

A Review of Explanation Methods for Bayesian Networks

Carmen Lacave
Dept. Computer Science
University of Castilla-La Mancha
Paseo de la Universidad, s/n
13071 Ciudad Real, Spain
E-mail: clacave@inf-cr.uclm.es

Francisco J. Díez
Dept. Artificial Intelligence
UNED
Senda del Rey, 9
28040 Madrid, Spain
E-mail: fjdiez@dia.uned.es

February 21, 2002

Abstract

One of the key factors for the acceptance of expert systems in real world domains is the capability to explain their reasoning [BS84, HD90]. This paper describes the basic properties that characterise explanation methods and reviews the methods developed up to date for explanation in Bayesian networks.

1 Introduction

Expert systems originated in the 1970s as computer programs capable to imitate human experts and even substituting them when necessary. One of the essential qualities of real experts is their ability to communicate their knowledge and explain their reasoning. This ability is especially important in the case of expert systems, not only for tracing the performance during the construction and evaluation of the system, but also for justifying their results when the system is deployed in an operating environment. In fact, an experiment performed at the MYCIN project showed that physicians are very reluctant to accept the advice of a machine if they do not understand

how it was obtained [TS84].

In the decades that followed, i.e., the 1980s and 1990s, the main goal of artificial intelligence shifted from *imitating* natural intelligence to *supporting* human beings in a synergistic way. In fact, Clancey [Cla93] points out that the authors' note of the *Knowledge Acquisition* journal said: "The key issue is not *artificial* intelligence, but how to extend *natural* intelligence through knowledge-based systems." In this context, explanation in expert systems becomes even more important, because human-computer collaboration requires mutual understanding: machine models must take into account human cognitive processes and at the same time must make artificial reasoning understandable for human users. It is surprising, however, that the amount of research devoted to this subject has been relatively small, compared to other areas of artificial intelligence. There are only isolated pieces of work, seldom used in real-world applications; in fact, the explanation capability of most of today's expert systems and packages is even poorer than that of MYCIN, which is often regarded as the first expert system.

There was another shift in the 80s, due to the need of uncertain reasoning in virtually all applications of expert systems, especially in medicine. It was shown that MYCIN's model of certainty factors was inconsistent and might lead to wrong results. The model of PROSPECTOR was not any better, and Dempster-Shafer theory, besides lacking a universally accepted interpretation, was impractical for real-world problems—at least until graphical methods came on the scene. Only fuzzy-set methods seemed to be useful for building expert systems in uncertain domains, but due to the lack of a clear semantics and of a consistent methodology, many researchers considered those methods as a heterogeneous set of ad-hoc solutions without a firm theoretical basis. In the same decade, advances in Bayesian Networks (BNs) [CGH97, Jen96, Pea88] and influence diagrams [HM84, Sha86b] showed theoretically and empirically that it was feasible to build probabilistic expert

systems without introducing unrealistic assumptions of independence. The advantages of these models are that they have a clear interpretation, easily combine subjective estimated and statistical data, and can be justified on theoretical grounds. In contrast, the main disadvantage is that their reasoning method follows a normative approach, and consequently, the explanation of inference is more difficult than in the methods that try to imitate human reasoning. Therefore, the need for explanation methods is even more important in Bayesian networks than in heuristic expert systems.

The aim of this paper is to analyse what has been done up to date and what remains to be done in the field of explanation in Bayesian networks. (In a future work we will review explanation methods for heuristic expert systems.) Therefore, after introducing very briefly Bayesian networks in Section 2.1 and explanation in Section 2.2, we study the fundamental properties of explanation in Section 3. In the light of such properties, we review the methods of explanation for Bayesian networks proposed in the literature (Section 4), and conclude by pointing out possible lines for future research (Section 5).

2 Preliminaries

2.1 Bayesian networks

A Bayesian network consists of an acyclic directed graph (ADG),¹ whose nodes represent random variables, together with a conditional probability distribution for each node X_i given its parents,² $P(x_i|pa(x_i))$. The conditional probability for a node without parents is just its prior probability $P(x_i|\emptyset) = P(x_i)$. These probabilities can be obtained from statistical data

¹A cycle is a closed directed path $X_i \rightarrow X_j \rightarrow X_k \rightarrow \dots \rightarrow X_i$. A restriction of Bayesian networks is that their graphs cannot contain cycles.

²A node X_i is a parent of X_j if and only if there is a directed link $X_i \rightarrow X_j$. It is usual to represent a configuration of the parents of X_i by $pa(x_i)$.

(for instance, from a database), from the literature on the specific domain or by the judgment of human experts.

The joint probability represented by a Bayesian network is

$$P(x_1, \dots, x_n) = \prod_i P(x_i | pa(x_i)) \quad (1)$$

This distribution satisfies the *d-separation* property [Pea88] and its equivalent, the Markov property [Jen96, Nea90], which states that a node is independent of its non-descendants in the graph given its parents. Roughly speaking this property implies that a link $X \rightarrow Y$ in a Bayesian network represents a probabilistic dependence between X and Y , while the lack of a link represents a probabilistic independence.

A *finding* is a piece of information that states with certainty the value of a random variable; a finding may be, for example, that the patient is a male; other findings might be that he is 54 years old, that he has a fever, that he does not usually have headaches, etc. The set of findings is called *evidence*, \mathbf{e} . Probabilistic reasoning consists of computing the posterior probability of the unobserved variables given the evidence; for instance, $P(x_i | \mathbf{e})$ or $P(x_i, x_j, x_k | \mathbf{e})$. This process is usually called *evidence propagation*, and is based, more or less explicitly, in the application of Bayes theorem.

2.2 Defining explanation

From the origins of philosophy and science, many researchers of both disciplines have tried to determine the meaning of explanation, connecting it with comprehension and description. According to *The Concise Oxford Dictionary of Current English*, an explanation is: “A statement or circumstance that explains something | A declaration made with a view to mutual understanding or reconciliation |” For the *Webster’s New World Dictionary*, some meanings of the term *explain* are: “to make clear, plain or understandable | to give the meaning or interpretation of; expound | to account for;

state reasons for; to give an explanation.”

After these definitions, we can conclude that explaining consists of **exposing something** in such a way that it is

understandable for the receiver of the explanation, which implies that he/she improves his/her knowledge about the object of the explanation; and

satisfactory as far as it covers the receiver’s expectations.

In the field of expert systems, there are different concepts of explanation. For instance, the rule-based expert system MYCIN [BS84], the first artificial program able to explain its reasoning, could show **how** it had obtained some conclusion (i.e., which rules it had applied to deduce a certain proposition) or **why** it was requesting additional information from the user (i.e., which rule it was trying to trigger). For some authors, such as Pearl [Pea88], the best explanation is the most probable assignment of values to a set of variables—the process of obtaining this kind of explanation is also called *abduction* [CS94, G98, San91, Shi91]. Other systems try to offer a simple but comprehensible report about the domain and the reasoning [CRG⁺96, CTJ89, HHNF86, LSF88, RP85, Str87, Swa83, WS89, WT92]. Finally, in the most sophisticated methods, explanation constitutes an ‘intelligent’ dialogue with the system user through natural language and by way of interactive methods [CMM94, Caw91, Caw93, Caw94, CGL⁺93, Fie01, Moo94].

3 Properties of explanation

In this section we define ten properties of explanation, which constitute the framework for the analysis of the methods studied in Section 4. Some of these properties are taken from [Wic89], and we have classified them in three categories (see Table 1):

content: **what** to explain;

communication: **how** the system interacts with the user;

adaptation: to **whom** the explanation is addressed.

Please note that these categories correspond to the definition of explanation given in Section 2.2: exposing something (*content*) in such a way that it is understandable (*communication*) for the receiver of the explanation, which implies that he/she improves his/her knowledge about the object of the explanation, and it is satisfactory as far as it covers the receiver's expectations (*adaptation*). We briefly describe them in the following subsections.

3.1 Content

One of the most important aspects of explanation is related to *what* is going to be explained, i.e., what an explanation must include to be understandable to the user. This question has a difficult answer because it depends on several issues such as the focus of explanation, the purpose, the level of detail and causality, which we are going to analyse.

3.1.1 Focus of explanation

In the case of Bayesian networks (and indeed for any expert system), there are three basic issues that can (and must) be explained: the knowledge base, the reasoning process performed by the system to obtain (or not) a conclusion and the evidence propagated, if any. For this reason, we can classify explanation methods in three groups, differing according with the focus taken on explanation:

Explanation of **evidence**: It consists of determining which values of the unobserved variables justify the available evidence. This process is

Content	Focus	evidence / model / reasoning
	Purpose	description / comprehension
	Level	micro / macro
	Causality	causal / non-causal
Communication	User-system interaction	menu / predefined questions / natural language dialogue
	Display of explanations	text / graphics / multimedia
	Expressions of probability	numeric / linguistic / both
Adaptation	User's knowledge about the domain	no model / scale / dynamic model
	User's knowledge about the reasoning method	no model / scale / dynamic model
	Level of detail	fixed / threshold / auto

Table 1: Properties of explanation methods.

usually called *abduction*, and it is based on the (usually implicit) assumption that there is a causal model—see Section 3.1.4. In this context, *an explanation* is a configuration of the unobserved variables, and the goal of the inference process is to obtain the most probable explanation (MPE) or the k most probable explanations. In general, the variables that take the value “present” or “positive” in the MPE are considered to be the causes that *explain* the evidence. The purpose of this kind of explanation is basically to offer a diagnosis for a set of observed anomalies. For instance, in medical expert systems, an explanation consists of determining the disease or diseases that explain the evidence: symptoms, signs, test results, etc.

Explanation of the **model**: This kind of explanation is sometimes called *static* [HD90] and consists of displaying (verbally, graphically or in frames) the information contained in the knowledge base. One of its objectives is to assist human experts in the *construction* of expert systems; for instance, Elvira’s explanation facility has been used to debug HEPAR-II, a medical Bayesian network [LOD01]. The other objective of explaining the model is to offer a novice user some knowledge about the domain for *instructional purposes*.

Explanation of **reasoning**: This is sometimes known as *dynamic explanation* [HD90] and may provide three kinds of justifications:.

The **results obtained** by the system and the **reasoning process** that produced them. During the *evaluation* of the expert system, some of the test cases are properly solved by the system and some of them are not. In the cases properly solved, the evaluators (knowledge engineers and human experts) can check that all the steps in the way to a solution were correct (an incorrect argument might lead to a correct solution by accident). In the

case that the expert system makes a mistake, the explanation of the reasoning process is an invaluable tool for isolating and correcting the wrong pieces of information in the knowledge base. During the *deployment* of an expert system, the explanation capability is very useful for convincing the user of the correctness of the results; in particular, physicians are very reluctant to accept the advice of a machine if they can not understand how it was obtained [TS84]. In the case of *tutoring systems*, the explanation capability is essential for improving the student's skills.

The **results not obtained** by the system, despite the user's expectations. It may also be of interest, especially for educational systems, to explain why the system did not produce a certain conclusion expected by the user; in particular, which findings oppose such a conclusion and which findings would be necessary to support it.

Hypothetical reasoning, i.e., what results the system would have returned if one or more given variables had taken on different values from those observed.

3.1.2 Purpose

There are two different goals an explanation capability might try to accomplish, which leads to two types of explanation:

Description: this kind of explanation consists in showing the underlying *knowledge base* (in the case of model explanation), or providing further details on the *conclusions* or displaying *intermediate results* (in the case of reasoning explanation).

Comprehension: in this case, the explanation tries to make the user to understand the implications of the model or the conclusions of the

system and/or the relation between them. For instance, how each finding affects the conclusion, individually or in conjunction with other findings.

3.1.3 Levels of explanation

Sember and Zukerman [SZ90] proposed another classification of explanation methods for Bayesian networks, thought it might also be applied to other types of expert system:

Micro level: a micro-level explanation consists of a detailed justification of the variations produced in a particular node as a consequence of the variations in its neighbours. In the case of a rule-base expert system, the micro level would consist of analysing the variables contained in a certain rule, or the rules containing a certain variable.

Macro level: this level of explanation analyses the main lines of reasoning (the paths in the Bayesian network) that lead from the evidence to a certain conclusion. A macro-level explanation for a traditional expert system would consist of one or several chains of rules.

3.1.4 Causality

From a mathematical point of view, a BN is just a model for representing probabilistic dependencies and independencies; in this case, a link, considered by itself, has no meaning. However, when a BN is built as a model of a real world system, a link $A \rightarrow B$ is *causal* when A is a cause of B , i.e., when there is a mechanism by which the value taken on by A influences the value of B . A BN is said to be causal when all of its links are causal. The reasons for using causal models in artificial intelligence, especially in probabilistic expert systems, are the following:

Human beings tend to interpret events in terms of cause-effects relations [KST82, PH88]. Therefore, causal models are easier to construct and modify [Hen89, PH88], and also more easily understood by users [Dru93, Sue92].

The identification of *invariant*³ causal relationships in a domain allows the prediction of effects of both spontaneous causes (usually corresponding to random variables) and actions (sometimes called manipulations or interventions) [Pea99].

Causality and probability are closely related, because causality normally implies a pattern of probabilistic inter-dependencies, which provides clues about causality. In fact, a necessary condition for establishing the presence of causality is statistical correlation [Dru93].

In the same line, the axiomatic properties of Bayesian networks (*d*-separation and the Markov property) correspond to probabilistic dependencies and independencies that appear in causal domains [D99].

There exist canonical probabilistic models (noisy OR, noisy MAX, noisy AND, etc.) based on the interpretation of the parents of a node as causes or conditions for that node and on the assumption of independence of causal interactions, although they might be used as mere probabilistic models. These models reduce the number of parameters of the network, simplify the acquisition of knowledge and some of them lead to a more efficient computation [DD00].

Causal Bayesian networks support certain qualitative reasoning patterns, which can be identified in order to explain the results of inference [Dru93, Dru96, HD90]. Some of these patterns are specific of certain

³“Invariant” means that when one mechanism is subjected to changes, the others remain intact.

canonical models; for instance, *explaining away* is a phenomenon typical of the noisy OR.

Finally, the concept of explanation is very closely-tied to the notion of causation; in fact, one of the modalities of scientific explanation consists of finding the causes of the observed facts.

In summary, some of the explanation methods for Bayesian networks are specifically designed for causal networks, or even for specific canonical models, while other methods are general, in the sense that they do not assume a causal interpretation of the network.

3.2 Communication

A crucial aspect of explanation is the way in which it is offered to the user. This depends, basically, on the methods of interaction between the user and the system and the manner in which it is presented. Of particular interest in the case of Bayesian networks is also the way in which probabilities are expressed.

3.2.1 User-system interaction

In some systems, users can request explanations by selecting some variables or options from a certain menu. In other systems, users can pose questions; for instance, in MYCIN the available questions were “how [did the system arrived at a conclusion]” and ”why [is the system asking for this piece of information].” Finally, other systems try to offer a natural language dialogue by analysing and building a model of the conversation, in order to “understand” the context of each question and answer.

On the other hand, some systems only allow the user to ask for an explanation after the program has presented its conclusions, while other

systems allow the user to interrupt the process of inference and request an explanation.

3.2.2 Display of explanations

We can distinguish how explanations are presented to the user, in the following ways:

Verbally, by using text and numbers.

Graphically, in two forms:

by using bar diagrams, pie charts, plots or any other graphical tools that represent the relation between values or variables, the evolution of the probabilities associated to the variables during the evidence propagation, etc.

by using the Bayesian network's own graph of , so that changes in the probability distributions are indicated by adding text, numbers or signs, or by colouring the nodes and/or the links.

Multimedia, integrating the above ways, text, numbers and graphs, in a hypertext and multimedia environment that combines interactive explanations with images, video and sound.

3.2.3 Expression of probability

Expressions of probability may be *numerical* or *quantitative*, such as 0.98 or 72.6%, and *linguistic* or *qualitative*, such as “seldom”, “very likely” or “almost sure”. Similarly, when comparing two probabilities it is possible to have quantitative expressions, such as “the odds between A and B are 10^{-3} ”, or qualitative expressions, such as “ B is much more probable than A .” There is a significant amount of research on the assignment of linguistic

expressions of probability to numeric values and vice versa—see [Dru89, Dru93] for a review, and also Section 4.3.1 in this paper.

3.3 Adaptation

Explanation always means explaining something to *somebody*. Therefore, one of the key features of an effective explanation is the ability to address each user’s specific needs and expectations, which essentially depends on the knowledge he/she has. There are three issues to be considered independently.

3.3.1 User’s knowledge about the domain

Some explanation methods do not take into account the variability of domain knowledge for different users. Therefore, explanations are generated for a hypothetical user having a certain knowledge: some explanation methods are intended for beginners, while others assume the user is an expert.

Other methods rely on a **static user model**, i.e., they do not consider that the user’s knowledge increases as he/she interacts with the system. In many cases, the user model often reduces to using a **scale** with two or three categories: “novice, advanced, expert.” In theory, an expert system could have the capability of classifying users automatically, but in practice, the user generally has to classify him/her in this scale at the beginning of his/her interaction with the system.

Finally, other explanation methods possess a **dynamic model** for each user, which explicitly represents his/her knowledge and changes as he/she proves to have learned new concepts or relations during his/her operation with the expert system. This model allows the explanation module to generate the explanation adapted to the knowledge that the user has at each moment.

3.3.2 User's knowledge about the reasoning method

Analogously, it would be possible for an explanation method to take into account the user's knowledge about the reasoning method employed by the expert system. In the case of a Bayesian network, the explanation generated for a user that is familiar with the concepts of prevalence, prior/posterior odds and likelihood ratios should be very different from the explanation generated for a user who has never heard about them. As in the case of knowledge about the domain, the possibilities rank from not considering the variations in user's knowledge about the reasoning method to having a dynamic model that explicitly represents the user's knowledge at each moment.

3.3.3 Level of detail

The level of detail of an explanation is closely related to the user's knowledge about the domain and about the reasoning method, but can also be established independently. For instance, it is possible for an expert system that does take into account the user's knowledge to assign an importance factor to each item (each rule or each variable, for instance) and an importance threshold, so that the explanation module only displays those items above the threshold. By lowering the threshold, the user can increase the level of detail, and vice versa. This way it is possible to offer different levels of detail without having a user model. Obviously, instead of having a fixed importance factor for each item, it would be desirable to dynamically adjust the importance factors as a function of the user's domain knowledge, the available evidence and the interaction with the system. Similarly, the aspects of explanation related to the reasoning method should also be offered at different levels of detail.

4 Explanation methods for Bayesian networks

In this section we review the techniques developed for generating explanation in Bayesian networks, in the light of the ten properties discussed above. We divide the methods into three groups, according to the focus of explanation (see Section 3.1.1): explanation of evidence, explanation of the model, and explanation of reasoning.

4.1 Explanation of evidence: abduction

We first analyse the methods of abduction, which focus on explaining the evidence. As mentioned in Section 3.1.1, in this context an explanation \mathbf{w} is an assignment of values to all the variables in a certain subset \mathbf{W} of the variables of the network. Since the values of observed variables are known with certainty, only unobserved variables are the object of scrutiny in abductive methods. The goal of abduction is to find the most probable explanation (MPE), i.e., the configuration \mathbf{w} with the maximum a posteriori probability $P(\mathbf{w}|\mathbf{e})$, where \mathbf{e} is the available evidence. Some methods are able to find the k most probable explanations. When \mathbf{W} includes all the unobserved variables, the process is known as *total abduction*; otherwise, it is called *partial abduction*.

In principle, the goal of the methods we review in this section is just to find the MPEs, without trying to justify why they are more probable than the others, i.e., the purpose of these methods is **description**, not comprehension. The distinction between micro and macro **level** makes no sense in this case, because there is no analysis of the reasoning process. The interpretation of a configuration (an assignment of values) as an explanation implicitly assumes that there is a **causal** model, so that the variables that take on the value “present” in the MPE are those that *explain* the observed anomalies. However, since these methods are purely mathematical, they can be applied to any network, independently of whether it is causal or not.

With respect to user-system communication, research on abduction in BNs usually limits itself to finding the MPEs, without paying attention to the interaction with and adaptation to the user. For this reason, most of the criteria for analyzing explanation methods (Section 3) do not apply to abduction. At most we can mention that the probability of each explanation is given **numerically**; no attempt has been done to express qualitatively the probabilities.

We expose in chronological order the most relevant works on abduction. Please note that all these methods only work for discrete variables.

4.1.1 Pearl's π - λ propagation

The first work on explanation in Bayesian networks [Pea88, Chapter5] was a method of total abduction based on the property that, for each value x of a given variable X in \mathbf{W} , there is a best explanation for the rest of the variables, $\mathbf{W} \setminus X$. This property is similar to the principle of optimality in dynamic programming. Therefore, the value of X in the MPE can be obtained by finding the best explanation for $\mathbf{W} \setminus X$ and then choosing the best value of X . Since the computations are local for each node, the process can be designed as an exchange of π - λ messages among nodes. The algorithm has linear complexity for the case of polytrees and exponential complexity for networks with loops. However, the search for the MPE may also be very complicated for polytrees, since even variables that are not interesting for the hypothesis must be taken into account. A shortcoming of this method is that it can only find the two MPEs ($k=2$) [Nea90]. As a consequence, this method is only adequate when the two MPEs are much more probable than the rest of the explanations.

4.1.2 Linear restrictions system

Santos [San91] proposed a method of total abduction that transforms a Bayesian network into an equivalent linear restriction system. A *restriction system* $L(W)$ is a tuple (Γ, I, ψ) where Γ is a set of variables, I is a finite set of inequalities defined over the variables of Γ , and ψ is a function from $\Gamma \times \{\text{true}, \text{false}\}$ to \mathbb{R} . The restrictions guarantee that each variable takes only one value and that the probability of a configuration of a set of variables is calculated by the corresponding set of conditioned probabilities. The construction of the restriction system is performed in linear time with respect to the number of variables of the Bayesian network. The search of all solutions involves solving a sequence of restriction systems in which each system is derived from the previous one. One of the disadvantages of this method is that, like Pearl's, it assigns values to all variables, including those that are not interesting for the hypothesis.

4.1.3 Irrelevance in partial abduction

Shimony [Shi91] and Suermondt [Sue92] (see Section 4.4.1) addressed the problem of obtaining explanations consisting of only relevant variables. In the case of partial abduction, Shimony proposed three definitions of irrelevance:

statistical independence: a variable X should not be part of the explanation \mathbf{w} if it does not affect the probability of the evidence: $P(\mathbf{e}|\mathbf{w}, x) = P(\mathbf{e}|\mathbf{w})$; it is said that X is irrelevant;

δ -independence: it is less restrictive than the former and it indicates that a fact X is irrelevant if given facts (W) are independent of it with a tolerance of δ : $|P(\mathbf{e}|\mathbf{w}, x) - P(\mathbf{e}|\mathbf{w})| \leq \delta$, and

quasi-independence, which means that a parent node is relevant only if its contribution to the probability of the node it explains is greater than

its own prior probability.

The algorithm for finding relevant explanations relies on the fact that the ancestors of a node V that have not been observed must remain without assigning them any value if they do not affect the probability of V (because as we have mentioned they can not explain V). Furthermore, if a node V is not an ancestor of node N , then V can not explain N because it is not a possible cause of N . Please note again the assumption of a causal network in this method.

4.1.4 Graphs of weighted Boolean functions

Charniak and Shimony [CS94] proved that the problem of abduction in BNs is equivalent to finding the minimum cost assignment for the variables of a certain *weighted Boolean function acyclic directed graph* (WBFADG) that results from a transformation of the Bayesian network together with evidence e . The WBFADG is solved by a best-first algorithm that, applied successively, produces the enumeration of the assignments in increasing order of cost, i.e., the explanations in decreasing order of probability. Two disadvantages of this method are its computational and conceptual complexities.

4.1.5 Approximate partial abduction

Finally, Gámez [G98] has made a detailed study about abduction, both total and partial. Among his most relevant contributions are:

he has studied how to adapt clustering algorithms (clique tree propagation) for partial abduction;

he has developed new approximate algorithms that can be used in those cases in which exact algorithms are inefficient. Those algorithms are based on *genetic algorithms* and *stochastic annealing* [dCGM99],

he has proposed some criteria to simplify explanations [dCMG99]. Those are:

normality. The simplification is based on showing the user only those values that are not the *usual* value of the variable. For example, in medicine, the usual value for the variable `Meningitis` is `absent`.

probabilistic independence, by analyzing the graph that represents the Bayesian network. He has developed an algorithm for detecting some kinds of independencies.

relevance. This implies that the omitted information is “almost” irrelevant (in the sense of statistical independence) for the observed facts.

he has proved that the criteria based on relevance and independence produce better results in those cases in which is not easy to decide which is the *usual* value of a variable.

4.2 Explanation of the model

In Section 3.1.1 we mentioned the advantages that the explanation of the model, also known as *static explanation*, presents in the phase of building the Bayesian network and also with educational purposes. We classify the methods according to the way they present their explanations.

4.2.1 Graphical display of the network

The most direct and intuitive way of showing the information embodied in a Bayesian network is to display the corresponding graph: each node is represented by an oval containing the name (or a short description) of the associated variable, and links are drawn as arrows. Some of the earlier tools that permitted the graphical edition and visualization of Bayesian networks are Analytica, IDEAL [SB90], DAVID [Sha86a], HUGIN [AOJJ90]

and PATHFINDER [Hec91]. However, when the size of the network grows large, it becomes more and more difficult to read the graph. For this reason, other packages, such as Analytica and GeNIe [Dru99], offer the possibility of defining submodels. When contracted, a submodel is represented by a special type of node; when expanded, it displays all the nodes and the links it contains.⁴

In any case, this kind of explanation is only a **description** of the model, and it is not intended to improve **comprehension**. The interaction with the user is done by means of **menus**, which give access to the properties of nodes, links and conditional probability tables. The parameters of the model appear in **numerical** form. In the tools we have examined, there is no possibility of **adaptation** to the user, except for minor formatting options.

4.2.2 Verbal description of the network

Druzdel [Dru93, HD90] proposed a method for translating the qualitative and quantitative information of a Bayesian network into linguistic expressions. There are patterns for indicating the prior probability of a node, such as “Cold is very unlikely (p=0.08)” and for comparing probabilities: “Cold is slightly less likely than [having a] cat (0.08/0.10).”.

In this method, **causal relations** play an important role. In particular, there are explanation patterns for describing the leaky noisy OR, a model that assumes the independence of causal interactions: “Cold very commonly (p=0.9) causes sneezing. Allergy very commonly (p=0.9) causes sneezing. Cold does not affect the tendency of allergy to cause sneezing, and vice versa. There are also other unlikely (p=0.1) causes of sneezing.” Relations of conditional dependence or indepen-

⁴The URL's for HUGIN, Analytica, GeNIe and other Bayesian network tools can be found at <http://www.ia.uned.es/~fjdiez/bayes/#software>.

dence can be expressed by sentences like this: “Given sneezing, cold and allergy are independent.”

The purpose of this explanation model lies between **description** and **comprehension**. Explanations are offered at the **micro level**. User-system **interaction** is not described in the references. The presentation of explanation is in the form of **text** containing **linguistic** and **numeric** expressions of probability.

Other systems, like B2 (see Section 4.4.1) and Elvira (see Section 4.2.4) are also capable of offering verbal explanations of the model.

4.2.3 Menu-driven navigation

Díez [D94], in his expert system DIAVAL, developed a method for explaining both the model and the reasoning. His method distinguishes several types of links and nodes corresponding to the different types of **causal** influence: in the noisy OR gate, the parents of a node are considered as *causes* in the strict sense, while in the general model parents are regarded as *factors* that influence the child variable.

User-system interaction is based on a system of windows and **menus**, which offer different options:

For each *node*, the user can view: definition, prevalence (prior probability), posterior probability, list of causes or factors, conditional probability table (for nodes with parents), list of effects (children) and definition.

For each *parameter*, which is a node with an associated measure (generally from echocardiography), the user can see: measured value, intervals, pathophysiological meaning, and formula (for those parameters calculated from others).

For the *links* that make part of a noisy OR the options are: cause, effect, sensitivity and specificity (only for binary variables), efficiency (the

probability that the cause produces the effect) and a text that explains the associated causal mechanism.

By visiting the parents (causes) and children (effects) of nodes, the user can navigate across the network. Explanation fits into the **micro level**, because explanation focuses on one node or link at each moment. Explanations are displayed as **text** inside different windows, and probability is only expressed **numerically**.

4.2.4 Static explanation in Elvira

Elvira⁵ is an environment for the edition and evaluation of Bayesian networks and influence diagrams, developed as a research project among several Spanish universities. Following Druzdzel's proposal [Dru93] (see above), Elvira offers **verbal** explanations at the **micro level**. The main difference with his method is that in Elvira the nodes are classified into several categories, such as *symptom* or *disease*. This information is used to generate model explanations. An example of an explanation could be: "The disease X has the following symptoms: s1, s2, ...". It is also capable of offering **graphical** explanations, such as nodes expansion, as well as verbal information about specific nodes or links. In a similar way to DIAVAL [D94], it also allows to navigate across the network. The interaction with the user is made by way of **windows** and **menus**.

Another useful option consists of automatically colouring the links of the network, in order to offer qualitative insight about the conditional probability tables [LOD01]. More specifically, given two ordinal discrete variables A and C such that there is a link $A \rightarrow C$, this link is said to be **positive** if higher values of A lead to higher values of C for any configuration of B , where B represents the set of the other parents of C . The definition of **negative link** and **null link** are analogous. When the influence is neither

⁵Information about Elvira can be found at <http://www.ia.uned.es/~elvira>.

positive nor negative nor null, then it is said to be **unknown** [Wel90a]. Typical orderings of values of a variable are $+a > -a$, *present* > *absent*, *severe* > *moderate* > *mild* > *absent*, *positive* > *negative*, etc. If A and C are binary variables, the above definition implies that link $A \rightarrow C$ is positive if and only if $P(+c|+a, b) > P(+c|-a, b)$ for each value b of B . If variable A represents a cause or a risk factor for C , or C is a test that detects A , then influence $A \rightarrow C$ is in general positive. In causal networks, most of the links are positive.

4.2.5 Assisted construction of medical Bayesian networks

We conclude this section on static explanation by mentioning the system MEDICUS, developed by Folckers et al. [FMST96, SCFT96] as a tool for the construction of explanation models in which the knowledge is complex and uncertain, as in medicine. The main contribution is that it is a shell in which the user can develop a Bayesian network for representing a certain domain, independently of his/her **level of knowledge** about probability. This is done by means of **micro-level** explanations generated by a linguistic editor of the model, together with a **graphic** editor of the Bayesian network. The explanations are expressed both **quantitatively** and **qualitatively** and are presented **verbally** and **graphically**.

4.3 Explanation of reasoning at the micro level

In this section we describe the methods used to explain the reasoning at the micro level in Bayesian networks. They are characterised by the fact that at each moment they focus only on one variable, usually called *focal hypothesis*, for generating the explanations.

4.3.1 Verbal description of variations in probability

Experimental research showed that human beings understand linguistic expressions of probability better than numerical data. For this reason, Elsaesser [EH89] proposed a method for generating linguistic explanations by means of a template based on Polya’s *shaded inductive patterns*, which are a set of heuristic rules for describing changes in probabilities.

For example, one of those rules, compatible with Bayes formula is: “*If $A \rightarrow B$ and B is true, then the existence of A is more credible.*” The template is filled with natural language terms that represent probabilistic data and their variations. The set of expressions denoting probability vary between those that reflect high probabilities, like *almost certain* (for values between 0.99 and 0.91) or *highly probable* —for the interval [0.90, 0.82)—to those that designate small values, like *improbable* (for the range 0.18 to 0.09) or *highly improbable* (for the range 0.08 to 0.01). There are also expressions for defining changes from probability p_1 to probability p_2 ; for instance, when $p_2/p_1 \geq 5$, the expression is “a great deal more likely”; when $2.5 < p_2/p_1 \leq 5$, the expression is “much more likely,” etc.

Later, an experiment by Elsaesser and Henrion [Els90] proved that the linguistic expressions assigned by people to different pairs (p_1, p_2) depends on the difference $p_2 - p_1$ more than on the ratio p_2/p_1 or the odds ratio $(p_1/(1 - p_1))/(p_2/(1 - p_2))$.

Clearly, the translation of numerical probabilities into linguistic expressions makes **no** assumption on **causality**. In this method, intended for users that are not familiar with probability theory, there is no explicit **user model**.

4.3.2 Graphical display of probabilities

There are several software tools for processing Bayesian networks that offer the possibility of showing graphically the variations in probability of certain

variables, such as HUGIN, Netica or Elvira. The basic idea common to all of them is to show the variations of probability by plotting a bar proportional to the probability of each state of a node, together with its numerical value. The purpose of this kind of explanation facility is the **description** of the reasoning process and it is useful both for **causal** and **non-causal** Bayesian networks. Probabilities are expressed **quantitatively** and **qualitatively**.

However, Elvira differs from other tools in that it is able to simultaneously store and display several evidence cases [LAD00, LOD01]. The user is allowed to navigate across the set of evidence cases, saving them in files, generating new cases, expanding or contracting the selected nodes, modifying the inference options, etc. This facility permits the user to observe the variations in the probability of each variable due to the sequential introduction of new findings, and to perform “what-if” analysis by introducing different hypothetical findings, even for the same variable. Moreover, when evaluating the impact of evidence, the user can know if the probability of a certain node has increased or decreased with respect to the previous case or to another fixed case, selected by him/her, because nodes are coloured depending on the changes on its probabilities.

The main objective of Elvira’s graphical explanation facility is the **comprehension** of the reasoning process and it is useful both for **causal** and **non-causal** Bayesian networks. The interaction with the user is performed by **windows** and **menus**. Probabilities are expressed **quantitatively** and **qualitatively**, by means of coloured bars. There is a rudimentary adaptation capability consisting in controlling the expansion and importance thresholds.

4.3.3 Explanation of local updates in polytrees

In a polytree (a network without loops), the posterior probability for variable B can be obtained by [Pea88]

$$Bel(b) = \alpha \lambda(b) \pi(b) \tag{2}$$

where $\lambda(b)$ is the support from the effects (the children) of B and $\pi(b)$ is the causal support. Sember and Zukerman [SZ90] developed an explanation method for justifying the value of $Bel(b)$ in terms of $\pi(b)$ and $\lambda(b)$. The user's expectations are featured by the changes produced in $\pi(b)$, since it represents causal information. If $\pi(b)$ increases, decreases or does not change, then the reader expects $Bel(b)$ to increase, to decrease or remain unchanged, respectively. When the expectation is met, their method generates the following explanation: “*The belief in B has augmented due to an increase in its causal support.*” In those cases when the user's expectation is not met, the explanation consists of identifying which values of $\pi(b)$ and $\lambda(b)$ caused the deviation.

The purpose of this method is **comprehension** and it assumes a **causal** Bayesian network. Explanations are presented as **text** and variations in the probability are expressed **linguistically** (“has augmented/has decreased”). Sember and Zukerman do not describe the user-system **interaction** nor any possibility of **adaptation**. The explanations require the user to be familiar with the reasoning method.

This method is direct and intuitive in the case of binary variables, but becomes more complicated for multivalued variables. The main limitation of this method is that it only works for polytrees, and therefore can seldom be applied to real-world problems, whose models almost always have loops.

4.3.4 Analysis of local variations of probability in DIAVAL

As we said in Section 4.2.3, the expert system DIAVAL [D94] had a method for explaining the reasoning, which allows the user to select each diagnosis and open a menu with different options, including the possibility of visiting the parents and children of the nodes, as described in Section 4.2.3. This way the user can investigate which neighbours of a node have increased or decreased their probability. Additionally, if a node Y is the child of a noisy OR/MAX gate, the probability that each parent X_i has produced Y can be computed as $P(+x_i|\mathbf{e}) \times c_i$, where c_i is a parameter of the OR gate that represents the probability that X_i , when being present, produces Y .

The purpose of this explanation method for **causal** Bayesian networks is mainly to **describe** the results of inference and includes specific explanation patterns for the noisy OR/MAX gates. User-system **interaction** is implemented by a system of windows and menus. Probabilities are expressed **numerically**. There is no user model and no adaptation.

4.3.5 Analysis of the impact of evidence on a variable

PATHFINDER One of the first systems that included some kind of explanation was PATHFINDER [Hec91]. The method basically consists in discriminating between two diseases or two groups of diseases, by showing how each value v_i of a certain variable V affects its probability distribution. To do this, the system draws a graphic bar proportional to the likelihood ratio for each value v_i of the selected variable in favor of one disease D_1 relative to the other D_2 given the evidence e . The measure unit used is the *evidence weight* proposed by Good [Goo77]

$$\log \left(\frac{P(v_i|D_1, e)}{P(v_i|D_2, e)} \right) \quad (3)$$

The purpose of PATHFINDER explanations is the **description** of reasoning and can be applied to both causal and non-causal networks. The

presentation of the explanation is **graphical** and the interaction with the user is led by way of menus, windows and dialogue boxes. Probabilities are expressed by **numbers** and there is no adaptation and the user must have some knowledge of probability theory.

INSITE Another method was Suermondt’s INSITE [Sue92, SC93], whose main objective is to identify the findings that influence the posterior probability of a certain hypothesis, as well as the paths through which the evidence flows, which we will describe in section 4.4.1. The influence of evidence \mathbf{e} on a certain variable D is measured by a cost-function. Suermondt [Sue92] analyses different cost functions and concludes that the most suitable function is cross-entropy,

$$H(P(D|\mathbf{e}); P(D)) = \sum_i \left[p(d_i|\mathbf{e}) \lg \left(\frac{p(d_i|\mathbf{e})}{p(d_i)} \right) \right], \quad (4)$$

where d_i represents the possible values of D . The influence of individual findings or subsets of findings is determined by a *sensitivity analysis* which computes the cost of omission of such findings.

The purpose of INSITE is the **comprehension** of reasoning. and can be applied to both causal and non-causal networks. Interaction with the user is led by way of menus, windows, dialogue boxes, buttons and by selecting nodes or arcs in the network display. Explanations are presented as a combination of **graphics and text**. In the latter case, the probabilities are expressed both by **numbers** and by **linguistic expressions**.

4.4 Explanation of reasoning at the macro level

The methods analysed above are characterised by generating macro level explanations, which typically implies that they attempt to explain the main paths in the Bayesian network through which the evidence flows.

4.4.1 Quantitative analysis of reasoning chains

NESTOR NESTOR, developed by Greg Cooper [Coo84], is one of the first Bayesian networks. It included a facility for explanation which offered two possibilities: to *compare* how two diagnostic hypothesis account for the evidence, and to *critique* a hypothesis with respect to all other possible diagnoses. In both cases, it could generate a **verbal** explanation for describing **qualitatively** the chains among the evidence and a given hypothesis, which basically consisted of translating into English the causal links of each chain. On the other hand, in the case of the “compare” command, the user could visualize, how each finding affected the relative probability of two hypothesis selected by the user. In the case of the “critique” command, the user could visualize the relative probability of the hypothesis with respect to all other hypothesis. The explanation method was defined only for **causal** networks. Explanations were displayed both **graphically** and **verbally**. Probabilities were expressed **numerically**. The interaction with the user was lead by way of commands (predefined questions). It had no adaptation capability and no user model.

Graphical display of the weight of evidence In the same line, Madigan et al. [MMA97] propose a **graphical method** that consists in showing how the evidence is propagated through a causal Bayesian network. The impact of the evidence on a binary node H is measured, as in PATHFINDER, by computing Good’s *evidence weight*

$$W(H : \mathbf{e}) = \log \left(\frac{P(\mathbf{e} | +h)}{P(\mathbf{e} | -h)} \right) \quad (5)$$

and it is displayed on the network’s own graph by varying the colour and thickness of nodes and links, by drawing different kind of links, etc.

The method also provides macro explanations for two types of questions. The first one is: “What is the relative importance of each finding f about

the variable of interest H ?” This is answered by visualizing $W(H : f)$ in a graphical evidence balance sheet, in which all findings are shown together with their importance for the hypothesis. Please note that the importance of each finding depends on the order in which the user sets their values.

The second type of question is: “Why does a concrete node have so much influence on a given variable?” The answer consists of showing the relevant paths, in a way similar to Suermondt’s *chains of reasoning*.

The purpose of this method is the **comprehension** of the reasoning process. Although Madigan et al. speak of **causal** Bayesian networks, their method also works in non-causal networks. User-system interaction is **menu-driven** and probabilities are **qualitatively** codified by means of colours and thicknesses. No adaptation capability has been developed for this method.

INSITE Suermondt’s method (see Section 4.3.5) also looks for relevant *chains of reasoning*, i.e., paths from \mathbf{e} to D that are computationally related to D given evidence \mathbf{e} . He analyses the *strength*⁶ of each of them and whether it conflicts or not with the inference result. The chains are presented graphically to the user by shading the conflicting ones and highlighting the consistent ones. Chains can also be described verbally. Moreover, INSITE defines a method for detecting conflicts among different variables.

As we said, the purpose of INSITE is the **comprehension** of reasoning and the method can be applied to both causal and non-causal networks. Explanations are presented as combinations of **graphics and text**, expressing probabilities by means of **numbers** and **linguistic expressions**. The method has no user model but the level of detail can be adapted to the needs of a given user.

One interesting aspect of INSITE is its independence of the algorithm

⁶The strength of a chain is viewed as the importance of the chain in obtaining the inference result.

of evidence propagation. In contrast, its main shortcoming is the computational complexity, that grows exponentially with the number of nodes and arcs.

BANTER Based on Suermondt’s method INSITE, Haddaway, Jacobson and Kahn [HJK94a, HJK94b, HJK97] have developed BANTER, a tool for decision-support and training. The tool is aimed at medicine, although it works on any network consisting of hypotheses, observations and diagnostic tests. Given evidence, BANTER can offer the probability of a hypothesis or select the most informative test for confirming or discarding a hypothetical diagnosis. When used as an educational tool, BANTER randomly generates scenarios and asks the user to make a diagnosis or select a test.

BANTER can also generate verbal explanations (by using a modification of Suermondt’s method) by identifying the most influential pieces of evidence and by selecting the strongest⁷ and shortest paths between the evidence and the hypothesis.

Both methods share the purpose of **comprehension** of the reasoning process. They differ in that BANTER is designed for **causal** Bayesian networks and only generates **verbal** explanations, while INSITE can also display graphic explanations. In BANTER, the expressions of probability are only **quantitative**. None of them offers any adaptation to the user. As BANTER is especially conceived for novice or inexperienced users, it does not require any knowledge of Bayesian networks, but only some familiarity with the application domain and an elementary understanding of probability.

B2 However, there are some problems when BANTER is used as an educational tool: it is not capable of making hypothetical reasoning (see Section 3.1.1) and it sometimes offers irrelevant information for the conclusion,

⁷The strength of a path C given evidence e is defined as $Strength(C) = \min |(P(n) - P(n|e))|, \forall \text{ node } n \in C$

which can be misleading for the user. On the other hand, verbal explanations are difficult to understand. In order to solve these deficiencies, McRoy et al [MHA97a, MHA97b, MLPHH96] have developed B2 as an extension of BANTER. It generates both **graphical** and **verbal** explanations in natural language, in a more consistent way than BANTER does. It generates explanation of the **reasoning**. This is done by the representation of a discourse model by way of a propositional semantic network. This allows the system to generate only relevant information, since it is capable of doing some reasoning about the content and the structure of its interaction with the user. Therefore, the objective is **comprehension** of the **reasoning**. It also offers the capacity for doing **hypothetical reasoning**. As in BANTER, the probabilities are expressed **quantitatively** and there is no user **adaptation**. User-system **interaction** is performed by windows and questions are expressed in **natural language**.

Probability of evidence Dittmer and Jensen [DJ97] focus on analysis tools for generating macro explanations in Bayesian networks, but instead of analysing the paths in the network, they compute the probability of evidence in order to detect conflicts among data, and perform a *sensitivity analysis* to determine how changes in the evidence affect the conclusion. They have applied them to the BOBLO Bayesian network for determining if the pedigree assigned to the cattle is correct. The main objective of this method is **description** and it is useful both for causal and non-causal models. We have found no reference in relation to user-system interaction or the possibility of adaptation.

4.4.2 Qualitative explanations

Druzdzel and Henrion [Dru93, Dru96, DH93b, DH93a, HD90] also proposed another explanation method similar to the previous one but it is based on

the *qualitative* analysis of the reasoning chains. It consists of transforming a causal Bayesian network into a qualitative probabilistic network (QPN) [Wel90a, Wel90b], in which the relation between two adjacent nodes is denoted as positive (+), negative (−), null (0) or unknown (?); there are also relations that involve more than two nodes, such as additive [Wel90a, Wel90b] or multiplicative synergies [DH93b, DH93a]. The main advantage of QPN’s is that they simplify the construction of models, because they do not require the elicitation of numerical parameters; as a consequence, their main disadvantage is the lack of precision in the results, especially because very often the combination of “positive” and “negative” influences leads to “unknown” influences. The motivation for this explanation method is that people usually reason and explain their reasoning in terms of qualitative relations.

The qualitative propagation algorithm proposed by Druzdzel and Henrion [DH93b] consists in exchanging messages among neighbour nodes. Evidence nodes are marked with + or −. Unobserved nodes are initially marked with 0. The sign of the message (+, −, 0 or ?) from X_i to X_j is determined by the product of the sign of X_i and the sign of the link.

The objective of explanation is to determine the qualitative impact that each finding f has produced on a certain variable of interest V and to find out the active paths from F to V . There are three kinds of elementary explanations corresponding to the three kinds of elementary qualitative inferences:

predictive inference goes from causes to effects, i.e., in the direction of the links; the explanation for this kind of inference relative to a link $A \rightarrow B$ is of the type “ A may cause B .”

diagnostic inference goes from effects to causes, i.e., in the opposite direction of the arcs; the explanation may be “ B is evidence for A .”

intercausal inference which analyses the qualitative impact of the evi-

dence for a variable A over another variable B when both have an influence on a third variable C , about which there is independent evidence (for example, C has been observed); the explanation generated in this case is of the kind: “ A and B may each cause C ; as A explains C , there is evidence against B .”

If there are several findings and several target variables, the process will be repeated several times. Similarly, if there are several active paths from finding f to target variable V , explanation is sequenced, since both human reasoning and natural language are essentially sequential.

The purpose of this method is the **comprehension** of the reasoning process. It assumes that the network is **causal**. Explanations are displayed as **text**, which includes **linguistic** expressions of probability. As in the case of scenarios, the authors do not describe the user-system **interaction** nor any possibility of **adaptation**. In principle, these qualitative explanations do not require that the user to be familiar with probabilistic reasoning, although knowledge on the propagation method may help to understand the explanations.

Trade-offs resolution It is also worthwhile to mention the work done by Renooij and Van der Gaag et al. [RvdG99, RvdGPG00] in order to improve the ambiguous results that can be obtained when doing inference in a qualitative network. They have designed a formalism for trade-off resolution without making use of probabilistic information [RvdG99]. It is called *enhanced qualitative networks* and it is based on distinguishing between strong and weak influences, in contrast with qualitative networks which do not make this distinction. The way they do this is associating a relative strength with influences. The algorithm for sign-propagation in this kind of networks is a generalization of the one for qualitative networks although its main difference is that strong influences dominate over conflicting weak

influences.

Together with Parsons and Green [RvdGPG00], they have also designed other algorithm for calculating informative results for the node of interest, in order to solve its ambiguous sign after doing inference. It is based on focusing on the subnetwork formed by the chains between the observed node and the node of interest. Then, an informative result is constructed for the *pivot node*⁸, because its sign determines the sign of the node of interest. The ambiguity at the pivot node is solved in terms of the relative strengths of the influences between them as well as the signs of the node’s resolvers.

4.4.3 Scenario-based explanation

According to Druzdzel and Henrion [Dru93, DH90, HD90], a *scenario* is an assignment of values to variables that are relevant to a certain conclusion, ordered in such a way that they form a coherent story—a causal story, if possible—compatible with the evidence. An example of a scenario could be: *Age 5, Fever high, Exanthema present*. The use of scenarios is funded on psychological studies [KST82, PH88] showing that humans tend to interpret and explain processes by weighing up the most credible stories that include the hypothesis being demonstrated.

Although a scenario may contain all the nodes in the network, it is more reasonable to include only those nodes that are *relevant* for a certain task [DS94, LD97, Lud98, SGP97]. If there is a certain focal hypothesis H selected by the user, the relevant nodes are those that affect the posterior probability of H given the observed evidence \mathbf{e} . Otherwise, the relevant nodes are all those whose probabilities depend on \mathbf{e} . After selecting the relevant nodes, the scenarios are generated by some of the methods of partial abduction—see Section 4.1. The explanation consists of showing the

⁸The pivot node is a node that separates the part of the relevant network that contains the trade-offs from the part that does not.

evidence, the most probable scenarios compatible with the hypothesis and those incompatible with the hypothesis, and a comparison of the probabilities of the most probable scenarios.

The purpose of scenario-based explanation is the **comprehension** of the reasoning process, even though the propagation algorithm is not based on scenarios. The method is mainly designed for **causal** Bayesian networks. Explanations are presented as a **text** that describes each scenario in natural language. The probabilities of scenarios are given **numerically**. Druzdzel and Henrion do not describe the user-system **interaction** nor any possibility of **adaptation**. In principle, these explanations do not require that the user to be familiar with probabilistic reasoning, although knowledge of the methods involved will certainly help to understand the explanations.

4.5 Other research

Finally, we describe another methods related to the generation of explanations for Bayesian networks.

4.5.1 Selection of diagnoses

Most of the Bayesian network tools limit the process of inference to displaying the posterior probabilities of the values of each variable. This is clearly insufficient for practical expert systems, especially when the network contains a great number of nodes. For this reason, the explanation capability of the expert system PATHFINDER [Hec91] (see Section 4.4.1) basically consists in showing only how the evidence affected all the possible diseases. For this purpose, an ordered list with all the possible diseases and its associated probabilities is displayed. These probabilities define the order of the list. The main goal of this method is **description**. The presentation of the explanation is **verbal** and the interaction with the user is led by way of menus, windows and dialogue boxes. Probabilities are expressed by **numbers** and

there is no user adaptation.

Similarly, DIAVAL [D94] also implemented a method for selecting the most probable and relevant diagnoses. Relevance in DIAVAL is a measure of the importance from the medical point of view: for instance, a disease like mitral stenosis is more “relevant” than an intermediate pathophysiological variable, like left atrium hypertension; during the construction of the Bayesian network, each node V is subjectively assigned two factors: the positive (negative) relevance factor applies when the most probable value is $+v$ ($-v$). The system only displays the diagnoses whose posterior probability and relevance exceed both the *certainty threshold* and the *relevance threshold*.

The purpose of this method is to **describe** the results of inference although it offers some rudimentary assistance for the **comprehension** of reasoning. The method assumes that the model is **causal**. The **interaction** among the user and the system is based on windows and menus. Probabilities are expressed **numerically**. There is no user model, and the only adaptation capability is that the user can control the **level of detail** in the selection of diagnoses by modifying the relevance and certainty thresholds.

4.5.2 Nice Argument Generator

Among the ongoing projects, the current research of Zukerman, Conachy and Corb [MKZ98, ZKM98, ZMK98] stands out. Their method focuses on the use of Bayesian networks to generate arguments expressed in natural language. In their latest investigations, they propose the utilization of Bayesian networks as the basis of the NAG system (Nice Argument Generator), which serves to analyse and generate ‘good enough’ reasoning to convince the user. So, given a goal proposition introduced by the user, the context in which it is made, and a grade of credibility, the system generates the arguments

to justify it; moreover, it analyses the reasoning provided by the user and prepares other arguments to refute it when necessary. Both the user model and the normative model are represented by Bayesian networks. The new arguments are added to the normative model when they are created. Since both networks can be too complex, there is a mechanism to focus on the relevant subnetworks, over which reasoning will be performed. The intersection of both network structures defines the reasoning graph. In order to analyse it, a propagation over both networks is made, using the probabilistic information associated with each model.

5 Conclusions

Explanation of reasoning is one of the key factors for the success of expert systems. As we discussed in Section 3.1.1, explanation capability is crucial for debugging a model, for convincing the user that the results are correct, and for educational purposes.

There are two main approaches for building expert systems. The *heuristic approach* tries to mimic the reasoning process of human experts, generally by using rules and structured objects. In contrast, the *normative approach* is based on the probability and decision theory; in practice, it amounts to using Bayesian networks for diagnostic expert systems and influence diagrams for decision-support systems. There are both theoretical and empirical studies indicating that the normative approach leads to more accurate and robust expert systems, but on the other hand the explanation capability is even more necessary, because normative reasoning methods are more foreign to human beings than heuristic methods.

In neither of the approaches are there explanation methods satisfactory for the end-user. In fact, there are only isolated proposals, partial solutions insufficient to constitute a standard method suitable for all the expert systems that use similar reasoning techniques.

In this paper (Section 3) we have discussed the basic properties of explanation methods and summarized them in Table 1. In the light of them, we analysed the explanation methods brought up to date for Bayesian networks by considering three kinds of explanation: explanation of evidence (Section 4.1), explanation of the model (Section 4.2), and explanation of the reasoning (Sections 4.3 and 4.4). The first kind correspond to the philosophic concept of explanation as a search for the causes of a certain phenomenon. The second and third kinds of explanation correspond to the concept of explanation used in artificial intelligence: the processes that help the user understand the model and the performance, respectively, of an expert system.

If we come back to Section 3 and look at Table 1 as a checklist of the features that an explanation capability should offer, we realize that all the methods currently available suffer from serious limitations. A significant part of the research limits itself to theoretical models that have not even been implemented in prototype applications. And in the methods implemented so far, the interaction with the user amounts, in the best cases, to displaying a menu with a few options. No explanation method has even tried to offer the possibility of a dialogue with the user.

In the same way, no explanation method for Bayesian networks takes into account the user's knowledge, not even by using the distinction between novice and experienced users, or the possibility of fixing a threshold that controls the level of detail. (In heuristic expert systems, there are several explanation methods that used either a novice-advanced-experienced scale or a dynamic model that explicitly represents the user's knowledge at each moment.)

The conclusion of this paper is that there is a lot of research to be done in the area of explanation in Bayesian networks. We trust that this paper may stimulate some researchers to build on the current methods (by remedying

the shortcomings of one of them or by combining different methods) or to explore new approaches aimed at the problems thus far unaddressed. In our opinion, the most promising lines are the study of causal methods—in particular, canonical models—and the application of user models, both for the domain knowledge and the knowledge about the reasoning method.

6 Acknowledgements

This work has been partially supported by the Spanish CICYT under project TIC-97-1135-C04. We thank Marek Druzdzel and the anonymous reviewers of The Knowledge Engineering Review for their comments on this paper.

References

- [AOJJ90] S.K. Andersen, K.G. Olesen, F.V. Jensen, and F. Jensen. HUGIN. A shell for building belief universes for expert systems. In *Readings in Uncertainty*, pages 332–337. 1990.
- [BS84] B.G. Buchanan and E.H. Shortliffe, editors. *Rule-based expert systems: The MYCIN experiments of the Stanford Heuristic Programming Project*. Addison-Wesley, Reading, MA, 1984.
- [Caw91] A. Cawsey. Generating interactive explanations. In *Proceedings of the Ninth National Conference on Artificial Intelligence (AAAI-91)*, pages 86–91, Anaheim, California, 1991.
- [Caw93] A. Cawsey. User-modelling in interactive explanations. *Journal of User Modelling and User Adapted Interaction*, 3(3):221–247, 1993.
- [Caw94] A. Cawsey. Developing an explanation component for a knowledge-based system: Discussion. *Expert Systems with*

- Applications*, 8(4):527–531, 1994. Special issue on explanation.
- [CGH97] E. Castillo, J. M. Gutiérrez, and A. S. Hadi. *Expert Systems and Probabilistic Network Models*. Springer Verlag, New York, 1997. Versión española: *Sistemas Expertos y Modelos de Redes Probabilísticas*, Academia de Ingeniería, Madrid, 1997.
- [CGL⁺93] A. Cawsey, J. Galliers, B. Logan, S. Reece, and K. S. Jones. Revising beliefs and intentions: A unified framework for agent interaction. In A. Sloman et al, editor, *Proceedings of the Conference of the Society for Artificial Intelligence and Simulation of Behaviour*, pages 130–139, Amsterdam, 1993. IOS Press.
- [Cla93] W. J. Clancey. Notes on “Heuristic classification”. *Artificial Intelligence*, 59:191–196, 1993.
- [CMM94] G. Carenini, V. Mittal, and J. Moore. Generating patient-specific interactive natural language explanations. In *Proceedings of 18th Symposium on Computer Applications in Medical Care (SCAMC 94)*. McGraw-Hill, 1994.
- [Coo84] G. Cooper. *NESTOR: A computer-based medical diagnostic aid that integrates causal and diagnostic knowledge*. PhD thesis, Stanford University, Stanford, CA, 1984.
- [CRG⁺96] B. De Carolis, F. De Rosis, F. Grasso, A. Rossiello, D. Berry., and T. Gillie. Generating recipient-centered explanations about drug prescription. *Artificial Intelligence in Medicine*, 8:123–145, 1996.
- [CS94] E. Charniak and S.E. Shimony. Cost-based abduction and MAP explanation. *Artificial Intelligence*, 66(2):345–374, 1994.

- [CTJ89] B. Chandrasekaran, M.D. Tanner, and J.R. Josephson. Explaining control strategies in problem solving. *IEEE Expert*, 4:9–24, 1989.
- [D94] F.J. Díez. *Sistema Experto Bayesiano para Ecocardiografía*. PhD thesis, UNED, Department of Informática y Automática, 1994.
- [D99] F. J. Díez. From causal graphs to Bayesian networks. In *Proceedings of the Workshop on Conditional Independence Structures and Graphical Models*, Toronto, Canada, 1999.
- [dCGM99] L. M. de Campos, J.A. Gámez, and S. Moral. Partial abductive inference in bayesian belief networks using a genetic algorithm. *Pattern Recognition Letters*, 20:1211–1217, 1999.
- [dCMG99] L. M. de Campos, S. Moral, and J.A. Gámez. Simplifying explanations in bayesian belief networks. Technical Report DECSAI-99-02-01, Dept. Ciencias de la Computación e Inteligencia Artificial, Universidad de Granada, Spain, 1999.
- [DD00] F. J. Díez and M. Druzdzel. Canonical probabilistic models. Part I: Knowledge engineering. Technical Report, Decision System Laboratory, University of Pittsburgh, 2000. In preparation.
- [DH90] M. Druzdzel and M. Henrion. Using scenarios to explain probabilistic inference. In *Working Notes of the AAAI-90 Workshop on Explanation*, pages 133–141, Boston, MA, July 1990.
- [DH93a] M. Druzdzel and M. Henrion. Belief propagation in qualitative probabilistic networks. In N. Piera Carrete and M.G. Singh, editors, *Qualitative Reasoning and Decision Technologies*, pages 451–460. CIMNE, Barcelona, Spain, 1993.

- [DH93b] M. Druzdzel and M. Henrion. Efficient reasoning in qualitative probabilistic networks. In *Proceedings of the 11th National Conference on Artificial Intelligence*, pages 548–553, Washington, July 1993.
- [DJ97] S.L. Dittmer and F.V. Jensen. Tools for explanation in Bayesian networks with application to an agricultural problem. In *Proceedings of the First European Conference for Information Technology in Agriculture*, Copenhagen, Denmark, 1997.
- [Dru89] M. Druzdzel. Verbal uncertainty expressions: Literature review. Technical Report CMU-EPP-1990-03-02, Department of Engineering and Public Policy, Carnegie Mellon University, May 1989.
- [Dru93] M. Druzdzel. *Probabilistic reasoning in decision support systems: From computation to common sense*. PhD thesis, Department of Engineering and Public Policy, Carnegie Mellon University, 1993.
- [Dru96] M. Druzdzel. Qualitative verbal explanations in Bayesian belief networks. *Artificial Intelligence and Simulation of Behaviour Quarterly*, 94:43–54, 1996. Special Issue on Bayesian Belief Networks.
- [Dru99] M. Druzdzel. GeNIe: A development environment for graphical decision-analytic models. In *Proceedings of the 1999 Annual Symposium of the American Medical Informatics Association (AMIA-1999)*, page 1206. Washington, D.C., November 1999.

- [DS94] M. Druzdzel and H.J. Suermondt. Relevance in probabilistic models: Backyards in a small world. In *Working Notes of the AAAI-94 Workshop on Fall Symposium Series: Relevance*, pages 60–63, New Orleans, LA, November 1994.
- [EH89] C. Elsaesser and M. Henrion. Explanation of probabilistic inference. In L.N. Kanal, T.S. Levitt, and J.F. Lemmer, editors, *Uncertainty in Artificial Intelligence*, volume 3, pages 319–328. Elsevier Science Publishers, Amsterdam, 1989.
- [Els90] C. Elsaesser. Verbal expressions for probability updates. How much more probable is much more probable. In L.N. Kanal, T.S. Levitt, and J.F. Lemmer, editors, *Uncertainty in Artificial Intelligence*, volume 5, pages 387–400. Elsevier Science Publishers, Amsterdam, 1990.
- [Fie01] Armin Fiedler. Dialog-driven adaptation of explanations of proofs. In *Proceedings of the Seventeenth International Joint Conference of Artificial Intelligence*, pages 1295–1300, Seattle, Washington, August 2001.
- [FMST96] J. Folckers, C. Mbus, O. Schrder, and H. Thole. An intelligent problem solving environment for designing explanation models and for diagnostic reasoning in probabilistic domains. In *Intelligent Tutoring Systems*, pages 353–362, 1996.
- [G98] J.A. Gámez. *Inferencia Abductiva en Redes Causales*. PhD thesis, Departamento de Ciencias de la Computación e Inteligencia Artificial, Universidad de Granada, 1998.
- [Goo77] I.J. Good. Explicativity: A mathematical theory of explanation with statistical applications. In *Proceedings of the Royal Statistical Society*, 354, pages 303–330, London, 1977.

- [HD90] M. Henrion and M. Druzdzel. Qualitative propagation and scenario-based approaches to explanation of probabilistic reasoning. In *Proceedings of the 6th conference on Uncertainty in Artificial Intelligence*, pages 17–32, 1990.
- [Hec91] D. E. Heckerman. *Probabilistic Similarity Networks*. MIT Press, Cambridge, MA, 1991.
- [Hen89] M. Henrion. Some practical issues in constructing belief networks. In *Uncertainty in Artificial Intelligence*, volume 3, pages 161–173. Elsevier, Amsterdam, 1989.
- [HHNF86] E. Horvitz, D. Heckerman, B. Nathwani, and L. Fagan. The use of a heuristic problem-solving hierarchy to facilitate the explanation of hypothesis-directed reasoning. In *MEDINFO 86*, pages 27–31. Elsevier Science Publishers, 1986.
- [HJK94a] P. Haddawy, J. Jacobson, and C. Kahn. An educational tool for high-level interaction with Bayesian networks. In *Proceedings of the 6th IEEE International Conference on Tools with Artificial Intelligence*, 1994.
- [HJK94b] P. Haddawy, J. Jacobson, and C. Kahn. Generating explanations and tutorial problems from Bayesian networks. In *Proceedings of the 18th Annual Symposium on Computer Applications in Medical Care*, pages 770–774, 1994.
- [HJK97] P. Haddawy, J. Jacobson, and C. Kahn. BANTER: A Bayesian network tutoring shell. *Artificial Intelligence in Medicine*, 10:177–200, 1997.
- [HM84] R. A. Howard and J. E. Matheson. Influence diagrams. In R. A. Howard and J. E. Matheson, editors, *Readings on the*

- Principles and Applications of Decision Analysis*, pages 719–762. Strategic Decisions Group, Menlo Park, CA, 1984.
- [Jen96] F.V. Jensen. *An Introduction to Bayesian Networks*. UCL Press, 1996.
- [KST82] D. Kahneman, P. Slovic, and A. Tversky, editors. *Judgement under uncertainty: Heuristics and biases*. Cambridge University Press, Cambridge, UK, 1982.
- [LAD00] C. Lacave, R. Atienza, and F. J. Díez. Graphical explanations in Bayesian networks. In *Proceedings of the First Symposium on Medical Data Analysis, ISMDA 2000*, pages 122–129, Frankfurt, Germany, September 2000. Springer.
- [LD97] Y. Lin and M. Druzdzel. Computational advantages of relevance reasoning in Bayesian belief networks. In *Proceedings of the 13th Conference on Uncertainty in Artificial Intelligence*, pages 342–350. Morgan Kaufmann, 1997. uai97.
- [LOD01] C. Lacave, A. Oniško, and F. J. Díez. Debugging medical Bayesian networks with Elvira’s explanation capability. In *Working notes of the workshop on Bayesian Models in Medicine, AIME 2001*, pages 47–52, Cascais, Portugal, July 2001.
- [LSF88] C.P. Langlotz, E.H. Shortliffe, and L.M. Fagan. A methodology for generating computer-based explanations of decision-theoretic advice. *Medical Decision Making*, 8(4):290–303, 1988.
- [Lud98] K. Ludwig. Functionalism, causation and causal relevance. *Psyche. An interdisciplinary journal of research and consciousness*, 4(3), 1998.

<http://psyche.cs.monash.edu.au/v4/psyche-4-03-ludwig.html>.

- [MHA97a] S. McRoy, S. Haller, and S. Ali. B2: An interactive tool for explaining Bayesian reasoning in natural language. In *Conference Proceedings of Energy-Week 97*, volume V, pages 153–159, Houston, TX, January 1997. Special Session on Computers in Engineering.
- [MHA97b] S. McRoy, S. Haller, and S. Ali. Uniform knowledge representation for language processing in the B2 system. *Journal of Natural Language Engineering*, 3(2), 1997.
- [MKZ98] R. McConachy, K. Korb, and I. Zukerman. Deciding what not to say: An attentional-probabilistic approach to argument presentation. In *Proceedings of the Twentieth Annual Meeting of the Cognitive Science Society*, pages 669–674, Madison, Wisconsin, 1998.
- [MLPHH96] S. McRoy, A. Liu-Perez, J. Helwig, and S. Haller. B2: A tutoring shell for Bayesian networks that supports natural language interaction. In *Working notes of the AAAI 96 Spring Symposium on Artificial Intelligence*, Stanford, CA, March 1996.
- [MMA97] D. Madigan, K. Mosurski, and R. Almond. Graphical explanations in belief networks. *Journal of Computational and Graphic Statistics*, 6(2):160–181, 1997. Special issue on Bayesian belief networks.
- [Moo94] J. Moore. *Participating in Explanatory Dialogues: Interpreting and Responding to Questions in Context*. MIT Press, Cambridge, MA, 1994.

- [Nea90] R.E. Neapolitan. *Probabilistic Reasoning in Expert Systems: Theory and Algorithms*. Wiley-Interscience, New York, 1990.
- [Pea88] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, San Mateo, CA, 1988.
- [Pea99] J. Pearl. Reasoning with cause and effect. Technical report, Department of Computer Science. University of California, Los Angeles, California, 1999. <http://bayes.cs.ucla.edu/jphome.html>.
- [PH88] N. Pennington and R. Hastie. Explanation-based decision making: Effects of memory structure on judgment. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 14(3):521–533, 1988.
- [RP85] J.A. Reggia and B.T. Perricone. Answer justification in medical decision support systems based on Bayesian classification. *Computers in Biology and Medicine*, 15(4):161–167, 1985.
- [RvdG99] S. Renooij and L. van der Gaag. Enhancing QPNs for trade-off resolution. In *Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence*, pages 559–566, Stocholm, Sweden, 1999. Morgan Kaufmann.
- [RvdGPG00] S. Renooij, L. van der Gaag, S. Parsons, and S. Green. Pivotal pruning of trade-offs in qpns. In *Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence*, pages 515–522, Stanford, California, 2000. Morgan Kaufmann.
- [San91] E. Santos. On the generation of alternative explanations with implications for belief revision. In *Proceedings of the 7th Con-*

- ference on Uncertainty in Artificial Intelligence*, pages 339–347. Morgan Kaufmann Publishers, San Mateo, CA, 1991.
- [SB90] S. Srinivas and J. Breese. IDEAL: A software package for analysis of influence diagrams. In *Proceedings of the 6th Workshop on Uncertainty in Artificial Intelligence*, pages 212–219, Cambridge, MA, 1990.
- [SC93] H.J. Suermondt and G. Cooper. An evaluation of explanations of probabilistic inference. *Computers and Biomedical Research*, 26:242–254, 1993.
- [SCFT96] O. Shcrder, C.Mobus, J. Folckers, and H. Thole. Supporting the construction of explanation models and diagnostic reasoning in probabilistic domains. Technical report, OFFIS Institute, Oldenburg, Germany, 1996.
- [SGP97] D. Subramanian, R. Greiner, and J. Pearl. The relevance of relevance. *Artificial Intelligence*, 97(1-2):1–5, 1997.
- [Sha86a] R.D. Shachter. DAVID: Influence diagram processing system for the macintosh. In *Proceedings of the Workshop on Uncertainty in Artificial Intelligence*, pages 243–248, Philadelphia, PA, 1986. Morgan Kaufmann.
- [Sha86b] R.D. Shachter. Evaluating influence diagrams. *Operations Research*, 34(6):871–882, 1986.
- [Shi91] S.E. Shimony. *A Probabilistic Framework for Explanation*. PhD thesis, Department of Computer Science, Brown University, 1991. Technical Report CS-91-57.
- [Str87] T.M. Strat. The generation of explanations within evidential reasoning systems. In *Proceedings of the 10th International*

- Joint Conference on Artificial Intelligence*, pages 1097–1104, Milan, Italy, 1987.
- [Sue92] H.J. Suermondt. *Explanation in Bayesian Belief Networks*. PhD thesis, Department of Computer Science. Stanford University, Stanford, CA, 1992. STAN-CS-92-1417.
- [Swa83] W. Swartout. XPLAIN: A system for creating and explaining expert consulting programs. *Artificial Intelligence*, 21:285–325, 1983.
- [SZ90] P. Sember and I. Zukerman. Strategies for generating micro explanations for Bayesian belief networks. In *Uncertainty in Artificial Intelligence*, volume 5, pages 295–302. Elsevier Science Publishers, 1990.
- [TS84] R. L. Teach and E. H. Shortliffe. An analysis of physician’s attitudes. In B. G. Buchanan and E. H. Shortliffe, editors, *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*, chapter 34, pages 635–652. Addison-Wesley, Reading, MA, 1984.
- [Wel90a] M. Wellman. Fundamental concepts of qualitative probabilistic networks. *Artificial Intelligence*, 44:257–303, 1990.
- [Wel90b] M. Wellman. Graphical inference in qualitative probabilistic networks. *Networks*, 20:687–701, 1990.
- [Wic89] M.R. Wick. The 1988 Workshop on Explanation. *Artificial Intelligence Magazine*, 10(3):22–26, 1989.
- [WS89] M. Wick and J.R. Slagle. An explanation facility for today’s expert systems. *IEEE Expert*, 4:24–36, 1989.

- [WT92] M.R. Wick and W.B. Thompson. Reconstructive expert system explanation. *Artificial Intelligence*, 54:33–70, 1992.
- [ZKM98] I. Zukerman, K. Korb, and R. McConachy. Perambulations on the way to an architecture for a nice argument generator. In *Notes of the ECAI-96 Workshop on Gaps and Bridges: New Directions in Planning and Natural Language Generation*, pages 31–36, Budapest, Hungary, 1998.
- [ZMK98] I. Zukerman, R. McConachy, and R. Korb. Bayesian reasoning in an abductive mechanism for argument generation and analysis. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*, pages 833–838, Madison, Wisconsin, 1998.