On Belief Networks and Diagnosis

M.L. Wessels

UU-CS-1995-24
July 1995

Utrecht University

Department of Computer Science

Padualaan 14, P.O. Box 80.089,
3508 TB Utrecht, The Netherlands,
Tel. : ... + 31 - 30 - 531454

On Belief Networks and Diagnosis

M.L. Wessels

Technical Report UU-CS-1995-24
July 1995

Department of Computer Science
Utrecht University
P.O.Box 80.089
3508 TB Utrecht
The Netherlands

ISSN: 0924-3275

On Belief Networks and Diagnosis

Maria L. Wessels
Utrecht University, Department of Computer Science
P.O. Box 80.089, 3508 TB Utrecht, The Netherlands
e-mail: mlw@cs.ruu.nl

Abstract

In artificial intelligence, the belief network framework for reasoning with uncertainty in knowledge-based systems is becoming increasingly popular. The framework merges probability theory and expert system technology, to arrive at a powerful formalism that is both intuitively appealing and theoretically well-founded. Since its introduction, more and more knowledge-based systems are being developed using this framework, most notably in the area of diagnosis. In this paper, we survey the state-of-the-art in the field of diagnostic belief network applications. The work of Poole and Provan [14] on various concepts of diagnosis serves as a guideline. We review how these concepts are effectuated within the belief network framework in real-world applications.

1 Introduction

In the area of artificial intelligence, the interest for knowledge-based systems or expert systems has been growing rapidly during the past few decades. As more and more experience in building and using expert systems was gained, it became apparent that the early expert systems lacked facilities for dealing with uncertain or incomplete information. This observation has led to the development of a new research area in artificial intelligence called *reasoning with uncertainty* or *plausible reasoning* [1].

Probability theory provides the mathematical principles for rational inference under uncertainty. However, when probability theory is applied straightforwardly to reasoning with uncertainty, serious complexity problems are encountered: for example, explicitly representing a joint probability distribution requires exponential space in the number of variables discerned, and computing probabilities from the joint probability distribution takes exponential time. Throughout the 1960s and early 1970s, some medical systems were developed using a Bayesian model, in which simplifying assumptions were made to overcome the computational problems [2; 3; 4]. However, the expressiveness of these models was shown to be too limited when applied to larger domains, since the simplifying assumptions are not easily satisfied anymore. In response to these problems, in the 1970s various modifications of probability theory were designed for incorporation in mostly rule-based expert systems, an example of which is the certainty factor model [5]. Although these modifications were popular due to their computational simplicity, they were criticized because of their incorrectness from a mathematical point of view.

Halfway through the 1980s, the theory of *Bayesian belief networks* was introduced for reasoning with uncertainty in knowledge-based systems. The belief network framework merges probability theory and expert system technology, to obtain a powerful formalism that is both intuitively appealing and theoretically well-founded. The formalism provides for a concise representation of knowledge concerning a joint probability distribution on a set of variables discerned in a domain. In addition, it provides a set of algorithms for efficient reasoning with knowledge represented in the formalism. Since its introduction, substantial progress has been made in belief network research, and more and more knowledge-based systems are being developed using this framework, most notably in the area of diagnosis. By now, the belief network framework has made its entrance into a diversity of domains, for example, in therapy monitoring [6; 7; 8], robot monitoring [9], computer vision [10; 11], forecasting [12], and information retrieval [13].

From way back, diagnostic problem solving has been a focus of attention in artificial intelligence research. Within diagnostic problem solving, various different concepts of diagnosis are in use. The concept of diagnosis employed determines to a large extent a system's behaviour. Poole and Provan [14] have analyzed six concepts of diagnosis that are frequently used in diagnostic problem solving. The earliest applications of the belief network framework were also built for diagnostic problem solving, and still a major part of its application concerns diagnosis. This has motivated us to investigate the state-of-the-art in the field of diagnostic belief network applications. In our overview, we have built on the work of Poole and Provan. We discuss which of the concepts of diagnosis have been employed in diagnostic belief network applications and why others have not been used.

The paper is organized as follows. In Section 2 we provide an introduction into the belief network framework. Section 3 discusses the various concepts of diagnosis discerned by Poole and Provan, and examines them in the context of a belief network. Section 4 reviews some well-known practical applications of the belief network framework and classifies them according to the concept of diagnosis they use. The paper ends with some concluding remarks in Section 5.

2 The Belief Network Framework

In this section we review the belief network framework. We first explain how uncertain information is modelled in a belief network. Then we briefly review some of the algorithms that are in use for probabilistic inference with a belief network. To conclude, we describe how the belief network framework can be extended to arrive at an adequate problem solver.

2.1 Knowledge Representation

The belief network framework offers a knowledge representation formalism that allows for modelling uncertain or incomplete information in the form of a joint probability distribution on a set of variables discerned in the domain at hand. The formalism provides for a concise representation of a joint probability distribution by separating the knowledge concerning independencies between variables of the domain at hand and the numerical probabilities involved explicitly. The independencies between the variables discerned are represented by an acyclic digraph. Each node in the digraph represents a variable that can adopt one of a set of values. The set of arcs of the digraph models an independency relation between the variables: absence of an arc between two variables means that these variables are (conditionally) independent.

Associated with each node of the digraph of a belief network is a table of (conditional) probabilities describing the influence of the values of the predecessors of the node on the probabilities of the values of the node itself. These tables provide all information necessary for uniquely defining a joint probability distribution that respects the independency relation portrayed by the digraph of the network [15].

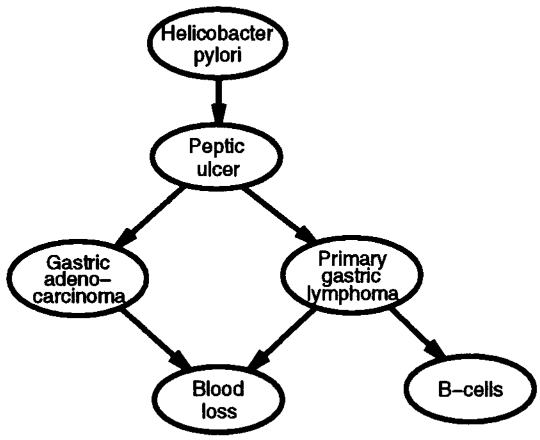

Figure 1: A diagnostic belief network

Figure 1 illustrates the belief network formalism for some medical knowledge concerning malignant gastric diseases. Although the domain is of a serious nature, the example should not be taken too seriously. Two malignant gastrical diseases are distinguished: primary gastric lymphoma and gastric adenocarcinoma. Both diseases can be caused by peptic ulcer. Recent research has shown that a bacterial infection by the helicobacter pylori may be one of the causes of peptic ulcer. The presence of B-cells is often seen in combination with gastric lymphoma. Gastric lymphoma and gastric adenocarcinoma can both cause blood loss. The conditional probability tables associated with the nodes have been omitted in this example. It should be noted that the information modelled is incomplete. For example, peptic ulcer may have other causes not modelled by the example belief network.

2.2 Probabilistic Inference

The digraph of a belief network and its associated conditional probability tables uniquely define a joint probability distribution on the set of variables discerned. From the probability distribution defined by the belief network, any probability of interest can be computed. Also, when the value of a specific variable becomes known, the revised probability of each of the values of the other variables can be computed from the belief network. However, computing probabilities by simply applying the basic rules of marginalization and conditioning would require exponential time complexity. Many algorithms have been developed to reduce the computational complexity of probabilistic inference. Although these algorithms differ in implementational details, they all gain their efficiency by exploiting the independencies portrayed by the digraph of the belief network. The algorithms view the digraph of the belief network more or less directly as an "object-oriented" computational architecture. Each node is provided with a local memory in which associated probabilistic information is stored and a local processor that is able to perform simple probabilistic computations; the arcs serve as bidirectional communication channels. The idea now is that each node of the network is able to compute its probabilities by performing local computations only, employing the information stored in its local memory and the information it receives from its neighbouring nodes. The algorithms tend to behave polynomially under certain topological restrictions. However, the restrictions under which they achieve their polynomial behaviour are different. It should be noted that exact probabilistic inference with belief networks without any topological restrictions is NP-hard [16]. We now discuss two algorithms that are frequently used in applications of the belief network framework.

In Pearl's algorithm [15], each node in a belief network receives messages on marginal probabilities from its predecessors and on likelihoods from its successors. The information contained in the messages and the information from the probability table stored in the local memory enables a node to compute the probabilities for its values. Initially, the belief network is in equilibrium. When the value of a certain node becomes known, the messages that the node sends to its neighbours are updated to reflect the obtained information. The neighbouring nodes in turn update their probabilities and compute new messages to send to their neighbours. In this message passing process each node is visited only once. The belief network therefore reaches a new equilibrium after all nodes has been visited. For multiply connected belief networks, the algorithm described above cannot be applied straightforwardly, since the computational architecture now contains loops that may cause the messages to cycle infinitely. The method of loop cutset conditioning offers a solution to this problem [17]. It is based on a form of human reasoning called *reasoning by assumption.* The idea is that the loops in the digraph are cut by a subset of nodes, called the *loop cutset*, by instantiating these nodes. For each possible combination of values of the loop cutset nodes updated probabilities for a variable of interest are computed. The variable's probabilities then are computed by weighing these updated probabilities.

The algorithm of Lauritzen and Spiegelhalter [18] does not use the belief network directly as a computational architecture. A special computational architecture is constructed by translating the original digraph of the belief network into a decomposable graph, consisting of cliques and clique intersections. The original conditional probabilities are transformed into a set of marginal probability distributions associated with the cliques of the graph. The so-obtained decomposable belief network still represents the same joint probability distribution as the original belief network. The cliques of the decomposable belief network serve as the nodes of the computational architecture, with the marginal probability distributions stored in its local memories; the clique intersections serve as the communication channels. The message passing is performed much in the same way as in Pearl's algorithm. To prevent messages from cycling, the architecture is organized as a tree. The algorithm is therefore often called the *clique tree propagation* (CTP) algorithm.

Since it is not expected that a probabilistic inference algorithm can be designed that provides an exact answer for all belief network topologies in polynomial time, it is not surprising that approximate algorithms are an area of active research interest. Approximation algorithms for belief networks use simulation as a method for estimating probabilities of network nodes. The idea of such algorithms is to generate a large number of samples specifying a value for each node of the belief network; these samples are generated so as to reflect the probability distribution defined by the belief network. The algorithm subsequently performs a frequency count on the sampled values. The normalized results provide estimates for the probabilities of the values of the nodes in the belief network. Currently, several different simulation algorithms are in use for probabilistic inference with belief networks. They differ primarily in the method they use for selecting samples and for counting frequencies. The simulation algorithms generally have an 'anytime' property in the sense that they give an approximate answer at any time; given more time, their estimates improve. For details on simulation algorithms, we refer to Cousins et al. [19].

Simulation algorithms seem to accommodate a wider range of belief network topologies than exact algorithms. However, they do not give accurate answers for all belief networks in polynomial time either; this task has been proven to be NP-hard as well [20].

2.3 The Belief Network as Problem Solver

Problem solvers, whether human or computerized, typically apply reasoning strategies to solve the problems they are used for in the best way. For this purpose, a problem solver is provided with means for exerting control over reasoning. We distinguish between two types of control. The first type concerns the interaction between user and problem solver. This type of control strongly depends on the domain of application. It involves for example strategies for acquiring information about the problem at hand, and for suggesting actions that can be undertaken to solve the problem. The second type of control is aimed at reducing the computational cost of the reasoning activities themselves. This type of control may be domain-dependent as well as domain-independent. It involves for example focusing reasoning activities on a restricted part of the domain.

The belief network framework comprises algorithms for computing probabilities of interest efficiently, and for processing pieces of evidence. However, the framework lacks with regard to intelligent control over probabilistic reasoning. To enable the belief network framework to perform as an adequate diagnostic problem solver, several extensions to the framework have been proposed in literature. First, the belief network formalism has been enhanced to the formalism of influence diagrams [21]. An influence diagram allows for making decisions based on methods from decision theory. The influence diagram in itself, however, does not provide for controlling the complexity of probabilistic inference. Secondly, the belief network framework can be extended with a control layer [22]. The control layer offers algorithms for decision making and for complexity control from which a strategy for problem solving is built. Note that the control layer allows for the incorporation of a wide variety of controlling mechanisms, dependent or independent of the domain of application.

3 The Most Likely Diagnosis

For shaping diagnostic reasoning various definitions of the concept of diagnosis have been proposed. Poole and Provan have analysed several well-known definitions in their paper "What is the most likely diagnosis?" [14]. They have shown that different definitions of diagnosis have different qualitative meaning and as a consequence yield different system behaviour. Since we feel that Poole and Provan can be considered authorative in the field of diagnosis, we have chosen their analysis as a point of departure for investigating how different concepts of diagnosis are modelled in belief network applications. In this section, we provide a concise overview of the definitions of the *most likely diagnosis* given by Poole and Provan and discuss which of these definitions are used in the context of the belief network framework.

3.1 Definitions of Diagnosis

The various definitions of diagnosis distinguished by Poole and Provan are expressed in terms of hypotheses. A hypothesis may be looked upon as a proposition. The first definition departs from a *single-fault* assumption. Under this assumption a diagnosis is one from a set of mutually exclusive hypotheses. The most likely diagnosis is defined as the single hypothesis that has highest probability given all available evidence. It is referred to as the *most likely single-fault hypothesis*. In a medical context, where it seems appropriate to assume that every hypothesis models a single disease, this definition of diagnosis may be used only if a patient cannot suffer from more than *one* disease at a time.

The single-fault assumption is relaxed in the second definition of diagnosis. The most likely diagnosis again is defined as the hypothesis with highest probability given the evidence, but now the hypotheses considered need no longer be mutual exclusive. This diagnosis is referred to as the *most likely posterior hypothesis*. In a medical context, where we assume once more that every hypothesis models a single disease, the most likely diagnosis is the disease with highest posterior probability. Since the hypotheses considered need no longer be mutually exclusive, a patient may suffer from *more* than one disease at a time, but this is not expressed in the diagnosis that is yielded.

For the next definition of diagnosis, truth assignments are given to all hypotheses discerned. The conjunction of truth assignments having highest posterior probability is taken as the most likely diagnosis, and is called the *most likely interpretation*. While the most likely diagnoses of the previous definitions are single hypotheses, the most likely interpretation is a conjunction of hypotheses. Thus, in a medical context where a hypothesis models a single disease, the most likely interpretation may involve multiple diseases.

The following definitions build on a logical axiomatisation in addition to a probabilistic model of the problem domain. In these definitions, a diagnosis is a conjunction of truth assignments to a *subset* of hypotheses discerned in the domain. In the definition referred to as *probability of provability*, the logical model is used to *prove* conjunctions of hypotheses from the observed evidence. The probabilistic model is used to compute probabilities for the deduced conjunctions. The most likely diagnosis is then defined as the conjunction that has highest probability. In the other definition building on a logical axiomatization, the logical model is used to compose conjunctions of hypotheses that *explain* the observed evidence. As in the probability of provability definition, the probability model is used to compute the probability of each of the abducted conjunctions. The most likely diagnosis, called *covering explanation*, is defined as the conjunction that has highest posterior probability. The diagnoses computed by the provability and covering approaches have been shown to be identical under certain conditions [23].

The last definition of the most likely diagnosis discussed by Poole and Provan is referred to as a *utility-based explanation*. While the previous definitions consider a diagnosis with respect to collected evidence only, this definition also takes into account the purpose for which a diagnosis is required. The utility-based explanation is the hypothesis that is most useful for reaching this purpose. In a medical context, making a distinction between two diseases that require the same medical treatment is considered useless if diagnostic reasoning is aimed at restoring a patient to health. Poole and Provan argue that a definition of diagnosis should depend on what the diagnosis is used for, and therefore they recommend the utility-based explanation as the most appropriate definition of diagnosis. The utility-based explanation should not be viewed as a separate definition of diagnosis: each of the previous definitions can be supplemented with a utility-based component, since extra knowledge is added to reach this kind of diagnosis.

3.2 Belief Networks and Diagnosis

We now consider to what extent the various definitions of diagnosis are reflected in applications of the belief network framework. Before doing so, we have to introduce the different roles the nodes of a belief network play in diagnostic reasoning. A *hypothesis node* represents a variable whose values are hypotheses; an *evidence node* represents a variable whose value can be obtained by observation; and an *intermediate node* represents a variable not classified in either of the former two groups.

In a belief network, the definition of the most likely single-fault hypothesis is supported by a restriction of the number of hypothesis nodes. This number is restricted to one, forcing the hypotheses, corresponding to the values of the hypothesis node, to be mutually exclusive. The belief network computes the most likely single-fault hypothesis by processing all available evidence, and then computing the posterior probabilities of each of the values of the hypothesis node. The value responsible for the highest probability is taken as the most likely diagnosis.

For computing the most likely posterior hypothesis, the hypotheses considered in the domain need not be mutually exclusive. To support this definition of diagnosis, more than one hypothesis node is admitted in the belief network. The belief network computes the most likely posterior hypothesis by processing all available evidence, and computing the posterior probabilities of each of the values of each of the hypothesis nodes. The hypothesis node responsible for the highest probability yields the most likely diagnosis.

For computing a most likely interpretation a value is assigned to all hypothesis nodes of the belief network, in such way that the resulting conjunction of values has highest probability. The number of value assignments is exponential in the number of nodes of the network; however, it is not necessary to enumerate all assignments in order to find the most likely interpretation [15]. This definition of diagnosis does not let us be ignorant about the value of any of the hypothesis nodes of the belief network. In diagnostic belief network applications, however, value assignments to hypothesis nodes that are hardly related to the observed evidence may only be disturbing for the outcome of the most likely diagnosis [14]. This may explain why the most likely interpretation is almost never used in diagnostic belief network applications. For belief network applications outside the field of diagnosis this definition has turned out to be useful [25], and is known as the *most probable explanation* (MPE) [15] or *maximum a posteriori probability* (MAP) [24].

For computing probability of provability and covering explanations, a logical axiomatisation of a domain is required. Since the belief network framework has its foundation in probability theory, it does not provide in itself for a logical axiomatization. However, the independencies portrayed by the digraph of the belief network can be exploited for constructing a diagnosis that is a conjunction of values of a *subset* of hypothesis nodes [26] and hence the resulting diagnosis to some extent looks like the diagnosis yielded by the logical approaches. If the belief network is supplemented with a control layer, this layer could easily incorporate a logical axiomatisation of the problem at hand. The belief network could serve as the probabilistic model needed to compute probabilities of deduced or abducted conjunctions of hypotheses.

4 Diagnostic Belief Networks

From the various definitions of diagnosis discerned by Poole and Provan only two are commonly used in diagnostic belief networks: the most likely single-fault hypothesis and the most likely posterior hypothesis. Both definitions sometimes are supplemented with a utility-based component. In this section we illustrate how these definitions are effectuated in some well-known diagnostic belief networks. In each section we discuss one application at length, and briefly review other applications.

4.1 Most likely single-fault hypothesis

As we have argued in the previous section, the most likely single-fault hypothesis is supported in a diagnostic belief network by restricting the number of hypothesis nodes to one. Although the single-fault assumption may at first sight seem rather restrictive, there are several diagnostic belief network applications that successfully adopt this assumption.

PATHFINDER is a knowledge-based system based on the belief network framework that assists surgical pathologists with the diagnosis of lymph-node diseases. Surgical pathologists examine sections of lymph-node tissue microscopically for reaching a diagnosis. It is important to discriminate between malignant and benign diseases as a malignant disease requires immediate treatment. Therefore, a timely as well as accurate diagnosis is required. The specialist field of lymph-node diseases, however, is one of the most difficult areas in surgical pathology. General pathologists often need to refer to subspecialists, incurring delay and extra costs. The PATHFINDER system has been built to close the gap between the quality of diagnoses made by general pathologists and those made by subspecialists. The PATHFINDER project began in 1984 as a rule-based expert system. The most recent version of PATHFINDER is a knowledge-based system employing a belief network [27; 28]. The system has been made commercially available as INTELLIPATH for practicing pathologists and pathologists in training.

The PATHFINDER system incorporates the largest belief network built on the single-fault assumption. It comprises over 140 nodes. The hypothesis node of PATHFINDER represents 63 mutually exclusive disorders. The builders of PATHFINDER argue that in the domain of lymph-node pathology the assumption that diseases are mutually exclusive is appropriate, because co-occurring diseases almost always appear in different lymph nodes or in different regions of a lymph node [27]. Besides the hypothesis node, the belief network contains only evidence nodes. Their number of values range from 2 to 10. The hypothesis node itself has no predecessors, yet is predecessor of all but two evidence nodes. This topological feature has inspired a more efficient algorithm for exact probabilistic inference called *aggregation after decomposition* that combines loop cutset conditioning and CTP [17].

In a dialogue with the user, PATHFINDER passes through a diagnostic cycle. To this end, the belief network of PATHFINDER is embedded in a problem solving architecture. The user enters observed evidence as values for the appropriate evidence nodes. After processing theses values in the belief network, PATHFINDER displays a list of node-value pairs, and the differential diagnosis being a list of diseases ranked according to their posterior probabilities. On request, PATHFINDER recommends additional tests to perform and explains why a test is recommended. The builders of PATHFINDER have experimented with utility-based explanations for the most likely diagnosis. However, to be satisfactory patient-dependent utility assessments would be required. In an application where delay in diagnosis may be fatal, patient-dependent utility assessments are thought to be infeasible.

The other belief network applications in which a single-fault assumption is adopted are considerably smaller than PATHFINDER. QUALICON (*QUALIty CONtrol*) is a knowledge-based system for quality control in nerve conduction studies [29]. Nerve conduction studies are performed to investigate action potentials in muscles and nerves. Abnormal potentials can point to a nerve disease. Yet, they sometimes are caused by misplacement of electrodes. QUALICON checks the acceptability of a potential. In case of an unacceptable potential, it investigates whether the electrodes are placed incorrectly and provides the user with recommendations for replacement. The intended users of QUALICON are residents, fellows, and technologists in a hospital EMG laboratory. Given the intended user level, the builders of QUALICON argue that it is unlikely that two or more errors will occur simultaneously, making the single-fault assumption appropriate. The QUALICON belief network is small (less than 10 nodes) and singly connected. Probabilistic inference is performed with Pearl's algorithm.

DAACS (*Dump Analysis And Consulting System*) is an example diagnostic system for a non-medical domain [30]. The system analyzes errors that have caused an assembler language program to terminate execution abnormally. In DAACS, the diagnosis is computed in two phases: the prediagnosis phase and the actual diagnosis. In the prediagnosis phase data from a minidump are extracted and the symptoms which caused the dump to occur are identified. These symptoms are entered as evidence into the belief network. In the diagnosis phase, the belief network is used to compute the underlying cause of the error by collecting and processing additional evidence. The belief network of DAACS has been kept small intentionally by using separate belief networks for major classes of mutually exclusive symptoms. The separate networks are singly-connected and use Pearl's algorithm for probabilistic inference.

4.2 Most likely posterior hypothesis

A belief network for which the most likely posterior hypothesis definition of diagnosis is adopted typically comprises more than one hypothesis node in its digraph. Many belief network applications build on this definition.

MUNIN (*MUscle and Nerve Inference Network*) is an intelligent assistant for the diagnosis of muscle and nerve diseases or, more in specific, for the interpretation of electromyographic findings to this end [31]. The interpretation of an electromyography (EMG) requires considerable expertise that takes years to acquire. The knowledge-based system MUNIN has been developed to aid in EMG-interpretation. Although the MUNIN and QUALICON systems both are applications in the EMG field, they differ with respect to their focus of attention: MUNIN addresses diseases, while QUALICON addresses electrode misplacement. The MUNIN system is the earliest application of the belief network framework: its first prototype dates from as early as 1987 [32; 33; 34]. The MUNIN system presently consists of a diagnoser and a test planner. The diagnoser queries the posterior probabilities of the diseases computed by the belief network. From these probabilities the most likely posterior hypothesis is computed. However, in the future the diagnoser is also capable to yield a utility-based explanation. To this end, the system is supplemented with utilities reflecting severity and treatability of a specific disease, and the possible side effects of treatment. Since EMG examination takes time and is uncomfortable for a patient, it is desirable to restrict the examination to a limited number of tests. To this end, the test planner computes the expected benefit for each test and decides on the outcomes which test is to be performed.

The current version of the MUNIN network comprises about 1100 nodes and 2500 arcs; the number of values represented by one node can be up to 27. About 20 disorders are modelled in the network, including general muscle and nerve disorders, and local nerve disorders. Each disorder is described by a number of nodes, typically three. The number of nodes modelling evidence is about 150. The size of the belief network and the complexity of its probabilistic inference have inspired a modelling technique called *divorcing multiple parents*. The aim of divorcing multiple parents is to reduce the number of predecessors of a node by introducing intermediate nodes. Hence, the conditional probability tables reduce in size resulting in faster inference. Probabilistic inference in MUNIN is performed with the CTP algorithm.

Rather than attempting to cover the full range of neuromuscular disorders as does MUNIN, the PAINULIM (*PAINful or impaired Upper LIMb*) expert system focuses on diseases of the spinal cord or the peripheral nervous system giving rise to painful or impaired upper limbs [35; 36]. The PAINULIM belief network comprises 14 hypothesis nodes and 69 evidence nodes; it has 271 arcs. The belief network is decomposed into natural subnetworks according to the three major information sources for diagnostic recommendations: clinical examination, EMG, and nerve conduction. The network is said to be *multiply sectioned*. In a dialogue with the system, the user focuses on a subnetwork of interest. The computation of probabilities and the propagation of evidence is restricted to the subnetwork focused on. Through this restriction, the computational costs are reduced approximately by half. PAINULIM uses the CTP algorithm for probabilistic inference.

The QMR-BN (*Quick Medical Reference-Belief Network*) system is one of the largest diagnostic applications of the belief network framework [37; 38; 39]. QMR-BN is a probabilistic interpretation of the INTERNIST-1 system, that was developed in the early 1980s as a decision-support tool for general internal medicine.

As the domain of internal medicine is very complex, several simplifying assumptions have been made to model the complex knowledge of this domain and to burden the computational complexity of inference. These simplifications have led to a so-called *two-layered* belief network. In a two-layered belief network the nodes are either hypothesis nodes or evidence nodes, and the arcs are directed from hypothesis to evidence nodes only. This topological property of the network reflects the assumptions of marginal independence of the diseases and conditional independence of findings. The influence of multiple disorders on a single finding further is taken to equal the "sum" of their individual influences. This assumption considerably reduces the number of conditional probabilities needed. The QMR belief network has more than 4500 nodes and 40,700 arcs. The number of hypothesis nodes is 534. For a belief network of this size, exact probabilistic inference is not feasible; QMR-BN therefore uses a simulation algorithm called *likelihood weighting*. The QMR-BN system has been extended to an influence diagram, called QMR-DT (*Decision Theoretic*), in which test and treatment decisions based on utility models are provided for.

The ALARM (*A Logical Alarm Reduction Mechanism*) monitoring system is a prototype knowledge-based system for diagnosing abnormalities measured by an anesthesia monitor. It models a small subset of variables that anesthesiologists encounter in the operating room [40; 41]. ALARM refers to *A Logical Alarm Reduction Mechanism*. The ALARM system accepts measurements as values for the evidence nodes and provides warnings when measurements are outside their normal range. Furthermore, it displays the differential diagnosis being a list of possible causes of the abnormalities measured, ranked according to their posterior probabilities. The ALARM belief network has 37 nodes and 46 arcs. There are 8 hypothesis nodes, 16 evidence nodes and 13 intermediate nodes. A node has at most 5 values. The CTP algorithm is used for probabilistic inference. In the domain of anesthesiology it is important that a diagnosis is reached quickly when abnormalities are measured. To this end, a small subset of belief network instantiations, based on the likelihood that the user will enter certain evidence values or based on situations that require a rapid diagnosis, are stored in a cache associated with the belief network. These cases need not be recomputed by the belief network when these evidence values are actually obtained [41].

In general, the sizes of the belief networks supporting the most likely posterior hypothesis definition of diagnosis are larger than those adopting the single-fault assumption. The digraphs of these belief networks are all multiply connected. They generally use the CTP algorithm or an approximation algorithm for probabilistic inference, where the networks adopting the single-fault assumption use Pearl's algorithm. The complexity of the domains of application could be an explanation for these observations. The hypothesis nodes are mainly located at the top of the digraphs, that is, they have no predecessors. The evidence nodes tend to be located at the lower part of the digraph, that is, they have no or only few descendants. This is explained by the fact that the direction of the arcs in the digraph of a belief network to some extent expresses causality.

5 Conclusions

In this paper, we have reviewed the state-of-the-art in diagnostic belief network applications by describing how different concepts of diagnosis are effectuated within the belief network framework. To this end, we have built on the work by Poole and Provan [14] who have distinguished between six different concepts of diagnosis. In diagnostic applications of the belief network framework, the most likely single-fault hypothesis and the most likely posterior hypothesis are frequently encountered. The most likely single-fault hypothesis is based on the assumption that the hypotheses considered in the domain are mutually exclusive. In a belief network, this assumption is reflected in a restriction of the number of hypothesis nodes to one. The most likely posterior hypothesis allows for more than one hypothesis node. We have reviewed typical examples of diagnostic belief network applications built on these concepts of diagnosis. The belief network applications in which the single-fault assumption is adopted tend to be smaller than the applications using the concept of the most likely posterior hypothesis. The limited scope of the problem domains that allow for the single-fault assumption is an explanation for this observation.

The most likely interpretation, where values are assigned to all hypothesis nodes of the network, is almost never used in diagnostic belief network applications, because of the disturbing effect that value assignments to hypothesis nodes that are hardly related to the observed evidence could have on the diagnosis yielded. This concept of diagnosis is used more often in belief network applications outside the field of diagnosis.

The probability of provability diagnoses and covering explanations have not been found in belief network applications until now. These diagnoses, based on a logical axiomatization of the problem domain, generally are multiple-disorder diagnoses. The belief network framework in essence provides for multiple-disorder diagnosis. However, the two concepts of diagnosis do not provide a method to construct such diagnosis in the context of a belief network. The method of multiple-disorder diagnosis proposed in [26] could be seen as a first step towards such a method. It would be interesting to investigate how the diagnosis obtained by this method differs from the most likely posterior hypothesis, the probability of provability hypothesis, and the covering explanation.

The use of the utility-based explanation seems to be increasingly popular in diagnostic belief network applications. In the medical domain, utilities are often difficult to assess, since they differ strongly from patient to patient. However, an increasing number of belief network applications are used for diagnosis in technical domains, where a utility-based explanation could be easier to incorporate.

More and more belief networks are supplemented with a control layer or extended to an influence diagram to enable them to perform as adequate problem solvers. Most of the present belief network applications provide for selective evidence gathering and sometimes even treatment selection. Also, special-purpose methods for control over computational complexity are frequently seen. These methods include focusing strategies, caching, and combined inference algorithms.

The belief network has gained in popularity as is demonstrated by the increasing number of applications in a wide variety of domains [42]. Yet, control over reasoning still needs our attention to arrive at even more competent problem-solving behaviour.

References

[1] G. Shafer and J. Pearl. *Readings in Uncertain Reasoning.* Morgan Kaufmann, San Mateo, California, 1990.

[2] H.R. Warner, A.F. Toronto, L.G. Veasy, and R. Stephenson. A mathematical approach to medical diagnosis: application to congenital heart disease. *Journal of the American Medical Association*, vol. 177, pp. 177–183, 1961.

[3] G.A. Gorry and G.O. Barnett. Experience with a model of sequential diagnosis. *Computers and Biomedical Research*, vol. 1, pp. 490–507, 1968.

[4] F.T. de Dombal, D.J. Leaper, J.R. Staniland, A.P. McCann, and J.C. Horrocks. Computer-aided diagnosis of acute abdominal pain. *British Medical Journal*, vol. 2, pp. 9–13, 1972.

[5] E.H. Shortliffe and B.G. Buchanan. A model of inexact reasoning in medicine. In *Rule-Based Expert Systems. The MYCIN Experiments of the Stanford Heuristic Programming Project*, pp. 233–262. Addison-Wesley, Reading, Massachusetts, 1984.

[6] D.J. Spiegelhalter, A.P. Dawid, T.A. Hutchinson, and R.G. Cowell. Probabilistic expert systems and graphical modelling: a case study in drug safety. *Philosophical Transactions: Physical Sciences and Engineering*, vol. 337, pp. 307–428, 1991.

[7] S. Andreassen, R. Hovorka, J. Benn, K. Olesen, and E. Carson. A model-based approach to insulin adjustment. In M. Stefanelli, A. Hasman, M. Fieschi, and J. Talmon, editors, *AIME 91, Proceedings of the Third Conference on Artificial Intelligence in Medicine*, pp. 239–248. Springer-Verlag, Berlin, 1991.

[8] G.W. Ruthledge, G.E. Thomsen, B.R. Farr, M.A Tovar, J.X. Polaschek, I.A. Beinlich, L.B. Sheiner, and L.M. Fagan. The design and implementation of a ventilator-management advisor. *Artificial Intelligence in Medicine*, vol. 5, pp. 67–82, 1993.

[9] A.E. Nicholson and J.M. Brady. Sensor validation using dynamic belief networks. In D. Dubois, M.P. Wellman, B. D'Ambrosio, and P. Smets, editors, *Uncertainty in Artificial Intelligence, Proceedings of the Eighth Conference*, pp. 207–214. Morgan Kaufmann, San Mateo, California, 1992.

[10] T.S. Levitt, M.W. Hedgcock, J.W. Dye, S.E. Johnston, V.M. Shadle, and D. Vosky. Bayesian inference for model-based segmentation of computed radiographs of the hand. *Artificial Intelligence in Medicine*, vol. 5, pp. 365–387, 1993.

[11] M. Bibbo, P.H. Bartels, T. Pfeifer, D. Thompson, C. Minimo, and H.G. Davidson. Belief network for grading prostate lesions. *Analytical and Quantitative Cytology and Histology*, vol. 15, pp. 124–135, 1993.

[12] B. Abramson. ARCO1: An application of belief networks to the oil market. In B. D'Ambrosio, P. Smets, and P.P. Bonissone, editors, *Uncertainty in Artificial Intelligence, Proceedings of the Seventh Conference*, pp. 1–8. Morgan Kaufmann, San Mateo, California, 1991.

[13] P.D. Bruza and L.C. van der Gaag. Index expression belief networks for information disclosure. *The International Journal of Expert Systems: Research and Applications*, vol. 7, pp. 107–138, 1994.

[14] D. Poole and G.M. Provan. What is the most likely diagnosis? In P.P. Bonissone, M. Henrion, L.N. Kanal, and J.F. Lemmer, editors, *Uncertainty in Artificial Intelligence 6*, pp. 89–105. North Holland, Amsterdam, 1991.

[15] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufmann, San Mateo, California, 1988.

[16] G.F. Cooper. The computational complexity of probabilistic inference using Bayesian belief networks. *Artificial Intelligence*, vol. 42, pp. 393–405, 1990.

[17] H.J. Suermondt and G.F. Cooper. Probabilistic inference in multiply connected belief networks using loop cutsets. *International Journal of Approximate Reasoning*, vol. 4, pp. 283–306, 1990.

[18] S.L. Lauritzen and D.J. Spiegelhalter. Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society, Series B*, vol. 50, pp. 157–224, 1988.

[19] S.B. Cousins, W. Chen, and M.E. Frisse. A tutorial introduction to stochastic simulation algorithms for belief networks. *Artificial Intelligence in Medicine*, vol. 5, pp. 315–340, 1993.

[20] P. Dagum and M. Luby. Approximating probabilistic inference in Bayesian belief networks is NP-hard. *Artificial Intelligence*, pp. 141–153, 1993.

[21] 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*, pp. 720–763. Strategic Decisions Group, 1984.

[22] L.C. van der Gaag and M.L. Wessels. Selective evidence gathering for diagnostic belief networks. *AISB Quarterly*, no. 86, pp. 23–34, 1993.

[23] D.L. Poole. Representing knowledge for logic-based diagnosis. In *Proceedings of the International Conference on Fifth generation Computer Systems*, pp. 1282–1290, 1988.

[24] S.E. Shimony and E. Charniak. A new algorithm for finding MAP assignments to belief networks. In P.P. Bonisone, M. Henrion, L.N. Kanal, and J.F. Lemmer, editors, *Uncertainty in Artificial Intelligence 6*, pp. 185–193. Elsevier Science Publishers, Amsterdam, 1991.

[25] L.M.R. Eizirik, V.C. Barbosa, and S.B.T. Mendes. A Bayesian-network approach to lexical disambiguation. *Cognitive Science*, vol. 17, pp. 257–283, 1993.

[26] L.C. van der Gaag and M.L. Wessels. Multiple-disorder diagnosis with belief networks. In G.M. Provan, editor, *DX-94, Fifth International Workshop on Principles of Diagnosis*, pp. 343–351, 1994.

[27] D.E. Heckerman, E.J. Horvitz, and B.N. Nathwani. Toward normative expert systems. Part 1: The PATHFINDER project. *Methods of Information in Medicine*, vol. 31, pp. 90–105, 1992.

[28] D.E. Heckerman and B.N. Nathwani. Toward normative expert systems. Part 2: Probability-based representations for efficient knowledge acquisition and inference. *Methods of Information in Medicine*, vol. 31, pp. 106–116, 1992.

[29] Y. Xiang, A. Eisen, M. MacNeil, and M.P. Beddoes. Quality control in nerve conduction studies with coupled knowledge-based system approach. *Muscle and Nerve*, vol. 15, pp. 180–187, 1992.

[30] L.J. Burnell and S.E. Talbot. Incorporating probabilistic reasoning in a reactive program debugging system. In *Proceedings of the Ninth Conference on Artificial Intelligence for Applications*, pp. 321–327. IEEE Computer Society Press, 1993.

[31] S. Andreassen, F.V. Jensen, S.K. Andersen, B. Falck, U. Kjærulff, M. Woldbye, A.R. Sørensen, A. Rosenfalck, and F. Jensen. MUNIN - an expert EMG assistant. In J.E. Desmedt, editor, *Computer-Aided Electromyography and Expert Systems*, chapter 21, pp. 255–277. Elsevier Science Publishers, Amsterdam, 1989.

[32] S. Andreassen, M. Woldbye, B. Falck, and S.K. Andersen. MUNIN - a causal probabilistic network for interpretation of electromyographic findings. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, pp. 366–372, 1987.

[33] F.V. Jensen, S.K. Andersen, U. Kjærulff, and S. Andreassen. MUNIN - on the case for probabilities in medical expert systems - a practical exercise. In J. Fox, M. Fieschi, and R. Engelbrecht, editors, *AIME 87 Proceedings of the European Conference on Artificial Intelligence in Medicine*, pp. 149–160. Springer-Verlag, Berlin, 1987.

[34] K.G. Olesen, U. Kjærulff, F. Jensen, F.V. Jensen, B. Falck, S. Andreassen, and S.K. Andersen. A MUNIN network for the median nerve - a case study on loops. *Applied Artificial Intelligence*, vol. 3, pp. 385–403, 1989.

[35] Y. Xiang, B. Pant, A. Eisen, M.P. Beddoes, and D. Poole. Multiply sectioned Bayesian networks for neuromuscular diagnosis. *Artificial Intelligence in Medicine*, vol. 5, pp. 293–314, 1993.

[36] Y. Xiang, D. Poole, and M.P. Beddoes. Multiply sectioned Bayesian belief networks and junction forests for large knowledge-based systems. *Computational Intelligence*, vol. 9, pp. 171–220, 1993.

[37] M. Henrion. Towards efficient probabilistic diagnosis in multiply connected belief networks. In R.M. Oliver and J.Q. Smith, editors, *Influence Diagrams, Belief Networks and Decision Analysis*, chapter 17, pp. 385–409. John Wiley & Sons, New York, NY, 1990.

[38] M.A. Shwe, B. Middleton, D.E. Heckerman, M. Henrion, E.J. Horvitz, H.P. Lehmann, and G.F. Cooper. Probabilistic diagnosis using a reformulation of the INTERNIST-1/QMR knowledge base. The probabilistic model and inference algorithms. *Methods of Information in Medicine*, vol. 30, pp. 241–255, 1991.

[39] B. Middleton, M.A. Shwe, D.E. Heckerman, M. Henrion, E.J. Horvitz, H.P. Lehmann, and G.F. Cooper. Probabilistic diagnosis using a reformulation of the INTERNIST-1/QMR knowledge base. Evaluation of diagnostic performance. *Methods of Information in Medicine*, vol. 30, pp. 256–267, 1991.

[40] I.A. Beinlich, H.J. Suermondt, R. Martin Chavez, and G.F. Cooper. The ALARM monitoring system: a case study with two probabilistic inference techniques for belief networks. In J. Hunter, J. Cookson, and J. Wyatt, editors, *AIME 89 Proceedings of the Second Conference on Artificial Intelligence in Medicine*, pp. 247–256. Springer-Verlag, Berlin, 1989.

[41] E.H. Herskovits and G.F. Cooper. Algorithms for Bayesian belief-network precomputation. *Methods of Information in Medicine*, vol. 30, pp. 81–89, 1991.

[42] D.E. Heckerman. Literature search. Available as pub/dtg/bn-apps.ps on research.microsoft.com, 1994.