”PROMEDAS”
a prototype decision support system for medical diagnosis

SNN
Foundation for Neural Networks
Nijmegen, The Netherlands

UMCU
University Medical Centre Utrecht
Utrecht, The Netherlands

Summary

The use of patient-specific Decision Support Systems (DSS) may improve the quality and efficiency of health care, while reducing its costs at the same time. The adoption of such a system is largely compatible with the principles of ”Evidence Based Medicine” and patient oriented care.

”Promedas” is a prototype of a diagnostic DSS based on a large causal probabilistic network, using recently developed computational techniques. Within the next three years it will cover a significant diagnostic area in the field of haematology and endocrinology.

The system intends to support diagnosis making in the setting of the outpatient clinic and for educational purposes. Its target-users are general internists, super specialists (i.e. cardiologists, rheumatologists), interns and residents, medical students and others working in the hospital environment.

The system offers diagnostic advice. In active decision mode, it supports the diagnostic process by indicating the most useful next step in the diagnostic process. The system is intended to be transparent and disposes of several explanatory and clarifying facilities, including the availability of the appropriate references to the literature. Integration with a Hospital Information System and an Electronic Patient Record in the future will augment its acceptance and will facilitate its implementation.

Contents

1 Introduction

Patient-specific diagnostic Decision Support Systems (DSSs) might be extremely useful in health care because they are able to improve the availability and accessibility of knowledge, resulting in quality improvement, increase of efficiency and reduction of costs. However, up to now, these systems have not yet entered daily clinical practice for a variety of reasons.

Why use decision support in medicine?

Decisions made by physicians are arbitrary and highly variable (within one physician and between physicians) and often lacking explanation or ”rationalisation” [1, 2]. Problems in modern medicine are often very complex, but evidence for the best choice to be made is often lacking. Clinical examples of this phenomenon in diagnosis making are abundant and easy to understand.

The body of potentially useful knowledge that is relevant to even a relatively narrow diagnostic area may be too large to make the optimal (diagnostic) decision on the spot. Ironically, modern information technology (especially through the Internet) increases the amount of available knowledge even more, probably further complicating this situation. Moreover, individual patients need ”individualised” decisions, because their characteristics differ from the ”average” and because of their individual wishes[3]. Apparently, individualising the general results of research may be cumbersome and time consuming, while on the other hand modern medical practice demands efficiency, cost-effectiveness and high technical quality.

The derivation of diagnostic protocols is a main problem in health care. But in some environments diagnostic support is simply not likely to influence physician’s decisions, e.g. on a neurological intensive care unit, since the diagnosis is often obvious [4]. In contrast, general internal medicine covers an enormous range of sometimes relatively rare diagnostic categories. Hence the tendency of medicine to be differentiated in super-specialisations, (e.g. gastro-entrologists, rheumatologists, cardiologists. etc.). A diagnostic DSS covering general internal medicine may be appreciated by both generalists and super-specialists alike: by the generalist because this field of work typically covers a very broad range of diagnoses, by the super-specialist because

he/she may not feel completely at ease outside his/her specific field of expertise.

It is readily understandable that the above comprises an enormous task and challenge for modern medicine in general and individual doctors in particular, illustrating the need for decision support techniques. Obviously, automated DSSs may be very promising from a theoretical point of view.

What are the problems in current decision support systems?

The currently available systems (e.g. Meditel [5], QMR [6], Dxplain [7] and Iliad [8]) have not yet been very successful. Certainly their use is still not widespread and not established in daily routine. A variety of reasons may be responsible for this:

Lack of accuracy

Current systems that intend to cover a broad diagnostic domain of medicine lack diagnostic accuracy [9, 10]. This is not due to the method that is used, but rather due to the levels of detail (e.g. diagnostic categories at the level of ICD-9) and completeness in the knowledge base. Systems that are based on detailed modelling of knowledge may have a good performance. Up to know, however, such systems are restricted to a relatively narrow field [11, 12]. The crucial problem with a consistent detailed model covering a broad domain is that it would be computational intractable. The next section will discuss how can be dealt with this problem.

Lack of transparency

In the era of evidence based medicine the advice of “a black box” is unacceptable. An advice must be motivated and preferably accounted for on the basis of research published in the peer reviewed literature.

Users attitude

A subset of (potential) users may have a misunderstanding about what computers can and cannot do for them. Generally, DSSs need educated and responsible users, who are able to interpret the advise given and estimate its merit [13]. This, however, is not exclusively a matter of users attitude. Producers of decision support tools should take this issue into account as well, especially when designing the user interface and deciding which facilities are needed.

Lack of integration of information

Patient specific decision support needs input data from several sources. A DSS will generate new information (e.g. a diagnostic advise) through inference. For this it uses specific information about a patient, given “patient-independent” knowledge (e.g. about diagnosis making) stored in the knowledge base of the DSS. Integration of information, multiple

usability of patient data, integration of databases and knowledge bases are common problems when using a heterogeneous Hospital Information System (HIS). Unfortunately, the completeness of patient information and the accuracy and level of detail of diagnoses stored in a HIS is in general very poor [14].

Lack of a controlled terminology

This is a problem that might not even be solved completely in the near future. Most standard classification systems are at a general level [15, 16], thus lacking the required detail, or specialised [17] and therefore too limited to meet the needs for a DDS covering a broad domain.

Careful introduction

Introduction of a DSS should be done as careful and thorough as is use for drugs that are new on the market. After implementation, the use of a DSS will need meticulous surveillance resulting in further improvement. Maintenance is needed to keep up with the last results of research.

Why develop a new diagnostic decision support system?

In conclusion, modern medicine is in need of computerised decision aids both to meet its own high standards and to keep pace with the stage of development in other domains such as manufacturing or the services industry. Taking into account the need for decision support and *diagnostic* decision support in particular, we strongly believe that a diagnostic DSS for a broad medical domain is viable and, eventually, marketable.

To avoid a ”gold rush style” in the search for these tools, the foremost thing to do is the development of safe and sound methods. The expertise of our multidisciplinary group primarily focuses on three main parts of the methodology typically needed in the development of decision support tools:

- *Modelling and inference algorithms that are able to deal with large complex systems.*
- *Knowledge modelling in the medical domain.*
- *User aspects.*

2 The Computational Engine

A diagnostic decision support system offers diagnostic advice for a diagnostic problem regarding an individual patient. The system needs a representation of medical knowledge, i.e. a model, and it must be able to do reason (i.e. compute) with patient specific data on the basis of this model. Belief networks are typically well suited for the representation of medical knowledge and for reasoning with this knowledge. Unfortunately, the large and complex belief networks that are required to model a large medical domain are intractable for exact computation. Recently developed approximate methods, possibly combined with other techniques, may be the answer to this problem.

Belief networks

Reasoning in the medical domain is a typical example of reasoning with uncertainty. This uncertainty has different sources: uncertain, incomplete or even missing patient information, uncertainty in medical tests, and the inherent uncertainty in physiological processes. Clearly, the model on which a DSS is based should be able to deal with these uncertainties.

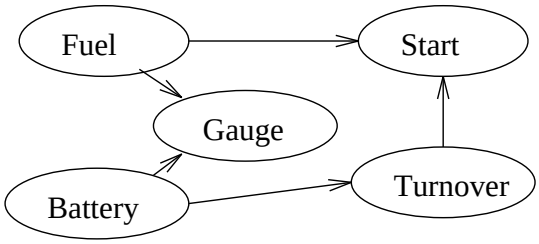

Figure 1: Belief network. Variables are represented by the nodes. Conditional independencies are represented by the arrow structure. For instance, the state of `Start` is independent of the state of `Gauge`, given the state of `Fuel` and `Turnover`.

The different DSSs that have been developed sofar use a variety of modelling approaches which can be roughly divided into two categories: rule-based approaches and probabilistic methods. The rule based approach can be viewed as an attempt to simplify the probabilistic approach in order to reduce computational complexity. The probabilistic approach has the advantage of mathematical consistency and

correctness. Belief networks [18] in particular provide a powerful and conceptual transparent formalism for probabilistic modelling.

A belief network – also called Bayesian networks or causal probabilistic networks – is a graph of nodes and arrows. See fig. 1. The nodes represent random variables whereas arrows between nodes represent direct influences. For each node a conditional probability table quantifying the effect of the parent nodes, i.e. the nodes that directly point to it. The total graphical structure specifies a probability distribution over the state space of all variables. Belief updating is done using the rules of probability (Bayes theorem).

Belief networks have several distinguishing features:

Transparency
: A belief network has an appealing, transparent and intuitively clear structure which can be graphically visualised. Expert knowledge can be made explicit, while users can have insight into how the system operates. Belief networks have a more modular representation of uncertain knowledge than rule-based systems. This makes them easier to maintain.

Accuracy
: Because all relations between variables are described by the rules of probability, there are no assumptions made by the methodology. The definition of the variables and the structure of the network contain all assumptions in the network. If the accuracy of the network is too much hindered by a particular assumption this can easily be removed by restructuring the network.

”Hidden” variables
: Large numbers of observable variables can be related in the model via ”hidden”, e.g. pathophysiological, variables. Besides the modelling advantages and the improvement of transparency, the use of this type of variables makes the network in general less complex and therefore less sensitive to over-fitting.

Learning
: Using standard learning algorithms for belief networks, the system can be fine-tuned using historical patient data and learn further ”from experience” on the basis of prospectively gathered patient data.

The progress that has been made during the last decade in exact computation in belief networks makes the argument in favour of rule based approaches less and less persuasive - at least for relatively small and simple models. Indeed, most modern approaches for medical diagnosis are based on the probabilistic approach. A drawback is that complex probabilistic models are intractable for exact computation.

New techniques to handle intractable networks

To deal with the complex belief networks that are required in medical modelling one has to rely on approximate computations. Recently, variational methods for approximation are

becoming increasingly popular [19]. An advantage of variational methods is that they provide guaranteed bounds on the level of approximations in contrast to stochastic sampling methods, which may yield unreliable results due to finite sampling times.

Promising results have been reported in an application of variational methods in the two-level QMR network (see fig. 2) [20]. Recently we have developed variational methods that are applicable to more complex (and more realistic) structures These methods exploit tractable substructures in the network, and therefore may give more precise bounds [21, 22, 23]. Variational methods are typically well applicable to large, detailed belief networks for medical diagnosis constructed by human experts, since these networks have typically a modular structure (see fig. 3)[24]. In "Promedas" we will combine exact methods, (new) variational methods and possibly other techniques to make the DSS as accurate as possible in a feasible amount of computing time.

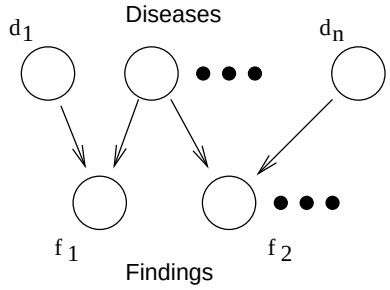

Figure 2: Two-level graphical structure of the QMR belief network. The dependencies between the diseases and their associated findings are modelled via noisy-OR gates. Despite the simple structure, this network is intractable.

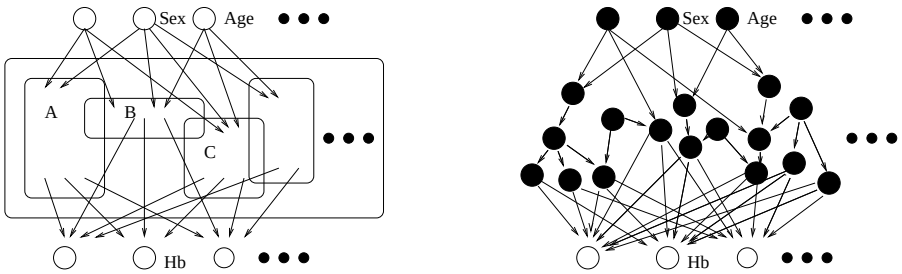

(a) Modular structure

(b) Graphical structure

Figure 3: Modular and graphical network structure. Left: modular structure of the network. A, B, C . . . represent (overlapping) sub-domains. Each sub-domain is modelled by a number of nodes (cf. right figure) representing variables that are relevant in that domain. The upper nodes, e.g. 'sex' and 'age' represent common ancestors of nodes in several sub-domains. The lower nodes, e.g. 'Hb' represent common children of nodes in several sub-domains (e.g. related to anaemia). Right: underlying graphical structure of same network. Filled circles: nodes in sub-domains and their common ancestors. Open circles: common children. The common children depend on their direct parents via noisy-OR gates.

3 The Medical Knowledge

Belief networks and their algorithms provide a powerful engine for medical decision making. The first step in the actual development of a diagnostic DSS is the definition and modelling of medical knowledge. The domain is defined, and the knowledge concerning this domain is acquired and finally represented in the network. The acquired knowledge must be well documented, in order to remain accessible and understandable.

Domain definition

A domain of medical knowledge may be defined in several ways:

- a) Problem oriented: e.g. anaemia or thyreotoxicosis, This is generally the way in which problems that prompt for clinical diagnosis making present themselves
- b) Textbook-style, e.g. a group of diseases that affect a specific organ system: disorders of the haematopoietic system, disorders of the respiratory system, endocrinology etc. Information in ”traditional” sources of information like textbooks is generally arranged in this way.

A problem oriented taxonomy of domains is most suitable for a diagnostic decision support system because it will best facilitate a patient-specific approach. ”Anaemia” is chosen as an example of a problem oriented domain definition that is relevant from a users point of view in a clinical setting. ”Anaemia” may serve as a model for general internal medicine (and probably medicine as a whole) because of its complexity.

A domain is subsequently subdivided in sub-domains, so called modules. A module consists of a group of diagnostic categories that share an important common pathogenetic or pathophysiological mechanism. The subdivision of a domain in modules is based on the information in textbooks and on experience, but in fact the subdivision is arbitrary and the modules are overlapping. Modules are convenient to keep information manageable in the process of designing the knowledge models.

Knowledge acquisition and representation

Domain knowledge is acquired from the literature[1].

After acquisition of knowledge from traditional sources the knowledge is ”translated” into a model represented by a causal probabilistic network. The knowledge is characterised by:

- A diagnostic repertoire, intended to be exhaustive in that context
- A repertoire of findings, including all tests that are potentially relevant for that sub-domain
- Causal relationships at a pathophysiological level with a high level of detail
- (Conditional) probability tables quantifying these causal relationships

The (conditional) probabilities are determined on the basis of data in the literature or on ”educated guesses” based on local statistics and experience if no data from the literature are available.

Auxiliary databases

Auxiliary databases are needed to store the domain knowledge and other relevant information that is not are not explicitely contained in the model. These databases will include in particular:

- The acquired domain knowledge and its sources in the literature.
- The definitions of terms and concepts that are used in “Promedas”, where possible, linked with corresponding terms in Unified Medical Language System. [2]
- Specifications of the discretisation of variables.
- Additional information about special circumstances that influence test results if not taken into account in the model.
- Local prevalences of diseases. Note that these have a direct impact on the prior probabilities on the model.
- Cases that are validated by means of consensus by a panel of domain experts. The case database can be filled prospectively. It can be used for education and research purposes, as well as to improve the system by retraining.

[1] For instance, Harrison’s principles of internal medicine, Cecil’s textbook of medicine, The medical clinics of North America, Up to Date, The Cochrane Library and relevant journal articles retrieved with the help of Medline.

[2] Ideally, the definitions and terms in “Promedas” would be conform a standard classification systems. Unfortunately, a standard classification systems covering the majority of the terms and concepts used in internal medicine at a sufficiently detailed level are not yet available.

Users will have (partly) access to these databases via help functions and therefore may better interpret the advise given by the system and estimate its merit. On the other hand, these databases are essential for maintenance of the system. In particular, cut-off points and local prevalences require adjustments depending on local circumstances.

4 The User Interface

Even if the DSS is equipped with a powerful engine, and loaded with deep knowledge, if the user isn't able -or willing- to interact with the system, the DSS is worthless. The user interface is of great importance. The first step in developing a user interface is to define the functionality of the system. With this in mind, a design can be made.

Functionality

"Promedas" disposes of a graphical user interface that can be used in several modes:

Diagnostic advise mode

The user seeks patient-specific diagnostic advice concerning a problem presented in clinical practice. Known findings regarding the case are put into the system, including historical data, signs and symptoms, laboratory results and results of imaging techniques. The system offers diagnostic advice by proposing a list of differentiated diagnostic possibilities along with their chances as well as a list of potentially involved pathophysiological processes based on the data.

The user makes a conclusion about which diagnoses and pathophysiology is most likely to be involved and which may be excluded because they seem very unlikely.

Active decision in clinical diagnosis making

The user seeks advice about which test (or set of tests) has the highest impact on the diagnostic process in a particular stage. The advice given includes the discriminative power of a test, but also its costs, its burden for the patient and the organisation and risks of the test involved.

These two modes can also be used for educational purposes, especially for interns and residents. In medical schools, students can use the system as well to study diagnosis making using the validated cases included in the database as virtual patients.

Design

The system is designed for use on a clinical workstation in the future, preferably integrated with other parts of the Hospital Information System and an Electronic Patient Record (EPR) , in order to enable data sharing. However, it may be used on a stand alone PC as well. Ideally it may be linked in addition with other sources of computerised information e.g. through the Internet, such as Medline.

Aiming for optimal acceptance of the system, the user interface of ”Promedas” has the following characteristics formulated from the user’s point of view:

- Attractive graphical design
- User friendliness, taking also in account users with relatively little computer experience
- Convenient to use in the work environment
- Easy importation of data (check lists, interactive), preferably automated if possible (especially lab results)
- Flexibility: there are several modes available depending on the domain and user preferences. Its use is basically suitable for any medical domain and adaptable to local standards and practice.

Starting-point for the design of the screen display of ”Promedas” was the current conventional hospital patient record, which use is wide-spread and accepted, if not obligatory. Data are organised in groups in a similar familiar way.

In case of an EPR, the screen display of ”Promedas” might be adapted accordingly.

5 Evaluation

Finally it is important to evaluate the usefulness of the system. During the development of the system, intermediate evaluation results are used to improve the system. When the system is fully developed, final "real life" evaluation results are used to assess the viability of the general method

During its development the validity and the performance of the system will be tested by the experts in the project team and by clinical experts in endocrinology and hematology at Utrecht University Hospital. The results will be used to improve the system performance.

Assessment procedures will be set up to evaluate the diagnostic performance of the system, and to compare it with other systems. A panel of external experts will determine a set of "gold standard" cases by reaching diagnostic concensus on a sufficiently large set of challenging cases from "real life".

Evaluation will include assessment of the performance of the system alone in comparison with the performance of a group of target users with and without using the advise of the system. The performance will also be compared with existing diagnostic DSSs (e.g. QMR and DXPlain).

"Promedas" will be implemented step-wise by installation at workstations at the outpatient clinic of interne medicine at UMCU and possibly one or two affiliated affiliated community hospitals. Its usefulness in daily clinical practice will be evaluated by closely monitoring and by structured pre-set questionnaires for (target) users.

6 Software Development

A C++/JAVA software package will be developed to implement “Promedas”
This package includes:

- A library of approximate, exact and hybrid inference methods
- A graphical interface for modelling of domain knowledge by domain experts
- A graphical user interface for diagnostic decision support that can be used on a clinical workstation by physicians
- A data interface for communication with the Hospital Information System to obtain patient data in electronic format
- A facility to retrain the network with prospective data.

A Case History

This validated case from clinical practice was presented to “Promedas” retrospectively:

Case description

A 67 years old male is seen because of type 2 diabetes at 3 months intervals at the outpatient department of diabetology starting 1986. He is currently being treated with insulin since 2 years. He has been diagnosed with incipient nephropathy because of microalbuminuria. He is treated for hypertension and for combined hyperlipidaemia. There are clinical signs and symptoms attributed to diabetic neuropathy since about one year.

Previous history

1967 accident (femur fracture, multiple ribs fractured)
1986 cholecystectomia

Current medication

Insulin, Simvastatin, Gemfibrozil, Enalapril

Smoking habits

Stopped smoking 10 years ago

Alcoholic beverages:

1 unit daily

Routine laboratory tests

Hb 7.6 mmol/L (decreased), creatinine 102 umol/L (normal), random glucose 11.2 mmol/L, HbA1c 8.1%,ALAT, AF, CK: normal

Prescheduled next visit

History during his next visit:

Slight loss of appetite during the past several months. Bloating after dinner. Defecation pattern somewhat irregular (since several years).
No worsening of symptoms of his neuropathy during the previous 6 months.

Physical examination

Body weight 83 kg (unchanged). Somewhat pale looking.
Not noteworthy otherwise.

Additional and repeated tests

Hb 6.9 mmol/L(decreased), MCV 102 fL (increased), Leukocyte (WBC) and Platelet counts both slightly decreased
Reticulocyte count: 2 promille (decreased)
Total bilirubin 20 umol/L (slightly increased), bilirubin direct ¡ 20% (normal)
LDH 724 U/mL (increased)
Cholesterol decreased from 6.6 mmol/L to 3.7 mmol/L during the previous 6 months
Triglycerides 1.8 mmol/L
ASAT, ALAT, AF, gammaGT: all within normal limits
TSH 2.3 mU/L (normal)
Ferritin and Folic Acid both normal
Vitamin B12 105 pmol/L (decreased)

Gastroscopy was normal

Next visit

Subsequently ordered additional tests

Periferal smear: showed hypersegmentation of neutrophils and macroovalocytic erythrocytes
Schilling test: 1st and 2nd phase both abnormal
Antiparietal cell antibodies: positive
Anti-intrinsic factor antibodies: positive
Fasting serum gastrin: increased

Diagnosis of the physician in charge:

Anaemia due to vitamin B12 deficiency, probably pernicious anaemia.

The diagnostic advise of "Promedas"

The diagnostic advise of "Promedas" for this case would have been Pernicious Anaemia (probability 93%). Other differential diagnostic options all have a probability below 5% and therefore may be considered as implausible. See the colour graph (fig. 4) of the causal probabilistic model of anaemia due to vitamin B12 deficiency.

For this diagnostic advise "Promedas" would have needed only part of the described tests that have been ordered by the physician in charge. Moreover, in active decision mode tests would have been ordered in a different, more efficient, hierarchy.

Figure 4: Probabilistic model of anaemia due to vitamin B12 defi ciency. Blue bars denote the clamped states of the known (measured) variables of this case. Red bars denote the state probabilities of the unknown variables, conditioned on the clamped states.

Hier komt de fi guur

B The "Promedas" Team

The development of "Promedas" is a joint project by the Foundation for Neural Networks (SNN) and the University Medical Centre Utrecht (UMCU). This project is financially supported by the Dutch Technology Foundation (STW).

The "Promedas" team members are

Dr. H.J. Kappen	SNN	Project management
Dr. W.A.J.J. Wiegerinck	SNN	Artificial Intelligence
M. Nijman, MSc	SNN	Software Engineering
Dr. J.P. Neijt, MD	UMCU	Internal medicine
E.W.M.T. ter Braak, MD	UMCU	Internal medicine
Dr. Y.L. O	UMCU	Medical Informatics
W.J.P.P. ter Burg, MSc	UMCU	Medical informatics

WWW:

http://www.mbfys.kun.nl/snn/Research/promedas/

Correspondence Address:

Foundation for Neural Networks
University of Nijmegen
PO Box 9101
6500 HB Nijmegen
The Netherlands
Tel.: +31-(0)24 3614241
Fax.: +31-(0)24 3541435
e-mail:bert@mbfys.kun.nl
http://www.mbfys.kun.nl/snn

References

[1] D. M. Eddy. The challenge. *Journal of American Medical Association*, 263:287–290, 1990.

[2] M. Berg. *Rationalizing Medical Work, Decision-support Techniques and Medical Practices*. MIT press, 1997.

[3] R.J. Lilford, S.G. Pauker, D.A. Braunholz, and J. Chard. Decision analysis and the implementation of research findings. *BMJ*, 317:405–409, 1998.

[4] B. Brigl, P. Ringleb, T. Steiner, P. Knaup, W. Hacke, and R. Haux. An integrated approach for a knowledge-based clinical workstation: Architecture and experience. *Methods of Information in Medicine*, 37:16–25, 1998.

[5] Meditel, Devon, Pa. *Meditel: Computer assisted diagnosis*, 1991.

[6] CAMDAT, Pittsburgh. *QMR (Quick Medical Reference)*, 1992.

[7] Massachussetts General Hospital, Boston. *DXPLAIN*, 1992.

[8] Applied Informatics, Salt Lake City. *ILIAD*, 1992.

[9] AS Elstein, CP Friedman, FM Wolf, G Murphy, J Miller, P Fine, P Heckerling, T Miller, J Sisson, S Barlas, K Biolsi, M Ng, X Mei, and T Franz. Performance of four computer-based diagnostic systems. *N-Engl-J-Med.*, 330(25):1792–6, 1994.

[10] ES Berner, JR Jackson, and J Algina. Relationships among performance scores of four diagnostic decision support systems. *J-Am-Med-Inform-Assoc.*, 3(3):208–15, 1996.

[11] D.E. Heckerman, E.J. Horvitz, and B.N. Nathwani. Towards normative expert systems: part I, the Pathfinder project. *Methods of Information in medicine*, 31:90–105, 1992.

[12] D.E. Heckerman and B.N. Nathwani. Towards normative expert systems: part II, probability-based representations for efficient knowledge acquisition and inference. *Methods of Information in medicine*, 31:106–116, 1992.

[13] J.J. Dijkstra. *On the use of computerised decision aids*. PhD thesis, State University of Groningen, 1998.

[14] W. Wiegerinck, W. ter Burg, E. ter Braak, P. van Dam, Y. O, Neijt J., and H Kappen. Inference and advisory system for medical diagnosis. Technical report, SNN, 1997.

[15] WHO. *International Classification of Diseases, 10th revision, Clinical Modification ICD-10*, 1992.

[16] J. Boersma, R.S. Gebel, and H. Lamberts. *International Classification of Primary Care, Short titles (vertaald in het Nederlands) en Nederlandse subtitels*. Nederlands Huisartsen Genootschap, 1993.

[17] WHO. *International Classification of Diseases for Oncology*.

[18] J. Pearl. *Probabilistic Reasoning in Intelligent systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, Inc., 1988.

[19] M. I. Jordan, editor. *Learning in Graphical Models*, volume 89 of *NATO ASI, Series D: Behavioural and Social Sciences*. Kluwer, 1998.

[20] T.S. Jaakkola and M.I. Jordan. Variational methods and the QMR-DT database. MIT Computational Cognitive Science Technical Report 9701, Massachusetts Institute of Technology, 1997.

[21] W. Wiegerinck and D. Barber. Mean field theory based on belief networks for approximate inference. In L. Niklasson, M. Bodén, and T. Ziemke, editors, *ICANN'98: Proceedings of the 8th International Conference on Artificial Neural Networks, Skövde, Sweden, 2-4 September 1998*, volume 2, pages 499–504, London, 1998. Springer.

[22] W. Wiegerinck and D. Barber. Variational belief networks for approximate inference. In H. La Poutre and J. van den Herik, editors, *Proceedings of the Tenth Netherlands/Belgium Conference on Artificial Intelligence (NAIC'98)*, pages 177–83, Amsterdam, 1998. CWI.

[23] D. Barber and W. Wiegerinck. Tractable undirected approximations for graphical models. In L. Niklasson, M. Bodén, and T. Ziemke, editors, *ICANN'98: Proceedings of the 8th International Conference on Artificial Neural Networks, Skövde,Sweden, 2-4 September 1998*, volume 1, pages 93–98, London, 1998. Springer.

[24] W. Wiegerinck *et. al.* Variational approximations in a broad and detailed probabilistic model for medical diagnosis. In H. La Poutre and J. van den Herik, editors, *Proceedings of the Tenth Netherlands/Belgium Conference on Artificial Intelligence (NAIC'98)*, pages 185–91, Amsterdam, 1998. CWI.