

Supporting the Construction of Explanation Models and Diagnostic Reasoning in Probabilistic Domains

Olaf Schröder

Claus Möbus

Jörg Folckers

Heinz-Jürgen Thole

OFFIS Institute, Oldenburg, Germany

E-Mail: {schroeder, moebus, folckers, thole}@informatik.uni-oldenburg.de

Abstract: MEDICUS (modeling, explanation, and diagnostic support for complex, uncertain subject matters) is an intelligent modeling and diagnosis environment designed to support the construction of explanation models and diagnostic reasoning in domains where knowledge is complex, fragile, and uncertain. MEDICUS is developed in collaboration with several medical institutions in the epidemiological fields of environmentally caused diseases and human genetics. Uncertainty is handled by the Bayesian network approach. In *modeling*, the user creates a Bayesian network for the problem at hand, receiving help information and explanations from the system. This differs from existing reasoning systems based on Bayesian networks, i.e. in medical domains, which contain a built-in knowledge base that may be used but not created or modified by the user. MEDICUS supports *diagnostic reasoning* by proposing diagnostic hypotheses and recommending examinations. In this paper we will focus on the modeling component of MEDICUS.

1. Introduction

Medical diagnosis is a reasoning and problem solving task that can be quite difficult [Barrows & Tamblyn 1980], [Boshuizen & Schmidt 1992], [Elstein et al. 1978], [Elstein & Bordage 1980], [Patel & Groen 1986]. One prerequisite is the availability of explanatory models of diseases, their etiologies, and symptoms associated with them. This is especially true in young medical subdomains with particularly complex, interrelated, fragile, and uncertain knowledge. Two examples are the epidemiology of diseases caused by environmental influences, like pollution, and of diseases caused by human genetic defects. Much of the knowledge in these domains is not yet available in a systematic way, and clear-cut taxonomies and explanatory models of diseases have not been developed yet. Still these domains are getting increasingly important. This is reflected by the fact that they receive increasing attention in medical science at university as well as in postqualification courses for physicians. In the domains of environmental medicine and human genetics, explanatory models of diseases are important i) in epidemiological research in order to systematize information from several resources, like epidemiological studies and clinical cases, ii) to provide a base for making e.g. environmental examinations more efficient, and iii) in the context of medical training and qualification. Thus the learner should have an opportunity to *actively construct* models of diseases, their possible causes, and the symptoms associated with them, and to evaluate the consequences of these models. In this way the learner acquires and uses the knowledge necessary for diagnostic reasoning. Furthermore, the learner should have an opportunity to *actively perform* diagnostic reasoning and to apply diagnostic strategies.

Existing computer-based systems supporting medical reasoning (i.e. MYCIN [Shortliffe 1976], CASNET [Weiss et al. 1978], PIP [Szolovits & Pauker 1978; 1993], INTERNIST [Miller et al. 1982], ABEL [Patil et al. 1981], NESTOR [Cooper 1984], MUNIN [Andreassen et al. 1987], PATHFINDER [Heckerman 1991], see also [Barahona et al. 1995]) are primarily aimed at proposing diagnostic hypotheses, given available clinical evidence, and to suggest further diagnostic evidence gathering steps, for example, for differential diagnosis (e.g., [Heckerman et al. 1992]). Some systems, like CASNET, also generate therapeutic recommendations and have some capability to explain their reasoning steps. But none of these systems is designed to support the *creation of explanatory models* for diseases and the training of diagnostic strategies. In cooperation with several medical institutions (Health Authority of Oldenburg, Documentation and Information Center for Environmental Issues,

Osnabrück, Medical Institute for Environmental Hygiene, Düsseldorf, Robert-Koch Institute, Berlin), we currently develop MEDICUS, an intelligent modeling and diagnosis environment. The aims of MEDICUS are

- to assist a user in developing a model of perceived causes, effects, and other relationships in a domain of interest. The user may be a learner in a training context or a professional interested in creating an explanatory model summarizing epidemiological hypotheses and findings.
- to assist a user in diagnostic reasoning tasks. Again the user may be a learner in a training context, or he / she may be a professional planning, executing and evaluating for example an environmental monitoring survey (i.e., chemical analysis of air in rooms).

MEDICUS differs from existing medical expert systems by being designed to support these two activities:

- *Model construction* is supported by a linguistic model editor based on a simplified natural language, and a graphical model editor for editing Bayesian networks. After creating an initial linguistic and/or graphical model of the domain of interest, the modeler may further specify, evaluate, and revise the model at a qualitative and quantitative level.
- *Diagnosis* will also be supported qualitatively (i.e., what information is necessary in order to support or differentiate between what hypotheses?) and quantitatively (i.e., how strongly do new facts affect diagnostic hypotheses?).

Although MEDICUS is developed within the mentioned fields of medicine, our intention is that it will be applicable in general in domains of uncertain and complex knowledge. The next section gives an overview of the design decisions for the system. The third section describes the modeling component of MEDICUS in some detail. The fourth section gives a brief sketch of the state of the diagnosis support component. Conclusions and directions of further work will be sketched in the closing section.

2. Design Decisions for MEDICUS

Model construction and diagnostic reasoning can be viewed as problem solving tasks. In order to create a system designed to support problem solving in a knowledge domain, design principles are required that are based on a theory of problem solving and knowledge acquisition. For training and knowledge communication contexts, we developed design principles leading to the concept of an *Intelligent Problem Solving Environment* (IPSE, [Möbus 1995]): The learner acquires knowledge while working on a sequence of problems, actively *testing hypotheses*. This means that the learner creates solution proposals, tests hypotheses about their correctness, and the system analyzes the proposals making use of an oracle or an expert knowledge base, and provides help and explanations. The psychological foundation of our IPSE approach is the ISP-DL Theory of knowledge acquisition and problem solving (i.e., [Möbus 1995]) which is influenced by [van Lehn 1988], [Newell 1990], [Anderson 1993], [Gollwitzer 1990]. Briefly, it states that new knowledge is acquired as a result of problem solving and applying weak heuristics in response to impasses. In contrast, existing knowledge is optimized if applied successfully. Furthermore there are four distinct problem solving phases: deliberating and setting a goal, planning how to reach the goal, executing the plan and evaluating the result. The ISP-DL Theory leads to several design principles for IPSE's [Möbus 1995]. For example, firstly, the theory states that the learner will appreciate help at an impasse. So the system should not interrupt the learner but offer help on demand. Secondly, feedback and help information should be available any time, aiming at the actual problem solving phase of the learner. Thirdly, the learner should be prevented from trapping into follow-up impasses. Thus help information should refer to the learner's pre-knowledge as much as possible.

The system to be described here is designed according to these criteria. Help information is or will be always available on demand. Planning a model is facilitated by the simplified-natural-language model editor allowing the learner to state her or his ideas in an informal way. The evaluation of models is supported qualitatively and quantitatively. Close correspondence to the learner's knowledge will be achieved by giving help that changes the learner's proposal as little as possible (minimal corrections and minimal completions).

From the beginning, computer-based support of medical reasoning had to face the problem of uncertainty of knowledge. Uncertainty was handled by heuristic approaches (for example, in MYCIN, CASNET, PIP, INTERNIST, or ABEL) as well as in a Bayesian, probability-based way (for example, in NESTOR, MUNIN, or PATHFINDER). We chose to handle uncertainty by the Bayesian network approach. A Bayesian network (e.g., [Neapolitan 1990], [Pearl 1988]) represents knowledge as a set of propositional variables and probabilistic

interrelationships between them by a directed acyclic graph. The variables are represented by the nodes of the graph, and the relations by directed arcs. The relations are conditional probabilities (each variable conditioned on its parents in the network) that define a joint probability distribution of the variables. The left of [Figure 1] shows a simple Bayesian network and the corresponding joint distribution. Independencies between variables are represented by omitting arcs, which simplifies the corresponding conditional distributions. For example, in the net on the right of [Figure 1], the variables "fever" and "sore throat" are independent given knowledge about "influenza" and "infection of throat". This means that the information "fever" is not relevant for the hypothesis "sore throat" (and vice versa) if it is already known whether the patient has influenza and a throat infection.

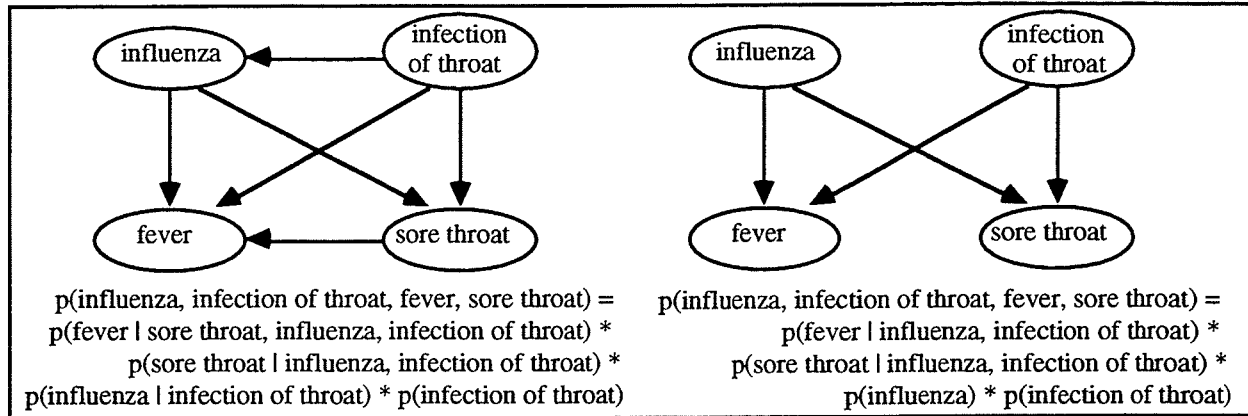


Figure 1: Simple Bayesian network without (on the left) and with (on the right) independencies

An important reason for choosing the Bayesian network approach is that it supports qualitative reasoning. A physician engaged in medical diagnosis proceeds in a highly selective manner [i.e., Elstein et al. 1978]. There is evidence that this selectivity can be explained by exploiting independencies that are also present in Bayesian networks. Our reviews of case studies in the domain of environmental medicine support this hypothesis. There is also evidence that qualitative reasoning as supported by Bayesian networks corresponds closely to human reasoning patterns [Jungermann & Thüring 1993], [Waldmann & Holyoak 1992], [Henrion 1987].

•*Alternative approaches.* The Dempster-Shafer Theory of evidential reasoning (i.e., [Gordon & Shortliffe 1990]) does not require a complete probabilistic model of the situation, so it can distinguish between uncertainty and equal certainties of events. A consequence of this is that belief and disbelief in an event do not have to sum to unity, leading to belief intervals from the positive belief in a set of possible events to their mere possibility. So the Dempster-Shafer Theory allows to represent beliefs, disbeliefs, and uncommitted beliefs, requiring two values (belief and plausibility) instead of one probability measure. This does not seem necessary at the current stage of our project. In addition, within the probability-based approach we can represent the uncertainty of probabilities by second-order probabilities [Cheeseman 1985], [Neapolitan 1990], [Pearl 1988]. So we consider the probability-based approach sufficient for our present aims, which is not meant to preclude other methods at later stages.

Fuzzy-Set Theory reasons with propositions that have vague meaning. Again there is a knowledge acquisition problem, i.e., the acquisition of many membership functions. But in our application domain, there are many vaguely specified concepts, like "severe headache" or "typical symptom" (see also the examples below). Therefore we work on integrating fuzzy concept descriptions into the Bayesian network approach by acquiring conditional distributions for fuzzy relations.

3. Modeling with MEDICUS

This section is organized in three steps: First we describe how the user may create an initial model. Then we show how MEDICUS assists the user in the qualitative revision of the model. The third subsection describes the quantitative specification of a model.

3.1 Initial Model Formulation

One of the two main goals of MEDICUS is to assist a user in developing a model of causes, effects, and other relationships in a domain of interest, where the user may be a learner in a training context or a professional. In

MEDICUS, the model is represented with a formal tool, Bayesian networks. The reason to use a formal tool is to have a precise base for reasoning and communication, and to be able to derive consequences (in-/dependencies, a posteriori distributions) which can be used for proposing recommendations, help, and modifications. At the same time it is necessary that the modeler is able to state her or his ideas in an informal way which s/he is used to. Therefore, we developed a simplified-natural-language model editor ("linguistic model editor"). After stating his model in this editor, the system can generate an initial graph automatically. Alternatively, the user may also create a graph directly in the graphical model editor.

[Figure 2] shows an example from the linguistic model editor with five sentences from the domain of environmental medicine: possible effects of benzol. Each sentence is placed in a sentence field. In order to create sentences, the modeler may select variable categories, relations, modifier, and logical junctions from a menu, and name them. The relations are classified based on i) probabilistic concepts of causality [Salmon 1984], [Suppes 1970] organized according to "kind of influence" (positive / negative) and "direction of influence" (forward, backward, or undirected), and ii) has-part / is-a hierarchies. Table 1 shows this taxonomy. Relations currently available in the linguistic model editor are marked by asterisks, but the modeler may introduce his or her own relations by specifying their kind and direction of influence. The sentences created by the modeler are checked by a definite clause grammar. Besides syntactical correctness, semantic restrictions are checked. The modeler receives feedback if the grammar finds errors.

If the modeler asks the system to create a graph representation for the model specified in the linguistic model editor, a graph is created in the graphical model editor [Figure 3]. The graph is an initial heuristic proposal which may have to be refined by the modeler qualitatively and quantitatively (see below). In creating the graph, nouns, that is, variable categories named by the user, are represented by nodes (propositional variables). Table 2 shows the propositions assigned to the variable categories. (If not specified otherwise by the user, the system creates binary variables by default.) The relations between nouns are represented by links as depicted in the rightmost column of Table 1. For relations describing undirected relations (like "corresponds to"), a dialog is evoked where the learner is asked to specify the direction, or to specify another variable as the common cause or effect of the corresponding variables.

When the graph is created, natural-language expressions of the conditional distributions for each node (resp. apriori distributions in case of root nodes) are created. They can be inspected by the user. Furthermore, the user may ask for an explanation of the relationship between the sentences in the model editor, and the graph. The explanation is based on the taxonomy of relations shown in Table 1. For example, the direction of a link is explained by the direction of influence of the verbal relation represented by that link.

3.2 Qualitative Model Revision

After the initial formulation of the model, it has to be analyzed and possibly revised on a qualitative level. In particular, it has to be verified that the dependencies and independencies implied by the graph correspond to the intentions of the modeler. As shown in [Figure 1], in Bayesian networks independencies are expressed by missing links. For example, the graph in [Figure 3] states that anaemia and leukopenia are independent, given benzol (that is, $p(\text{anaemia} | \text{benzol}) = p(\text{anaemia} | \text{benzol}, \text{leukopenia})$). This means that knowledge that a patient suffers from leukopenia is not relevant for the hypothesis that he suffers from anaemia, if it is known whether the patient is exposed to benzol. In contrast, if nothing is known about benzol, information about leukopenia is useful for the hypothesis "anaemia" ($p(\text{anaemia}) \neq p(\text{anaemia} | \text{leukopenia})$).

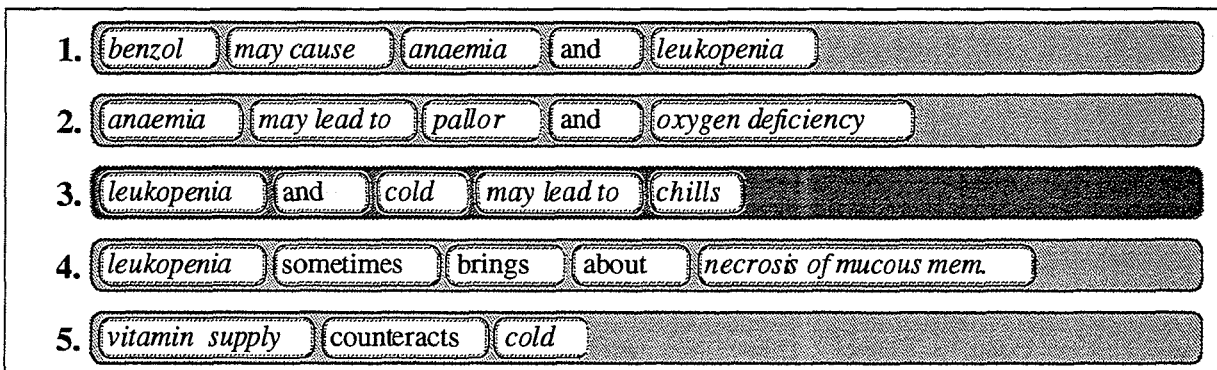


Figure 2: Five sentences created in the linguistic model editor

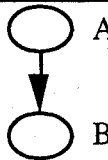
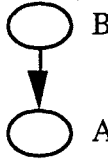
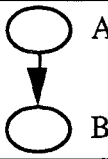
	Directed and undirected relations: kind of influence		Representation in the graph:
	positive	negative	
forward $t_A \leq t_B$	A causes B* A brings about B* A triggers B* A may lead to B ... $p(B A) > p(B)$	A counteracts B* A prevents B A suppresses B ... $p(B A) < p(B)$	 A B
backward $t_A \geq t_B$	A follows B A is consequence of B ... $p(A B) > p(A)$	A does not follow B A is suppressed by B ... $p(A B) < p(A)$	 B A
undirected	A corresponds to B* A occurs with B*	A and B are mutually exclusive	Dialog (see below)
	Is-a and Part-of hierarchies:		Representation in the graph:
	A is example for B* A contains B* $p(B A) = 1$	A is exemplified by B A is part of B $p(B \neg A) = 0$	 A B

Table 1: Taxonomy of relations used in the linguistic model editor

Variable category	Proposition
<person>	Person is <Person>
<state>	Person is in the state <state>
<event>	Person experiences the event <event>
<action>	Person performs the action <action>
<object>	Person has to do with the object <object>
<substance>	Person has to do with the substance <substance>

Table 2: Propositions assigned to variable categories

Similarly, "necrosis of mucous membrane" and "chills" are independent, given knowledge about leukopenia, but "leukopenia" and "cold" are dependent given "chills": If it is known that a patient suffers from chills, then new evidence that weakens the hypothesis "cold" will strengthen the hypothesis "leukopenia", and vice versa: Weakening one explanation for "chills" strengthens the other one. For example, if we learn that a patient suffering from chills has leukopenia, then we have an explanation for the chills, making the alternative explanation ("cold") less likely. Formally, conditional independence is described by the d-separation criterion [Pearl 1988].

In MEDICUS, we want the knowledge of the modeler to be acquired in a way that is at the same time comfortable to the modeler and informative for generating independence assertions. Therefore, a knowledge acquisition facility is currently developed that can be used for model construction and for model validation, that is, for verifying or rejecting the independencies inherent in the graph. The system offers a diagnostic dialog that proceeds in three steps:

1. For a case, the modeler specifies the initial data and symptoms, i.e., results of history taking (left window in [Figure 4]: for example, "benzol", "pallor", and "cold"). Next, he specifies a hypothesis (middle window in [Figure 4], for example "anaemia"). Thirdly, he specifies what information he considers relevant for his hypothesis, that is, what information he would look for next (right window in [Figure 4]: "oxygen deficiency" and "vitamin supply" in this case). Independency assertions are constructed from this dialog in the following way: Information *not* considered relevant to the hypothesis by the modeler, given the initial data and symptoms,

is independent of the hypothesis, because it is not considered informative for the hypothesis by the modeler. In [Figure 4], "chills" was *not* selected in the right window, so "chills" and "anaemia" are considered independent, given "benzol", "pallor", and "cold": $p(\text{anaemia} \mid \text{benzol}, \text{pallor}, \text{cold}, \text{chills}) = p(\text{anaemia} \mid \text{benzol}, \text{pallor}, \text{cold})$. Similarly, "leukopenia" and "necrosis of mucous membrane" were also not selected in the right window of [Figure 4], so $p(\text{anaemia} \mid \text{benzol}, \text{pallor}, \text{cold}, \text{leukopenia}) = p(\text{anaemia} \mid \text{benzol}, \text{pallor}, \text{cold})$, and $p(\text{anaemia} \mid \text{benzol}, \text{pallor}, \text{cold}, \text{necrosis of mucous membrane}) = p(\text{anaemia} \mid \text{benzol}, \text{pallor}, \text{cold})$.

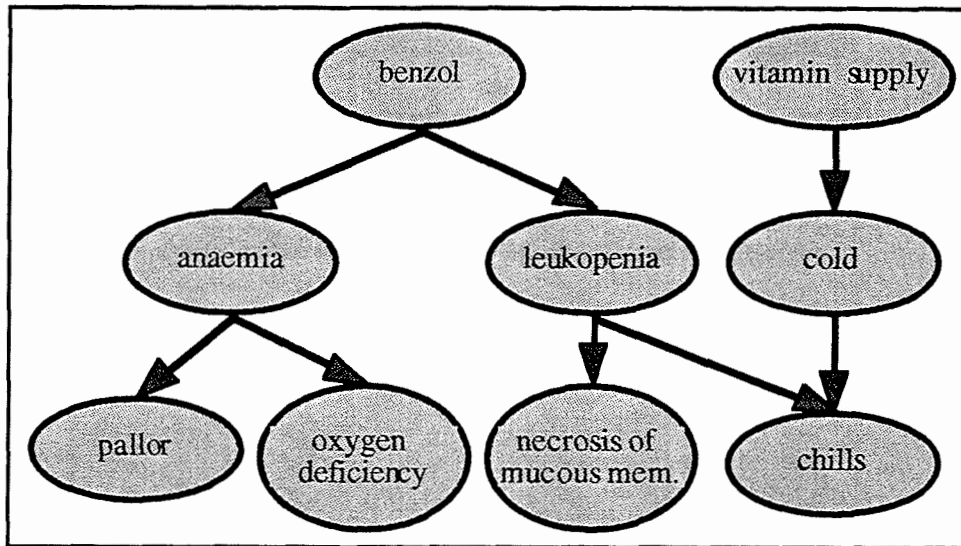


Figure 3: Graph representation in the graphical model editor generated for the sentences of [Figure 2]

Sentences: Needed information #0		
Patient data, Symptoms	Hypothesis	Additionally needed information
benzol	benzol	pallor
anaemia	anaemia	oxygen deficiency
leukopenia	leukopenia	cold
pallor	pallor	chills
oxygen deficiency	oxygen deficiency	necrosis of mucous mem.
cold	cold	vitamin supply

Figure 4: Diagnostic dialog for the acquisition of information about independencies

2. The modeler states the hypothesis that the graph is consistent with the information specified by her or him in the diagnostic dialog. The system analyzes this hypothesis using the d-separation criterion. If differences are found, a graph is constructed internally [Srinivas et al. 1990] from the dependence and independence assertions acquired in the diagnostic dialog. This internal graph is compared to the modeler's graph. This may lead to one of the following results: i) The modeler's graph and the in-/ dependencies acquired in the dialog are consistent, ii) links have to be removed from the graph in order to be consistent with the in-/ dependencies, iii) links have to be added to the graph, iv) links have to be removed from and added to the graph as well. After the dialog of [Figure 4] has taken place, the feedback for the graph in [Figure 3] is that a link has to be added.

3. On further request, the modeler may ask the system for modification proposals and an explanation of these proposals. For example, after the dialog of [Figure 4] has occurred, the system proposes for [Figure 3] to add a link between "anaemia" and "vitamin supply" because the modeler specified that "vitamin supply" is informative for "anaemia" given "benzol", "pallor", and "cold". The direction of the to-be-added link is not specified in the system's proposal because both directions are compatible with the information specified in [Figure 4]. The explanation presented to the modeler on request is shown in [Figure 5].

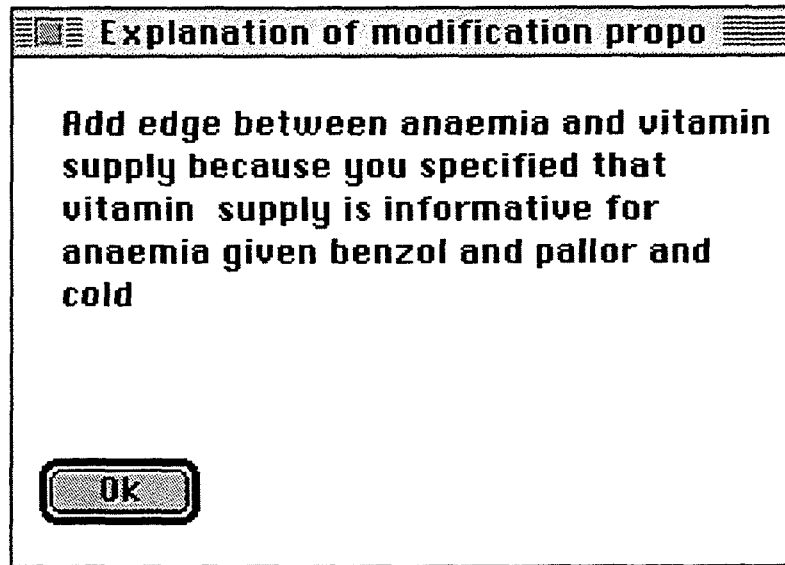


Figure 5: Explanation of the modification proposal for the graph in [Figure 3] after the dialog shown in [Figure 4] has occurred

3.3 Quantitative Model Specification

When the qualitative structure of the model is fixed, the modeler may quantify the net with apriori and conditional probabilities, enter evidences, and let the system generate posterior distributions. Like in ERGO and HUGIN, evidence propagation is implemented according to the algorithm of [Lauritzen & Spiegelhalter 1988]. As mentioned at the end of [Section 2], we want MEDICUS to generate the needed conditional probabilities from the verbal relations specified by the modeler in the linguistic model editor. This is part of our current work.

4. Supporting Diagnostic Reasoning

For a specified model, MEDICUS generates qualitative diagnostic recommendations in a preliminary way (without a utility model, [Heckerman et al. 1992]). For symptoms given, MEDICUS lists the currently most probable syndrome hypotheses, and it recommends diagnostically relevant symptoms and environmental factors to consider next. These recommendations have been demonstrated with a more realistic, multiply connected net containing about fifty variables to a community of environmental medicinal professionals.

5. Conclusions and Further Work

Together with our cooperation partners, diagnostic support will be applied to problems of planning and interpreting clinical and environmental investigations. Currently we create practically useable applications for problems of environmental monitoring and human genetics. With the Medical Institute of Environmental Hygiene, Düsseldorf, it is planned to apply our system to a large set of case data from environmental medicine. In this way it will be possible to construct a large and realistic network suitable for serious diagnostic training. One of our long-term goals is to establish MEDICUS within university and postqualification courses. So one of the next research goals is to give more detailed support and explanations for diagnostic reasoning. Another research goal is to enable collaborative or competitive modeling of several agents. Thus modeling will become a

group activity. The goal of this application will be to help structure cases and research results in environmental medicine to achieve a unified model, or if this is not possible, to pinpoint differences and contradictions.

6. References

- [Anderson 1993] Anderson, J.R. (1993). *Rules of Mind*. Hillsdale: Erlbaum.
- [Andreassen et al. 1987] Andreassen, S., Woldbye, M., Falck, B., & Andersen, S.K. (1987). MUNIN - A Causal Probabilistic Network for Interpretation of Electromyographic Findings. *Proceedings 10th IJCAI 87*, 366-372.
- [Barahona et al. 1995] Barahona, P., Stefanelli, M., & Wyatt, J. (eds) (1995). *Artificial Intelligence in Medicine*. Proc. 5th Conf. AIME 95. Berlin: Springer (LNAI 934).
- [Barrows & Tamblyn 1980] Barrows, H.S., & Tamblyn, R.M. (1980). *Problem-Based Learning: An Approach to Medical Education*. Springer.
- [Boshuizen & Schmidt 1992] Boshuizen, H.P.A., & Schmidt, H.G. (1992). On the Role of Biochemical Knowledge in Clinical Reasoning by Experts, Intermediates, and Novices. *Cognitive Science*, 16, 153-184.
- [Cheeseman 1985] Cheeseman, P. (1985). In Defense of Probability. *Proc. 9th IJCAI 85*, Los Angeles, 1002-1009.
- [Cooper 1984] Cooper, G.F. (1984). NESTOR: A Computer-Based Medical Diagnostic Aid that Integrates Causal and Probabilistic Knowledge. PhD Thesis, Medical Computer Science Group, Stanford Univ., CA (Report HPP-84-48).
- [Elstein & Bordage 1980] Elstein, A.S., & Bordage, G. (1980). Psychology of Clinical Reasoning., in G.C. Stone, F. Cohen, N.E. Adler, and Assoc. (eds), *Health Psychology - A Handbook*. San Francisco: Jossey-Bass, 333-367.
- [Elstein et al. 1978] Elstein, A.L., Shulman, L.S., & Sprafka, S.A. (1978). *Medical Problem Solving - An Analysis of Clinical Reasoning*. Cambridge: Harvard University Press.
- [Gollwitzer 1990] Gollwitzer, P.M. (1990). Action Phases and Mind-Sets. In E.T. Higgins & R.M. Sorrentino (eds): *Handbook of Motivation and Cognition: Foundations of Social Behavior*, Vol.2. 53-92.
- [Gordon & Shortliffe 1990] Gordon, J., & Shortliffe, E.H. (1990). The Dempster-Shafer Theory of Evidence. In G. Shafer & J. Pearl (eds), *Readings in Uncertain Reasoning*. Morgan Kaufman.
- [Heckerman 1991] Heckerman, D.E. (1991). *Probabilistic Similarity Networks*. Cambridge: MIT Press.
- [Heckerman et al. 1992] Heckerman, D.E., Horvitz, E.J., & Nathwani, B.N. (1992). Toward Normative Expert Systems: Part I - The Pathfinder Project. *Methods of Informatics in Medicine*, 31, 90-105.
- [Henrion 1987] Henrion, M. (1987). Uncertainty in Artificial Intelligence: Is Probability Epistemologically and Heuristically Adequate? In J.L. Mumpower, L.D. Philipps, O. Renn, V.R.R. Uppuluri (eds), *Expert Judgements and Expert Systems*. Berlin: Springer (NATO ASI Series F: Computer and Systems Science), 106-129.
- [Jungermann & Thüring 1993] Jungermann, H., & Thüring, M. (1993). Causal Knowledge and the Expression of Uncertainty. In G. Strube, K.F. Wender (eds), *The Cognitive Psychology of Knowledge*. Elsevier, 53-73.
- [Lauritzen & Spiegelhalter 1988] Lauritzen, S.L., & Spiegelhalter, D.J. (1988). Local Computations with Probabilities on Graphical Structures and their Application to Expert Systems. *J. of the Royal Statistical Soc.*, B50(2), 157-224.
- [Miller et al. 1982] Miller, R.A., Pople, H.E., & Myers, J.D. (1982). INTERNIST-I, an Experimental Computer-Based Diagnostic Consultant for General Internal Medicine. *The New England Journal of Medicine*, 307(8), 468-476.
- [Möbus 1995] Möbus, C. (1995). Towards an Epistemology of Intelligent Problem Solving Environments: The Hypothesis Testing Approach. In J. Greer (ed), *Proc. Art. Intelligence and Education AI-ED 95, AACE*, 138-145.
- [Neapolitan 1990] Neapolitan, R.E. (1990). *Probabilistic Reasoning in Expert Systems*. New York: Wiley.
- [Newell 1990] Newell, A. (1990). *Unified theories of cognition*. Cambridge: Harvard University Press.
- [Patel & Groen 1986] Patel, V.L., & Groen, G.J. (1986). Knowledge Based Solution Strategies in Medical Reasoning. *Cognitive Science*, 10, 91-116.
- [Patil et al. 1981] Patil, R.S., Szolovits, P., & Schwarz, W.B. (1981). Causal Understanding of Patient Illness in Medical Diagnosis. *Proc. 7th Int. Joint Conf. on Artificial Intelligence IJCAI 81*, Los Altos: Kaufman.
- [Pearl 1988] Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufman (2nd printing).
- [Salmon 1984] Salmon, W.C. (1984). *Scientific Explanation and the Causal Structure of the World*. Princeton Univ.
- [Shortliffe 1976] Shortliffe, E.H. (1976). *Computer-Based Medical Consultations: MYCIN*. New York: North-Holland.
- [Srinivas et al. 1990] Srinivas, S., Russell, S., & Agogino, A. (1990). Automated Construction of Sparse Bayesian Networks from Unstructured Probabilistic Models and Domain Information. In M. Henrion, R.D. Schachter, L.N. Kanal, J.F. Lemmer (eds), *Uncertainty in Artificial Intelligence 5*, Amsterdam: North Holland, 295-307.
- [Suppes 1970] Suppes, P. (1970). *A Probabilistic Theory of Causality*. Amsterdam: North-Holland.
- [Szolovits & Pauker 1978; 1993] Szolovits, P., & Pauker, S.G. (1978). Categorical and Probabilistic Reasoning in Medical Diagnosis. *Artificial Intelligence*, 11, 115-144, and (1993) *Artificial Intelligence*, 59, 167-180.
- [van Lehn 1988] van Lehn, K. (1988). Toward a Theory of Impasse-Driven Learning. In: Mandl, H.; Lesgold, A. (eds): *Learning Issues for Intelligent Tutoring Systems*. New York: Springer, 19-41.
- [Waldmann & Holyoak 1992] Waldmann, M.R., & Holyoak, K.J. (1992). Predictive and Diagnostic Learning Within Causal Models: Asymmetries of Cue Competition. *Journal of Experimental Psychology: General*, 121, 222-236.
- [Weiss et al. 1978] Weiss, S.M., Kulikowski, C.A., & Amarel, S. (1978). A Model-Based Method for Computer-Aided Medical Decision Making. *Artificial Intelligence*, 11, 145-172.

Acknowledgements

We thank Karsten Rommerskirchen for assisting in the implementation and in the mathematical work.