Medical Informatics Europe '96
J. Brender et al. (Eds.)
IOS Press, 1996

A heuristic method for diagnosing multiple diseases in complex medical domains modelled by causal probabilistic networks

Kristian G. Olesen and Steen Andreassen
Department of Medical Informatics and Image Analysis,
Institute of Electronic SystemsAalborg University,
Fredrik Bajersvej 7 D, DK-9220 Aalborg, Denmark
Phone +45 9815 8522, Fax: +45 9815 4008,
Email: kgo@miba.auc.dk & sa@miba.auc.dk

Abstract: We propose a method to manage the complexity of large expert systems based on causal probabilistic networks. The method operates by partitioning large domains into smaller subdomains. An iterative procedure operates through increasingly complex subdomains, where inference is limited to consider only plausible diagnoses. The method is introduced by a simple example and we report a case from MUNIN, where a patient with three diseases is diagnosed.

1. Introduction

A problem that have caused great difficulty in medical expert systems is the diagnosis of multiple diseases. Most systems are limited by the assumption that at most one disease is present, and the diagnostic process is thereby limited to find the most probable diagnosis. The results of such systems are typically presented as an ordered list, where the possible disorders are ranked by some numerical measure. The Pathfinder system for diagnosis of lymph-node pathology [1] is one of the few medical expert systems that have found their way into practical use. Pathfinder is modelled as a causal probabilistic network (CPN) [2, 3], where the diagnoses are states of a single discrete random variable. Because the states of a random variable are mutually exclusive each possible combination of two or more diseases have to be included as separate states of the variable. This may be a satisfactory solution in certain domains, but in general it leads to a combinatorial explosion. Alternatively each disease can be modelled as a single variable, but this will typically result in a complex CPN with many cycles. In the worst case these cycles will prevent the generation of a runtime system, and for large systems in general the result will be computationally intractable.

In rule-based systems [4, 5] diseases are modelled as facts with associated certainty factors. The certainty factors denote the current belief in the facts, and they are used to obtain a mutual ranking of the diseases. One of the disadvantages of certainty factors is that they are relative measures, as opposed to probabilities that have an independent interpretation. There are other concerns regarding certainty factors, and the strong modular nature of knowledge and inference makes it impossible to give a meaningful

interpretable measure for a combination of diseases based on the measures of individual diseases [6]. This, and other shortcomings connected to the use of rule-based systems, such as the lack of nonmonotonic reasoning (see e.g. [7]), render the construction and maintenance of rule-based systems for large complex domains with inherent uncertainty at least doubtful, and such systems will seldom posess the ability to diagnose multiple diseases.

We claim that a global coherent view of large domains is crucial and we see no obvious alternative to causal probabilistic networks to obtain this. In the MUNIN system [8, 9, 10] for diagnosis of neuromuscular diseases, we have modelled diseases as (groups of) separate variables. As pointed out this enables, at least in principle, the diagnosis of multiple diseases, but the cost is increased computational complexity. Despite various tricks to reduce this problem, such as e.g. divorcing of parents [11], the complexity apparently limits the general applicability of the formalism. In this paper we propose a method to manage the complexity of large systems by a partitioning of a large domain into a number of smaller subdomains. A global solution to the diagnostic problem is obtained by iterating over increasingly complex subdomains, where the inference is limited to consider only plausible diagnoses. We illustrate the method by reporting an example from MUNIN, where a patient with three diseases is diagnosed.

2. An example of medical reasoning

The problem of diagnosing multiple diseases arises when more diseases have (partially) overlapping findings. Consider as a simple example a domain where we have the possible diseases FLU and THROAT INFECTION and the symptoms FEVER, SORE THROAT and SORE MUSCLES. FEVER is an indication for both diseases, SORE THROAT is an indication for THROAT INFECTION and SORE MUSCLES is an indication for FLU. Figure 1A shows the initial state for a CPN model of this domain, created by the HUGIN system [12].

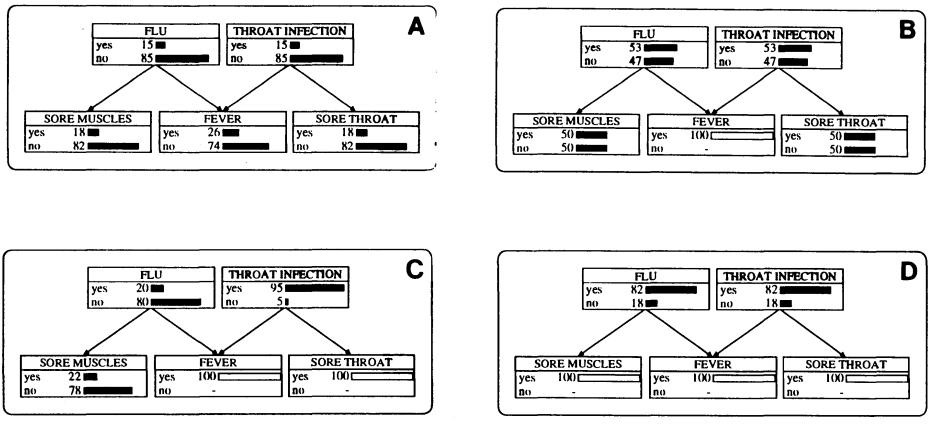

Figure 1. A CPN for a simple example with two diseases and three symptoms. A: Initial configuration of the system. B-D: Revised beliefs based on evidence of respectively FEVER, SORE THROAT and SORE MUSCLES.

If we meet a patient with FEVER we increase our belief in both FLU and THROAT INFECTION (Fig. 1B), and if we subsequently find that the patient has a SORE THROAT this confirms our belief that the problem is actually a THROAT INFECTION. This in turn explains the FEVER and we decrease our belief that the patient has got a FLU (Fig. 1C). This situation is known as nonmonotonic reasoning - although we meet positive evidence for a disease, the findings are explained away by a competing hypothesis. If, on the other hand, the patient also complains of SORE MUSCLES, he probably has both a FLU and a THROAT INFECTION (Fig. 1D), and in this case there is no need of explaining away, FEVER acts as positive evidence for both diseases. In this simple example there is no problem with the actual computation, but in realistic domains the number of diseases with overlapping findings is much more exhaustive, and it becomes impossible to create a computational structure for the complete domain. The point is that this is seldom necessary. Hardly any patient is examined for all possible symptons in a complex domain. If we alone have knowledge of a SORE THROAT in our simple example, we need not consider FLU, as this disease is independent of what is known in the domain. It is only when we have knowledge of common findings (which makes the diseases conditionally dependent) we need to consider diseases simultaneously.

3. A method for diagnosing multiple diseases

Our method for diagnosing multiple diseases stems from the line of reasoning outlined above. The approach is iterative, where the ith step considers combinations of i diseases. Thus in the first step we consider only single diseases, in the second step we consider all pairs of diseases on so on. The diseases considered in step i+1 are a chosen among those considered in step i. Different criteria for inclusion in the next step can be used, for example that the probability for the presence of the disease is higher than the prior probability for the disease, or that the probability for the disease exceeds some fixed level. In each cycle the inference becomes increasingly complex, but the number of possible diseases is reduced, such that the complexity becomes manageable. The procedure stops when the number of relevant dependent diseases is less than the number of the next cycle.

Models that include only a subset of the diseases are easily obtained from the complete model. We simply mark the diseases under consideration and generate a subdomain consisting only of the marked diseases and their descendants. In many common cases it is possible to gain further efficiency by various tricks for splitting domains up into several subdomains followed by a combination of the partial results to a global solution.

In the example in Fig. 1C we would consider FLU and THROAT INFECTION (one at a time) in the first cycle, find positive evidence for both, and thus continue with a second cycle where both the diseases are considered simultaneously. In the second cycle FLU is ruled out, and we conclude that the patient suffers from a THROAT INFECTION.

4. Diagnosing multiple diseases in MUNIN

We have implemented the proposed method for diagnosis of multiple diseases in the MUNIN system, designed to assist in the diagnosis of neuromuscular diseases. The diseases covered by MUNIN includes general muscle and nerve diseases and a range of local nerve lesions. The general muscle diseases proximal myopathy, myotonic dystrophy and myasthenia are modelled as single variables with up to seven states that characterise the nature of the disease (e.g. pre- or postsynaptic). Local nerve lesions are characterised by three variables that detail their severity, pathology and time course. The primary attribute of these diseases is the severity, it is only if this variable takes a value different from "no" that the disease is considered and the other attributes become relevant. Diffuse neuropathies describe general nerve diseases. They are characterised by the same attributes as the local nerve lesions and in addition by two variables describing type and distribution. This makes it possible to describe a large group of disorders, including such different diagnoses as Guillain Barre and motor neuron diseases.

The possible investigations included in MUNIN are grouped in so-called anatomical structures. Anatomical structures, which represent sensory nerves or motor nerves together with the muscles they innervate, are described by subnets with typically more than one hundred nodes. Each subnet consists of layers of variables. Some layers are directly interpretable, they describe e.g. pathophysiology or examinations, others are merely included for technical reasons. Addition of new anatomical units is relatively unproblematic since they add only linearly to the overall complexity, but addition of diseases creates cycles in the network, and these cycles increase the complexity exponentially within the subnets of the affected anatomical units. Two diseases can always be managed, but for most combinations of three or more diseases affecting the same anatomical unit, other methods, such as "cutset conditioning" [13], have to be exploited. This is practically impossible for the total model, but obtainable for the smaller subdomains. The good news is that patients with more than two diseases affecting the same anatomical unit are extremely rare.

The application of the method to a patient suffering from both a diabetic polyneuropathy and a bilateral carpal tunnel syndrome resulted in 17 possible diagnoses, each having an increased probability compared to the prior. Most of these were false positives, they were suggested because they were the only possible explanations of abnormal findings in the anatomical units they affect. In the second iteration 14 local nerve lesions were ruled out, the diffuse neuropathy explained away the findings that supported these lesions. In addition to the diffuse neuropathy two local lesions affecting respectively the right and the left median nerve were left. All proposed diagnoses remained after the third cycle. The next cycle is the fourth, thus the procedure stops and the final correct diagnosis consisting of three lesions is concluded.

5. Conclusion

We have presented a new method for diagnosis of multiple diseases in complex domains by causal probabilistic networks. The method have been implemented in the MUNIN system and we have found it to be the key to an extension of MUNIN to include all

commonly examined anatomical structures. The method is applicable in general, provided that no false negative conclusion is generated by competitive (possible false positive) conclusions to which it is compared. We have not encountered this situation in practice, and we conjecture that the method enables construction of large medical expert systems in other complex areas.

References

[1] D.E. Heckerman, E.J. Horvitz and B.N. Nathwani, Towards normative expert systems, part I: The Pathfinder project. *Methods of Information in Medicine* **31** (1992) 90-105.

[2] J. Pearl, Probabilistic reasoning in intelligent systems. Morgan Kaufmann, 1988.

[3] F.V. Jensen, Introduction to Bayesian networks. UCL Press, 1996.

[4] B.G. Buchanan and E.H. Shortliffe, Rule-based expert systems. Addison-Wesley, 1984.

[5] J. Giarratano and G. Riley, Expert systems - Principles and programming, 2nd edition. PWS Publishing Compagny, 1993.

[6] D.E. Heckerman and E.H. Shortliffe, From certainty factors to belief networks. *Artificial Intelligence in Medicine* **4** (1992) 35-52.

[7] S. Andreassen, F.V. Jensen and K.G. Olesen, Medical expert systems based on causal probabilistic networks. *International Journal of Biomedical Computing* **28** (1991) 1-30.

[8] 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 (Ed.), Computer-Aided Electromyography and Expert Systems. Elsevier, 1989.

[9] S. Andreassen, B. Falck and K.G. Olesen, Diagnostic function of the Microhuman prototype of the expert system MUNIN. *Electroencephalography and Clinical Neurophysiology* **85** (1992) 143-157.

[10] S. Andreassen, A.Rosenfalck, B. Falck, K.G. Olesen and S.K. Andersen, Evaluation of the diagnostic performance of the expert EMG assistant MUNIN. *Electroencephalography and Clinical Neurophysiology*, to appear.

[11] 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* **3** (1989) 385-403.

[12] S.K. Andersen, K.G. Olesen, F.V. Jensen and F. Jensen, HUGIN - A shell for building Bayesian belief universes for expert systems. Proceedings of the 11th International Joint Conference on Artificial Intelligence, 1080-1085. Morgan Kaufman, 1989. Reprinted in: G. Shafer and J. Pearl (Eds.), Readings in Uncertain Reasoning, Morgan Kaufman, 1990.

[13] J. Pearl, A constraint-propagation approach to probabilistic reasoning. In: L.N. Kanal and J.F. Lemmer (Eds.), Uncertainty in artificial intelligence, North-Holland, 1986.