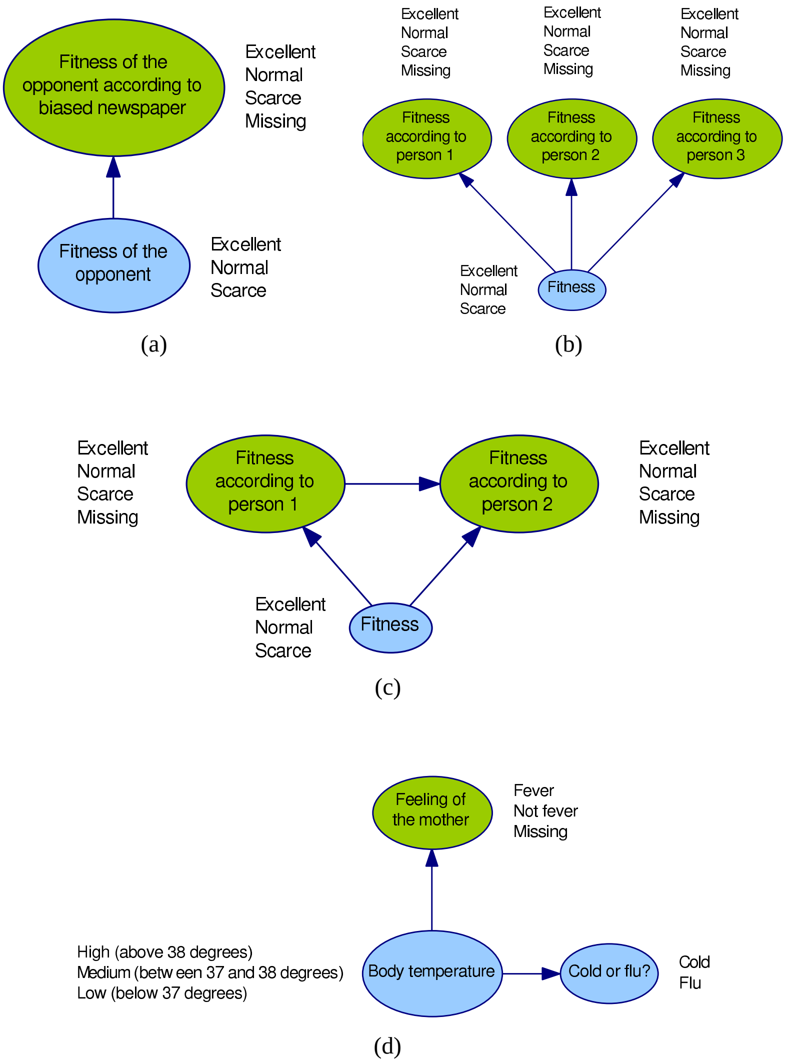


Figure 11: The networks of Examples 3.14–3.17.

*cannot return wrong (or missing) values. Accordingly, the response assessment leads to the following (precise) quantification of* $L$(*Home match observation* | *Home match*),

$$\begin{cases} P(H\,|\,H) = 1, & P(\neg H\,|\,H) = 0, & P(*\,|\,H) = 0, \\ P(H\,|\,\neg H) = 0, & P(\neg H\,|\,\neg H) = 1, & P(*\,|\,\neg H) = 0, \end{cases}$$

where $*$ denote the *missing* outcome.

As noted in this example, if the source of information delivers the correct value of the underlying latent variable in a perfectly reliable way, the observational process can be modelled through (degenerate) precise probabilistic quantification. It is sufficient to assign probability one to the outcome corresponding to the conditioning outcome, and probability zero to all the other outcomes (including *missing*). The manifest variable is therefore just a replica of the latent variable. Accordingly, the evidence about the manifest variable should

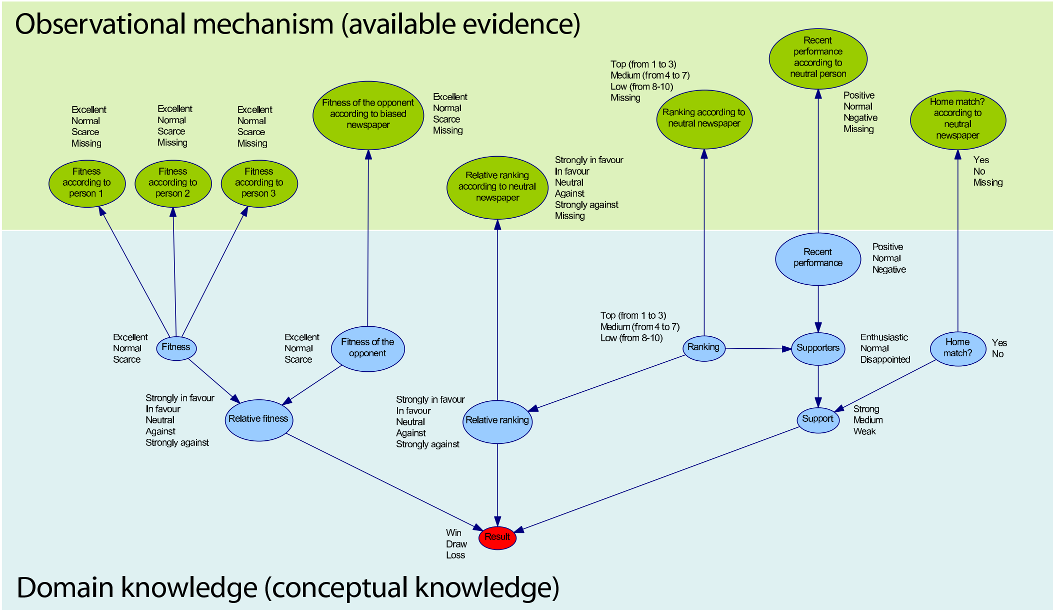


Figure 12: The CN modelling the football example.

be assigned also to the latent variable. Therefore, by exploiting the separation properties of CNs, we can remove the manifest variable from the network.

**Example 3.19 (MAR source of information)** *Consider the situation of Example 3.14. The source of information for the* Fitness of the opponent *is a newspaper. The outcome of the observation is* missing *because the newspaper does not report any information about the fitness of the Basel football team. Assume that this lack of information has no relation with the level of fitness of the team. Accordingly, the response assessment leads to the following quantification of* $L($*Fitness of the opponent observation* $|$ *Fitness of the opponent*$)$,

$$P(* \mid X) = P(* \mid N) = P(* \mid B),$$

*where* $X$, $N$ *and* $B$ *denote the outcomes* excellent, normal *and* bad. *Note that the same equality holds also in a situation where the newspaper has not been received by the observational process expert (and in this case all those probabilities are equal to one).*

The above example describes a source of information reporting a MAR missing outcome. Loosely speaking, MAR assumes that the missing values are produced by a *non-selective* process, i.e., a process that does not take into account the realized value of the latent variable to produce the missing observation. In other words, the fact that a variable has not been observed is uninformative about the underlying value of the latent variable. The probability of observing a missing value is therefore the same for each possible underlying value of the latent variable. A particular case of variable at the manifest level for which MAR holds is a variable that is never known (e.g., because we just do not observe it). In fact, in this case the probability of the variable being *missing* is one, independent of the underlying value of the latent variable. In any case, the MAR assumption models a condition of conditional independence between the manifest and the latent variable [2]. The arc connecting the latent with the manifest variable, and hence the manifest variable, can be removed from the network.

Summarizing, for both MAR and perfectly reliable sources, the manifest variables can be removed from the network, thus making an explicit modelling of the observational process unnecessary. Yet, these two cases reflect strong assumptions about the observational process, which are not necessarily satisfied in practice [34, 60]. In the rest of this section we illustrate through some other examples the response assessment procedure for more general observational processes.

**Example 3.20 (A malicious source of information)** *Consider the same situation as in the previous example. The* Fitness of the opponent *is reported by a newspaper, and this source of information reports the outcome* missing. *Yet, in this case, the newspaper has not reported any information about the fitness of the team, because the editors of the newspaper support the Basel football team and they prefer to not express negative judgements about their team.*

Unlike Example 3.19, in this example there is a relationship between the missingness and the realized value of the latent variable. This means that the MAR assumption is not tenable anymore. In these cases, an explicit quantification of the missingness process is therefore required. As an example, the fact that the editors of the newspaper support the Basel football team could be modelled by the following constraints,

$$P(*\,|\,B) > P(*\,|\,N), \qquad P(*\,|\,B) > P(*\,|\,X).$$

Yet, there are situations where an explicit modelling of the missingness process is not possible because we are completely *ignorant* about the relation between the latent and the manifest variable.

**Example 3.21 (Ignorance about the source of information)** *Consider the same situation of the previous example. Yet, in this case, no explanations are available in order to explain the lack of information about the fitness of the opponent. Perhaps the newspaper supports the Basel team as in the previous example, or it supports the Bellinzona team, or it decides to publish or not publish information about the fitness of the Basel team on the basis of same strange mechanism, which takes into account the actual fitness of the team.*

In a situation like the one in this example, there is not really any knowledge to be modelled by the missingness process. What should be modelled here is, in a sense, the condition of *ignorance* about the reasons for why the outcome of the manifest variable is missing. Imprecise probabilistic approaches are particularly suited for the modelling of these situations: the *vacuous* CS made of *all* the probability mass functions of a variable (i.e., the whole probability simplex) can properly model the ignorance about the state of the manifest variable conditional to any value of the latent variable. Accordingly, the only constraints we assign to the probabilities in the above example are the consistency conditions:

$$0 \leq P(*\,|\,B) \leq 1, \quad 0 \leq P(*\,|\,N) \leq 1, \quad 0 \leq P(*\,|\,X) \leq 1.$$

This represents the most conservative approach to the modelling of the missingness process, and should be adopted when no information is available about the reasons of the missingness of the information returned by the source.

The approach has been formalized and exploited by de Cooman and Zaffalon [22] in their *conservative updating rule* (CUR). A consequence of this approach with respect to updating is that the lower probability for a state of the variable of interest conditional on a CUR missing observation of a manifest variable is the minimum of the lower probabilities for this state conditional on all the possible values of the latent variable (corresponding to the CUR manifest variable). This is explained by the following example.

**Example 3.22** *In the football example, assume CUR for the variable* Fitness of the opponent observation*. Given the missing outcome of this variable, consider the lower probability of the state* win *(W) for the variable* Result*. CUR implies that this lower probability is*

$$\min\{\underline{P}(W \mid B), \underline{P}(W \mid N), \underline{P}(W \mid X)\}.$$

Accordingly, if a source of information reports a missing outcome and we want to model this missingness by CUR, we can either adopt the *vacuous* quantification as in Example 3.21 either consider all the possible explanations of the latent variable as in Example 3.22. The approach in Example 3.22 requires a number of evaluation exponential in the number of CUR variables, and it is therefore feasible only if few CUR variables are present, while in the other cases an approach as in Example 3.21 should be preferred.

In practice, we recommend to start with the response assessment for the observational process trying to model it explicitly. In this way, the observational process expert has the possibility of reasoning deeply about observations and missingness. If the expert remarks that the process either produces correct observations or missing values, then the observational process is a missingness process. In this case, the expert can continue by modelling explicitly the process, or he can assume some default assumption, like MAR or CUR. In the following example, we show how to specify completely the observational process for a single variable of the domain knowledge.

**Example 3.23 (A biased newspaper)** *Consider once more the newspaper of Example 3.14. Suppose that the newspaper has been described by the observational process expert through the following arguments,* the newspaper is cautious in expressing extreme opinions about the fitness state of the Basel football team. If it is not sure about the fitness state of the team, it prefers not to express itself. It prefers in general to express positive opinions. The probability of the newspaper expressing the right observation is higher than the probability of reporting something different or not expressing itself. *Suppose that the observational process expert decides to translate his arguments using the probability ratio scale in Figure 9 expressing the judgements given in Figure 13. Having expressed only arguments comparing different outcomes with the same conditioning event from the quantification of the* Reported fitness*, the CSs produced are automatically separately specified. We shorten* Fitness of the opponent *with* Fitness *and* Fitness of the opponent according to newspaper *with* Reported fitness*. The vertices of the CS* $K(\text{Reported fitness} \mid \text{Fitness=excellent})$ *can be obtained by solving the linear system of inequalities corresponding to the first table in Figure 13 and are the following, where we denote the possible values of* Fitness *and* Reported fitness *with*

$X$ *for* excellent*,* $N$ *for* normal*,* $B$ *for* bad *and* $*$ *for* missing*. The letter on the left hand side of the bar refers to* Reported fitness *and the letter on the right hand side of the bar refers to* Fitness*.*

| $P(*\,\vert\,X)$ | $P(X\,\vert\,X)$ | $P(N\,\vert\,X)$ | $P(B\,\vert\,X)$ |
|---|---|---|---|
| 0.05 | 0.73 | 0.18 | 0.04 |
| 0.02 | 0.77 | 0.19 | 0.02 |
| 0.05 | 0.75 | 0.18 | 0.02 |
| 0.02 | 0.75 | 0.18 | 0.05 |

*In the same way, the CS* $K(\text{Reported fitness}\,|\,\text{Fitness=normal})$ *is defined by the following vertices.*

| $P(*\,\vert\,N)$ | $P(X\,\vert\,N)$ | $P(N\,\vert\,N)$ | $P(B\,\vert\,N)$ |
|---|---|---|---|
| 0.09 | 0.09 | 0.78 | 0.04 |
| 0.09 | 0.04 | 0.78 | 0.09 |
| 0.08 | 0.08 | 0.75 | 0.09 |
| 0.09 | 0.05 | 0.82 | 0.04 |
| 0 | 0 | 1 | 0 |

*Finally, the CS* $K(\text{Reported fitness}\,|\,\text{Fitness=bad})$ *is defined by the following vertices.*

| $P(*\,\vert\,B)$ | $P(X\,\vert\,B)$ | $P(N\,\vert\,B)$ | $P(B\,\vert\,B)$ |
|---|---|---|---|
| 0.16 | 0.02 | 0.16 | 0.66 |
| 0.29 | 0 | 0.14 | 0.57 |
| 0.28 | 0.03 | 0.14 | 0.55 |
| 0.25 | 0 | 0.25 | 0.50 |
| 0.24 | 0.03 | 0.24 | 0.49 |
| 0.17 | 0 | 0.17 | 0.66 |

*To verify if the obtained CSs reflect the statements describing the newsletter, it is useful to consider the lower and upper probabilities implied by the above CSs. Consider for example the lower and upper probabilities when* Fitness *is* excellent*, we have,* $P(*\,|\,X) \in [0.02, 0.05], P(X\,|\,X) \in [0.73, 0.77], P(N\,|\,X) \in [0.18, 0.19], P(B\,|\,X) \in [0.02, 0.05]$*. We see that there is a non-negligible probability that the newsletter reports that the* Fitness *is* normal*, although the true fitness state is* excellent*. This reflects the caution of the newsletter in expressing extreme opinions. We can also verify, for example, if in this case* the probability of the newsletter expressing the right observation is higher than the probability of reporting something different or not expressing itself*, actually,* $\underline{P}(X\,|\,X) = 0.73$*. It is interesting to verify the behavior of the newsletter when the true underlying* Fitness *state is* bad*. We have,* $P(*\,|\,B) \in [0.16, 0.29]$*,* $P(X\,|\,B) \in [0, 0.03]$*,* $P(N\,|\,B) \in [0.14, 0.25]$*,* $P(B\,|\,B) \in [0.49, 0.66]$*. We see that in this case there is a non-negligible probability of* missing*. This reflects the cautiousness of the newsletter in expressing extreme opinions and its preference for positive opinions. If it is not sure about the true state and it supposes that the true state is negative (*bad*), then it prefers not to express itself.*

Under particular conditions, the quantification procedure described in the previous example can be significantly simplified, as shown in the following example.

**Example 3.24 (A simplified procedure)** *Consider once more the opinion of a single person about the* Fitness*. Suppose that in this case the level of expertise of the person is defined by the following qualitative statements:*

- *The probability of the person reporting the correct value of the* Fitness *is much higher than the probability of reporting something different.*
- *If he does not report the correct observation, he reports usually a value near to the correct one (small error).*
- *The probability of committing large errors or of not expressing any opinion are much smaller than the probability of reporting the correct value or of committing a small error.*

*The main difference with respect to the previous example is that the level of expertise of the person in this case is described with respect to the possible types of error that could be committed by himself, without referring to the particular underlying value of the* Fitness*. In this case, the quantification can be simplified, specifying only the possible types of error and comparing the probabilities of the different types of error. Suppose for example that we divide the possible observations into three categories:* correct *when the reported value corresponds to the true underlying value,* small error *when the expert reports a value that is close to the correct one and* large error *when the expert reports a value that is not close to the correct one. The classification of the possible combinations reported value/true value can be specified using a* similarity matrix*, like the one displayed in Figure 14.*

*Instead of filling three tables as in the preceding example, we can now fill only one table comparing the probabilities of the different types of error, as in the last table of Figure 14. Combining the information of the similarity matrix with the information in the latter table, we can (automatically) fill the three tables required for the quantification of the observation. Consider for example the second table, when we compare* excellent *with* normal *with as true underlying* Fitness *state* normal*. In this case* excellent *corresponds to a small error, while* normal *is the correct observation. The values to be inserted in this cell of the table are therefore the same as specified for the combination small error/correct in the latter table of Figure 14, i.e.,* $<<$ *and* $<$*.*

The assumption underlying the simplified approach is that we do not distinguish between errors of the same type. This assumption is reflected for example by the fact that for a combination of errors of the same type, we do not specify any value, like for example for the combination *excellent/bad* in the second table of Figure 14. It follows that this assumption is tenable only when we are dealing with sources of information whose attitude is independent of the underlying value and without preferences between the possible values. For example, this approach would not be justified in the case of the biased newspaper of the previous example. Actually, calculating the lower and upper probabilities when the underlying true *Fitness* state is *normal* we obtain $P(*\,|\,N) \in [0.04, 0.07]$, $P(X\,|\,N) \in [0.14, 0.25]$, $P(N\,|\,N) \in [0.50, 0.64]$, $P(B\,|\,N) \in [0.14, 0.25]$. We see that, as expected, the lower and upper probabilities assigned to the reported values *excellent* and *bad* are the same, because both are considered to be small errors in this case (second column of the similarity matrix in Figure 14). While, for the biased newspaper of the previous example, confounding *normal* with *excellent* was not the same as confounding *normal* with *bad*, because of its preference for positive opinions.

**Example 3.25 (Football example)** *The sources of information in the football example in Figure 12 can be defined combining the three examples above: the three sources of information about the* Fitness *and the source of information about the* Recent performance *are considered to be experts of the type described in Example 3.24. The newspaper reporting the* Fitness of the opponent *has been described in Example 3.23. The sources of information about* Ranking, Relative ranking *and* Home match *are perfectly reliable information sources. Finally,* Relative fitness, Supporters *and* Support *have no source of information.*

The response assessment of the observational process concludes the specification of the KBS, which can be regarded as a single CN over both the latent variables of the domain knowledge and the manifest variables associated with the outcomes of the observations. The KBS can now be used to extract posterior probabilistic information about any variable in the domain knowledge given some evidence specified according to the observational process, as shown in the next section.

True fitness state: excellent

| | Missing | | Excellent | | Normal | | Scarce | |
|---|---|---|---|---|---|---|---|---|
| Missing | | | | <<< | <<< | << | | |
| Excellent | >>> | | | | >> | >>> | >>> | |
| Normal | >> | >>> | <<< | << | | | >>> | >> |
| Scarce | | | | <<< | << | <<< | | |

True fitness state: normal

| | Missing | | Excellent | | Normal | | Scarce | |
|---|---|---|---|---|---|---|---|---|
| Missing | | | = | > | | <<< | = | > |
| Excellent | < | = | | | | <<< | | |
| Normal | >>> | | >>> | | | | >>> | |
| Scarce | < | = | | | | <<< | | |

True fitness state: scarce

| | Missing | | Excellent | | Normal | | Scarce | |
|---|---|---|---|---|---|---|---|---|
| Missing | | | >>> | | = | > | << | < |
| Excellent | | <<< | | | | <<< | | <<< |
| Normal | < | = | | | | | << | < |
| Scarce | > | >> | >>> | | > | >> | | |

Figure 13: The tables describing the newspaper reporting the *Fitness of the opponent* in Example 3.23.

## 4 Interaction with the Knowledge-Based System

The procedure described in the previous sections implements a KBS through a CN. From a theoretical point of view, the combination of the conditional independence relations depicted in the graphical structure of the network and the CSs defined for each variable defines a joint CS over the whole set of variables, called the *strong extension* of the CN [16], which models the knowledge of the domain expert and the observational process expert about the particular problem under consideration.

The kind of knowledge we typically want to extract from the KBS requires the *updating* of the CN, i.e., the calculation of the (posterior) probabilities for the hypothesis variable given some evidence reporting the values of some other variables in the network. In our particular case, the evidence used for updating consists in a value for each variable involved in the observational process, while the hypothesis variable is one of the variables in the

True fitness state: excellent

| | Missing | | Excellent | | Normal | | Scarce | |
|---|---|---|---|---|---|---|---|---|
| Missing | | | | <<< | << | < | < | > |
| Excellent | >>> | | | | > | >> | >>> | |
| Normal | > | >> | << | < | | | >> | >>> |
| Scarce | < | > | | <<< | <<< | << | | |

True fitness state: normal

| | Missing | | Excellent | | Normal | | Scarce | |
|---|---|---|---|---|---|---|---|---|
| Missing | | | << | < | | <<< | << | < |
| Excellent | > | >> | | | << | < | | |
| Normal | >>> | | > | >> | | | > | >> |
| Scarce | > | >> | | | << | < | | |

True fitness state: scarce

| | Missing | | Excellent | | Normal | | Scarce | |
|---|---|---|---|---|---|---|---|---|
| Missing | | | < | > | << | < | | <<< |
| Excellent | < | > | | | <<< | << | | <<< |
| Normal | > | >> | >> | >>> | | | << | < |
| Scarce | >>> | | >>> | | > | >> | | |

(a)

*True underlying value (latent variable)*

| *Observation* | Excellent | Normal | Scarce |
|---|---|---|---|
| **Excellent** | Correct | Small error | Large error |
| **Normal** | Small error | Correct | Small error |
| **Scarce** | Large error | Small error | Correct |

(b)

| | Correct | | Small error | | Large error | | Missing | |
|---|---|---|---|---|---|---|---|---|
| Correct | | | > | >> | >>> | | >>> | |
| Small error | << | < | | | >> | >>> | > | >> |
| Large error | | <<< | <<< | << | | | < | > |
| Missing | | <<< | << | < | < | > | | |

(c)

Figure 14: The tables describing the observation of the *Fitness* through a neutral person. The table comparing the probabilities of the different types of error (a), the similarity matrix classifying the possible combinations of true underlying value and observation with respect to the defined types of error (b), and the three tables obtained combining the information of the table below and the information of the similarity matrix (c).

domain knowledge which usually (but not necessarily) has been already identified during the construction of the KBS. For example, in the example in Figure 12, the evidence used for the updating consists of a value for each one of the variables in the observational process, while the hypothesis variable is the *Result*.

From a theoretical point of view, as illustrated in Antonucci and Zaffalon [6], a CN under strong independence is equivalent to a set of BNs. Accordingly, it would be possible to employ algorithms used for BNs updating[17] to update CNs, by simply updating each BN in the set. Yet, this approach is feasible only for very small and simple CNs: the number of BNs corresponding to a CN increases exponentially in the number of vertices of the CSs and variables in the network, being therefore too large in most of the cases. Thus, standard algorithms for BNs cannot be easily applied to CNs. Nevertheless, in the recent times, several exact and approximate procedures for CNs updating have been proposed,

[17] See for example Jensen and Nielsen [38].

see for example de Campos and Cozman [19] for exact algorithms and Antonucci et al. [4] for approximate algorithms. This makes it possible to update CNs of medium size in reasonable time (e.g., [1],[3]).

In this section, we show several examples of updating for the CN in Figure 12. In each example, we report the posterior probabilities computed by the exact algorithm proposed in de Campos and Cozman [19] and the approximate algorithm described in Antonucci et al. [4]. The computational times of the exact and approximate algorithms are very different: about one hour for the exact procedure and few seconds for the approximate (on a 2.8 GHz Pentium 4 machine). In practice, for small and medium sized models, the choice between exact and approximate algorithms depends on the time and the accuracy constraints of the particular problem.

**Example 4.1 (A favorable scenario)** *In this example we illustrate the case corresponding to the most favorable evidence for the Bellinzona football team. In the following table, we indicate in the right column the evidence provided by the sources of information about the latent variable on the left. For the latent variable* Fitness*, we report three values corresponding to the opinions of the three persons represented as sources of information in Figure 12.*

| Latent variable | Evidence |
|---|---|
| Fitness | excellent, excellent, excellent |
| Fitness of the opponent | bad |
| Relative ranking | strongly in favour |
| Ranking | top |
| Recent performance | positive |
| Home match | yes |

*The results of the updating are lower and upper probabilities, i.e., a probability interval, for the outcomes of the hypothesis variable* Result*. In the table below, the probability intervals produced with exact algorithms [20] and the probability intervals produced with the approximate algorithm called GL2U [4] are displayed. Note that the results produced by the approximate algorithm are a good approximation of the exact ones. This has actually occurred in all the examples illustrated in this section.*

| *Outcome* | *Exact* | *GL2U* |
|---|---|---|
| win | $[0.88, 1.00]$ | $[0.87, 1.00]$ |
| draw | $[0.00, 0.12]$ | $[0.00, 0.11]$ |
| loss | $[0.00, 0.12]$ | $[0.00, 0.12]$ |

Given the posterior probability intervals returned by the updating algorithms, we need now a criterion for identifying one or more outcomes of the hypothesis variable to be considered as the most plausible given the available evidence. To this aim, different approaches have been proposed in literature. In the case of CNs, as the main algorithms for updating return probability intervals, the *interval dominance* is the most common approach and the easiest to understand: the outcomes of the hypothesis variable whose upper probability is smaller than the lower probability of some other outcome are rejected. Figure 15(a) reports the (exact) posterior probability intervals for the example above. According to the interval dominance criterion, the outcomes *draw* and *loss* should be rejected, as their upper probabilities are smaller than the lower probability of *win*. Thus, *win* should be regarded as the only acceptable option. Other decision criteria, like for example *maximality* [57, chapter 3.9], could be also adopted. For a review of decision criteria with imprecise probabilities see Troffaes [53].

**Example 4.2 (An unfavorable scenario)** *In this example we illustrate a situation in which evidence is against the football team of Bellinzona.*

| Latent variable | Evidence |
|---|---|
| Fitness | bad, bad, bad |
| Fitness of the opponent | excellent |
| Relative ranking | strongly against |
| Ranking | low |
| Recent performance | negative |
| Home match | no |

*The results, which are depicted in Figure 15(b), are the following.*

| *Outcome* | *Exact* | *GL2U* |
|---|---|---|
| win | $[0.10, 0.39]$ | $[0.09, 0.37]$ |
| draw | $[0.09, 0.37]$ | $[0.08, 0.37]$ |
| loss | $[0.42, 0.81]$ | $[0.40, 0.83]$ |

*It follows that the result produced by the KBS in this case is* loss.

**Example 4.3 (An ambiguous scenario)** *In this example we illustrate a scenario in which the evidence suggests that that the two teams are equivalent from the point of view of ranking and fitness.*

| Latent variable | Evidence |
|---|---|
| Fitness | excellent, excellent, excellent |
| Fitness of the opponent | normal |
| Relative ranking | neutral |
| Ranking | medium |
| Recent performance | normal |
| Home match | yes |

*The results, which are depicted in Figure 15(c), are the following.*

| *Outcome* | *Exact* | *GL2U* |
|---|---|---|
| win | $[0.36, 0.75]$ | $[0.35, 0.74]$ |
| draw | $[0.12, 0.51]$ | $[0.12, 0.53]$ |
| loss | $[0.13, 0.52]$ | $[0.12, 0.53]$ |

*We see that the three probability intervals are overlapping. The situation in which the probability intervals associated with two outcomes are overlapping models a situation of doubt between the two outcomes in the light of the available evidence. In this case, the absence of a dominance means that the KBS is unable to produce a single outcome as result.*

The possibility of representing ambiguity and indecision, as in the example above, is one of the major advantages of CNs with respect to BNs. Actually, it occurs very often in practice to have situations of ambiguity, while a BN in this case would produce however a precise probability for each outcome and therefore a prediction, a CN refrains from producing a prediction highlighting a situation of ambiguous or poor evidence. This feature of CNs makes them suitable for modelling information fusion in case of contrasting or contradictory evidence, as shown in the following example.

**Example 4.4 (Partial indecision)** *In this example we illustrate a scenario where the fitness and the ranking of the two teams are similar, but the match is played in the home field.*

| Latent variable | Evidence |
|---|---|
| Fitness | normal, normal, normal |
| Fitness of the opponent | normal |
| Relative ranking | neutral |
| Ranking | medium |
| Recent performance | positive |
| Home match | yes |

*The results, which are depicted in Figure 15(d), are the following.*

| *Outcome* | *Exact* | *GL2U* |
|---|---|---|
| win | $[0.41, 0.57]$ | $[0.43, 0.58]$ |
| draw | $[0.32, 0.44]$ | $[0.35, 0.42]$ |
| loss | $[0.11, 0.24]$ | $[0.13, 0.22]$ |

*The outcome* loss *is dominated by* draw *(and also by* win*). According to the interval dominance criterion, we should therefore reject this option and regard this scenario as a condition of partial indecision between* win *and* draw.

**Example 4.5 (Disagreeing sources of information)** *In this example we illustrate a case in which the three persons, whose opinions are used to collect information about the* Fitness, *disagree.*

| Latent variable | Evidence |
|---|---|
| Fitness | excellent, medium, bad |
| Fitness of the opponent | bad |
| Relative ranking | strongly against |
| Ranking | low |
| Recent performance | negative |
| Home match | no |

*The results are the following.*

| *Outcome* | *Exact* | *GL2U* |
|---|---|---|
| win | $[0.10, 0.49]$ | $[0.10, 0.49]$ |
| draw | $[0.06, 0.51]$ | $[0.05, 0.51]$ |
| loss | $[0.28, 0.84]$ | $[0.26, 0.86]$ |

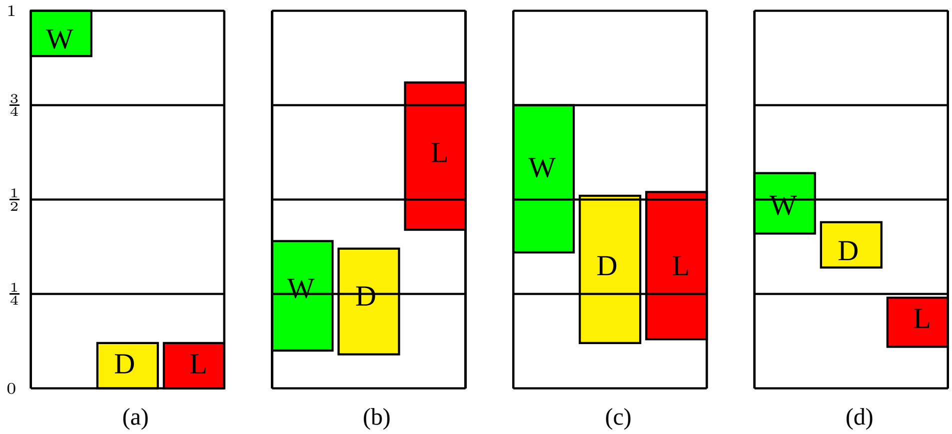


Figure 15: Histograms modelling the (exact) posterior intervals for the variable *Result* according to the scenarios described in Example 4.1 (a), Example 4.2 (b), Example 4.3 (c), and Example 4.4 (d). The colors red, yellow and green correspond respectively to three outcomes *win* (W), *draw* (D) and *loss* (L).

*The three probability intervals are overlapping also in this case, denoting a situation of doubt. The width of the intervals highlights a situation of great ambiguity.*

In general, the ambiguity, the contrast and the eventual contradiction of different sources of information have a negative effect on the width of the probability intervals produced. This is for example the case when some sources of information produce *missing* values, as shown in the following example.

**Example 4.6 (Missing values 1)** *In this example we illustrate the effect of* missing *values on the results and in particular on the width of the probability intervals produced for the hypothesis variable. Consider the following scenario.*

| Latent variable | Evidence |
|---|---|
| Fitness | excellent, excellent, excellent |
| Fitness of the opponent | excellent |
| Relative ranking | strongly in favour |
| Ranking | top |
| Recent performance | positive |
| Home match | no |

*The results are the following.*

| *Outcome* | *Exact* | *GL2U* |
|---|---|---|
| win | $[0.82, 0.96]$ | $[0.81, 0.98]$ |
| draw | $[0.03, 0.17]$ | $[0.01, 0.18]$ |
| loss | $[0.01, 0.19]$ | $[0.01, 0.18]$ |

*The evidence is clearly in favour of a* win *of the Bellinzona football team. Suppose now that only one person has expressed his opinion about the* Fitness *and the opinions of the other two persons are* missing, *then the results are the following.*

| *Outcome* | *Exact* | *GL2U* |
|---|---|---|
| win | $[0.72, 0.98]$ | $[0.70, 0.97]$ |
| draw | $[0.01, 0.25]$ | $[0.01, 0.29]$ |
| loss | $[0.01, 0.27]$ | $[0.01, 0.29]$ |

*We see that the prediction remains the same, but the probability intervals are larger, denoting a situation of increased uncertainty. The same holds if, instead of having two* missing *values, the three experts disagree completely, as in the previous example, the results in this case are the following.*

| *Outcome* | *Exact* | *GL2U* |
|---|---|---|
| win | $[0.65, 0.99]$ | $[0.70, 0.99]$ |
| draw | $[0.00, 0.31]$ | $[0.00, 0.29]$ |
| loss | $[0.01, 0.32]$ | $[0.00, 0.30]$ |

**Example 4.7 (Missing values 2)** *Finally, it is interesting to consider a case where the only source of information about the* Fitness of the opponent *is* missing. *Consider once more Scenario 1, the probability intervals for* Result *are the following.*

| *Outcome* | *Exact* | *GL2U* |
|---|---|---|
| win | $[0.88, 1.00]$ | $[0.87, 1.00]$ |
| draw | $[0.00, 0.12]$ | $[0.00, 0.11]$ |
| loss | $[0.00, 0.12]$ | $[0.00, 0.12]$ |

*Suppose now that the newspaper reporting the* Fitness of the opponent *does not report any information. In this case the value associated with this source of information is* missing. *The probability intervals associated with the* Result *are the following.*

| *Outcome* | *Exact* | *GL2U* |
|---|---|---|
| win | $[0.72, 0.98]$ | $[0.73, 1.00]$ |
| draw | $[0.01, 0.25]$ | $[0.00, 0.25]$ |
| loss | $[0.01, 0.27]$ | $[0.00, 0.27]$ |

*Also in this case the intervals are larger than in the original scenario, reflecting thus the increased uncertainty.*

# 5 Conclusions

In this chapter we have described a step-by-step procedure for building KBSs in the framework of CNs. Our approach satisfies several conditions which seems to be desirable for the implementation of a KBS. They are the following.

- **Imprecision**: CNs allow to model qualitative expert knowledge without introducing unrealistic or unnecessary assumptions.

- **Aggregation**: CNs allow to merge knowledge coming from different experts or sources of information, in a realistic way. This feature is particularly interesting when the fusion of different sources of information should be modelled.
- **Modularity**: the expert knowledge is specified separately for each variable in the CN. This feature allows to modify the quantification of a particular variable without the need of modifying the quantification of the others.
- **Transferability**: a CN can be easily represented in the memory of a computer. The qualitative structure can be passed as it is, while CSs can be represented through their vertices or through the systems of inequalities generating them.
- **Observation** and **separation**: by specifying in a separate way the observational process and the quantification of the expert knowledge through CSs it is possible to achieve a realistic description of the way the evidence is collected in practice. In particular, it is possible to model the fusion of sensors, especially when they are contrasting or contradicting, and the processes producing missing values. Furthermore, the same domain knowledge can be adapted to different practical situations by redefining only the observational process.
- **Reasoning**, **completeness** and **ontology**: the representation of the expert knowledge through the qualitative part of the CN and the lists of arguments help both the expert and other persons that have not directly participated to the development of the KBS to access easily the expert knowledge. This feature is particularly useful to improve during the debug of the system and also in order to favor the acceptance by human users of the outputs of the system. Furthermore, the structured specification of arguments for each variable in the CN is helpful for the expert to recall in mind all the knowledge that is relevant for the given problem and eventually for detecting contradictions, contrasting beliefs and partial ignorance.

Despite their desirable features, KBSs based on CNs are not yet widespread in the AI community. Actually, modelling based on CNs for KBSs has been limited in the past years by the lack of efficient inference algorithms. Nowadays the situation is different, and medium-sized CNs can be efficiently updated with good accuracy [1, 3]. More specifically, the particular elicitation procedure described in this chapter has been successfully adopted for the development of a KBS for airspace surveillance. This makes CNs a suitable alternative to other probabilistic frameworks for KBSs development, especially if the expert knowledge is expressed through arguments and the available evidence is characterized by many, even contradictory, sources of information and by non-trivial missingness processes.

## Acknowledgments

This work was supported by Armasuisse, Competence Sector Science + Technology Management, and by the Swiss NSF grants n. 200020-121785/1 and n. 200020-116674/1. The original manuscript of this chapter was reviewed by Fabio G. Cozman and Silja Renooij, whose insightful comments were very helpful to improve the presentation of the arguments.

## References


[1] A. Antonucci, R. Brühlmann, A. Piatti, and M. Zaffalon. Credal networks for military identification problems. *International Journal of Approximate Reasoning*, 50(4):666–679, 2009.

[2] A. Antonucci and A. Piatti. Modeling unreliable observations in Bayesian networks by credal networks. In *Proceedings of the Third International Conference on Scalable Uncertainty Management (SUM 2009)*, pages 28–39, 2009.

[3] A. Antonucci, A. Salvetti, and M. Zaffalon. Credal networks for hazard assessment of debris flows. In J. Kropp and J. Scheffran, editors, *Advanced Methods for Decision Making and Risk Management in Sustainability Science*. Nova Science Publishers, New York, 2007.

[4] A. Antonucci, Y. Sun, C. P. de Campos, and M. Zaffalon. Generalized loopy 2U: a new algorithm for approximate inference in credal networks. *International Journal of Approximate Reasoning*. To appear.

[5] A. Antonucci and M. Zaffalon. Equivalence between Bayesian and credal nets on an updating problem. In J. Lawry, E. Miranda, A. Bugarin, S. Li, M. A. Gil, P. Grzegorzewski, and O. Hryniewicz, editors, *Soft Methods for Integrated Uncertainty Modelling, Springer (Proceedings of the third international conference on Soft Methods in Probability and Statistics: SMPS 2006)*, pages 223–230, 2006.

[6] A. Antonucci and M. Zaffalon. Decision-theoretic specification of credal networks: a unified language for uncertain modeling with sets of Bayesian networks. *International Journal of Approximate Reasoning*, 49(2):345–361, 2008.

[7] D. Avis. A revised implementation of the reverse search vertex enumeration algorithm. In G. Kalai and G. Ziegler, editors, *Polytopes–Combinatorics and Computation*, pages 177–197. Birkhäuser, 2000.

[8] D. Avis and K. Fukuda. A pivoting algorithm for convex hulls and vertex enumeration of arrangements and polyhedra. *Discrete and Computational Geometry*, 8:295–313, 1992.

[9] R.P. Bagozzi and L.W. Phillips. Representing and testing organizational theories: A holistic construal. *Administrative Science Quarterly*, 27(3):459–489, 1982.

[10] A. Benavoli and A. Antonucci. Aggregating imprecise probabilistic knowledge. In T. Augustin, F.P.A. Coolen, and S. Moral, editors, *Proceedings of the 6th International Simposium on Imprecise Probability: Theories and Applications (ISIPTA'09)*, pages 31–40, Durham, 2009.

[11] P.G. Benson, S.P. Curley, and G.F. Smith. Belief assessment: An underdeveloped phase of probability elicitation. *Management Science*, 41(10):1639–1653, 1995.

[12] D. Borsboom, G. J. Mellenbergh, and J. van Heerden. The theoretical status of latent variables. *Psychological Review*, 110(2):203–219, 2002.

[13] G.J. Browne, S.P. Curley, and P.G. Benson. Evoking information in probability assessment: Knowledge maps and reasoning–based directed questions. *Management Science*, 43(1):1–14, 1997.

[14] G. Choquet. Theory of capacities. *Annales de l'Institut Fourier*, pages 131–295, 1953–1954.

[15] I. Couso, S. Moral, and P. Walley. A survey of concepts of independence for imprecise probabilities. *Risk Decision and Policy*, 5(2):165–181, 2000.

[16] F. G. Cozman. Credal networks. *Artificial Intelligence*, 120:199–233, 2000.

[17] F. G. Cozman. Graphical models for imprecise probabilities. *International Journal of Approximate Reasoning*, 39(2–3):167–184, 2005.

[18] A. Daneshkhah. Psychological aspects influencing elicitation of subjective probability. Beep working paper, Department of Probability and Statistics, University of Sheffield, 2004.

[19] C. P. de Campos and F. G. Cozman. Inference in credal networks through integer programming. In G. de Cooman, J. Vejnarova, and M. Zaffalon, editors, *Proceedings of the fifth International Symposium on Imprecise Probabilities (ISIPTA '07)*, pages 145–154, Prague, 2007.

[20] C. P. de Campos and F. G. Cozman. Inference in credal networks through integer programming. In G. de Cooman, J. Vejnarova, and M. Zaffalon, editors, *Proceedings of the fifth International Symposium on Imprecise Probabilities (ISIPTA '07)*, pages 145–154, Prague, 2007.

[21] G. de Cooman, F. Hermans, A. Antonucci, and M. Zaffalon. Epistemic irrelevance in credal networks: the case of imprecise Markov trees. In T. Augustin, F.P.A. Coolen, and S. Moral, editors, *Proceedings of the 6th International Simposium on Imprecise Probability: Theories and Applications (ISIPTA'09)*, pages 149–158, Durham, 2009.

[22] G. De Cooman and M. Zaffalon. Updating beliefs with incomplete observations. *Artificial Intelligence*, 159:75–125, 2004.

[23] S. Destercke, D. Dubois, and E. Chojnacki. Unifying practical uncertainty representations: (i) generalized p-boxes. *International Journal of Approximate Reasoning*, 49(3):649–663, 2008.

[24] S. Destercke, D. Dubois, and E. Chojnacki. Unifying practical uncertainty representations: (ii) clouds. *International Journal of Approximate Reasoning*, 49(3):664–677, 2008.

[25] M.J. Druzdzel and L.C. van der Gaag. Building probabilistic networks: Where do the numbers come from ? *IEEE Transactions on Knowledge and Data Engineering*, 12:481–486, 2000.

[26] D. Dubois and H. Prade. *Possibility Theory*. Plenum Press, New York, 1988.

[27] J.R. Edwards and R.P. Bagozzi. On the nature and direction of relationships between constructs and measures. *Psychological Methods*, 5(2):155–174, 2000.

[28] M.J. Eppler. A process-based classification of knowledge maps and application examples. *Knowledge and Process Management*, 15(1):59–71, 2008.

[29] B. Fischhoff and M. Bar-Hillel. Focusing techniques: A shortcut to improving probability judgement? *Organizational Behavior and Human Performance*, 34:175–194, 1984.

[30] Shafer G. *A Mathematical Theory of Evidence*. Princeton University Press, 1976.

[31] L.C. van der Gaag, S. Renooij, C.L.M. Witteman, B. Aleman, and B.G. Taal. How to elicit many probabilities. In K.B. Laskey and H. Prade, editors, *Proceedings of the Fifteenth Conference on Uncertainty*, pages 647–654, San Francisco, 1999. Morgan Kaufman Publishers.

[32] L.C. van der Gaag, S. Renooij, C.L.M. Witteman, B. M. P. Aleman, and B.G. Taal. Probabilities for a probabilistic network: A case-study in oesophageal cancer. *Artificial Intelligence in Medicine*, 25(2):123–148, 2002.

[33] R. Gill, M. Van der Laan, and J. Robins. Coarsening at random: characterisations, conjectures and counter-examples. In D. Y. Lin, editor, *Proceedings of the first Seattle Conference on Biostatistics*, pages 255–294. Springer, 1997.

[34] P. Grünwald and J. Halpern. Updating probabilities. *Journal of Artificial Intelligence Research*, 19:243–278, 2003.

[35] E.M. Helsper and L.C. van der Gaag. Ontologies for probabilistic networks. Technical Report UU-CS-2003-042, Institute of information and computing sciences, Utrecht University, 2003.

[36] R.A. Howard. Knowledge maps. *Management Science*, 35:903–922, 1989.

[37] D. Jenkinson. The elicitation of probabilities-a review of the statistical literature. Beep working paper, Department of Probability and Statistics, University of Sheffield, 2005.

[38] F. T. Jensen and T. D. Nielsen. *Bayesian Networks and Decision Graphs*. Springer, New York, 2007.

[39] D. Kahneman, P. Slovic, and A. Tversky. *Judgement under uncertainty: Heuristics and biases*. Cambridge University Press, New York, 1982.

[40] M. R. Klein and L. B. Methlie. *Knowledge-Based Decision Support Systems: With Applications in Business*. John Wiley & Sons, Inc., New York, NY, USA, 1995.

[41] A. Kovačec and B. Ribeiro. Convex hull calculations: a matlab implementation and correctness proofs for the lrs–algorithm. Technical report, Pr–Publicações do Departamento de Matemática, Preprint Number 03–26, Universidade de Coimbra, 2003.

[42] I. Levi. *The Enterprise of Knowledge*. MIT Press, London, 1980.

[43] R.J.A. Little and D.B. Rubin. *Statistical Analysis with Missing Data*. John Wiley & Sons, Inc., New York, NY, USA, 1987.

[44] M.W. Merkhofer. Using influence diagrams in multiattribute utility analysis–improving effectiveness through improving communication. In R.M. Oliver and J.Q. Smith, editors, *Influence Diagrams, Belief Nets and Decision Analysis*. Wiley, New York, 1990.

[45] S. Moral and A. Cano. Strong conditional independence for credal sets. *Annals of Mathematics and Artificial Intelligence*, 35(1–4):295–321, 2002.

[46] A. O' Hagan, C. E. Buck, A. Daneshkhah, J. E. Eiser, P. H. Garthwaite, D. J. Jenkinson, J. E. Oakley, and T. Rakow. *Uncertain Judgements: Eliciting Expert Probabilities*. John Wiley and Sons, Chichester, 2006.

[47] C. Ofir and S.K. Reddy. Measurement errors in probability judgements. *Management Science*, 42:1308–1325, 1996.

[48] O. Pourret, P. Naim, and B. Marcot. *Bayesian Networks: A Practical Guide to Applications*. Wiley, Chichester, 2008.

[49] S. Renooij and L.C. van der Gaag. From qualitative to quantitative probabilistic networks. In A. Darwiche and N. Friedman, editors, *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence*, pages 422–429, San Francisco, 2002. Morgan Kaufman Publishers.

[50] S. Renooij and C. Witteman. Talking probabilities: communicating probabilistic information with words and numbers. *International Journal of Approximate Reasoning*, 22:169–194, 1999.

[51] A. Skrondal and S. Rabe-Hesketh. *Generalized latent variable modeling: multilevel, longitudinal, and structural equation models*. Chapman and Hall/CRC, Boca Raton, 2004.

[52] S.E. Toulmin. *The uses of Argument*. Cambridge University Press, Cambridge, 1958.

[53] M. C. M. Troffaes. Decision making under uncertainty using imprecise probabilities. *International Journal of Approximate Reasoning*, 45(1):17–29, 2007.

[54] M. Uschold and M. Gruninger. Ontologies: principles, methods and applications. *Knowledge Engineering Review*, 11(2):93–136, 1996.

[55] E.F. Vail. Mapping organizational knowledge. *Knowledge Management Review*, 2(8):10–15, 1999.

[56] L.C. van der Gaag and E.M. Helsper. Experiences with modelling issues in building probabilistic networks. In Gómez-Pérez A. and Benjamins V.R., editors, *Knowledge Engineering and Knowledge Management: Ontologies and the Semantic Web, Proceedings of EKAW 2002*, pages 21–26, Berlin, 2002. Springer-Verlag.

[57] P. Walley. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, New York, 1991.

[58] P. Walley. Measures of uncertainty in expert systems. *Artificial Intelligence*, 83:1–58, 1996.

[59] M. Zaffalon. Conservative rules for predictive inference with incomplete data. In F. G. Cozman, R. Nau, and T. Seidenfeld, editors, *ISIPTA '05: Proceedings of the Fourth International Symposium on Imprecise Probabilities and Their Applications*, pages 406–415, Manno, Switzerland, 2005. SIPTA.

[60] M. Zaffalon and E. Miranda. Conservative inference rule for uncertain reasoning under incompleteness. *Journal of Artificial Intelligence Research*, 34:757–821, 2009.