Acquiring Knowledge from Linguistic Models in Complex, Probabilistic Domains

Olaf Schröder, Claus Möbus, Heinz-Jürgen Thole[1]
OFFIS Institute, Escherweg 2, D - 26121 Oldenburg, Germany
Telephone: + 49 441 798 2900, Fax: + 49 441 798 2155
E-Mail: {schroeder, moebus, thole}@informatik.uni-oldenburg.de

Working session: 12. Natural Language

Abstract. This paper describes an approach to acquire qualitative and quantitative knowledge from verbally stated models in complex, probabilistic domains. This work is part of the development of an intelligent environment, MEDICUS[2], that supports modelling and diagnostic reasoning in the domains of environmental medicine and human genetics. These domains are two yet new subdomains of medicine receiving increasing research efforts, but still consisting of largely fragile and uncertain knowledge. In MEDICUS, uncertainty is handled by the Bayesian network approach. Thus the *modelling task* for the user consists of creating a Bayesian network for the problem at hand. But since we want mathematically untrained persons to work with MEDICUS, the user may alternatively state propositions verbally and let the system generate a Bayesian network proposal. 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. The *diagnostic reasoning task* for the learner consists of using the network for stating diagnostic goals, and for proposing diagnostic hypotheses and examinations.

---

[1] We thank Jörg Folckers and Karsten Rommerskirchen for doing large parts of the implementation and for assisting in the mathematical work.

[2] Modelling, explanation, and diagnostic support for complex, uncertain subject matters

In this paper, we first give an overview of the aims and the actual implementation state of MEDICUS. Then we will focus on the modelling component and the central concern of this paper: One of the most difficult problems in the design of a domain model represented as a Bayesian network is to acquire the necessary qualitative and quantitative information from the modeller (of course this is a problem for other uncertainty formalisms too.):

• With respect to qualitative information, the dependence and independence relations implied by a Bayesian network have to be validated empirically. This can be achieved by deriving assertions from these relations and comparing them to the modeller's assertions. For example, the dependence and independence relations can be compared to descriptions of diagnostic procedures: Given a diagnostic hypothesis and some case information, what diagnostic information is considered next? What diagnostic information is *not* considered?

• With respect to quantitative information, apriori and conditional distributions have to be obtained in order to be able to use the network for diagnostic reasoning. But even domain experts are usually hesitant to specify numerical relationships. The possibility to work with intervals, as for example offered by Dempster-Shafer Theory, does not solve this problem. Rather, the modeller should be able to state propositions verbally. The system should be able to assign probabilities to these "fuzzy" relations and concepts. Existing approaches have been concerned with the empirical acquisition of the semantics of fuzzy terms concerning a single variable or proposition (for example: "It is *possible* that it will rain tomorrow."), but not with fuzzy *relations*. We attempt to close this gap by extending existing approaches to relations.

Keywords: Knowledge acquisition, hypotheses testing, reasoning with uncertainty, explanation, Bayesian networks

1. Introduction

Diagnosis and model building are reasoning and problem solving tasks that can be quite difficult. This is especially true in medical domains (Barrows & Tamblyn, 1980; Boshuizen & Schmidt, 1992; Elstein et al., 1978; Elstein & Bordage, 1980; Patel & Groen, 1986) where the knowledge is particularly complex, interrelated, fragile, and uncertain. Two examples of such domains are the epidemiology of diseases caused by environmental influences, like pollution, and of diseases caused by human genetic defects. In these domains, clear-cut taxonomies and explanatory models of diseases, or syndromes, have not been developed yet. Still these domains are getting increasingly important. This is for example reflected by the fact that currently many postqualification courses for physicians are established.

From the beginning of computer-based support of medical reasoning, the problem of uncertainty received central attention. It was handled by heuristic approaches (i.e. MYCIN, Shortliffe, 1976; CASNET, Weiss et al., 1978; PIP, Szolovits & Pauker, 1978; 1993; INTERNIST, Miller et al., 1982) as well as in a normative probability-based way (i.e., NESTOR, Cooper, 1984; MUNIN, Andreassen et al., 1987; PATHFINDER, Heckerman, 1991). The main aim of these systems is to provide diagnostic hypotheses, given the available evidence, and to suggest further diagnostic steps, for example, for differential diagnosis. Some systems, like CASNET, also generate therapeutic recommendations. But in spite of some capability to explain their reasoning steps, much of the reasoning and knowledge structures of these systems remains hidden to the user.

For the purpose of medical training, the recommendation of diagnostic hypotheses and investigations is important but not sufficient. Firstly, since diagnostic reasoning is a problem solving task, there is a need for students of medicine to *train* it (Barrows & Tamblyn, 1980). Thus learner should have an opportunity to *actively perform* diagnostic reasoning and to apply diagnostic strategies. Secondly, a prerequisite for diagnoses is a sound knowledge base about the phenomena in question. 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 collaboration with medical institutions (Health Authority of Oldenburg, Documentation and Information Center for Environmental Issues, Osnabrück, and Medical Institute for Environmental Hygiene, Düsseldorf, Robert-Koch-Institute, Berlin), we develop an intelligent modelling and diagnosis environment, MEDICUS, that differs from existing medical expert systems by being designed to support these two activities:

- *Model construction*. The modeller may create an initial model of the domain at interest and further specify, evaluate, and revise the model qualitatively and quantitatively.

• *Training of diagnostic strategies.* The learner may state diagnostic hypotheses and obtain diagnostic recommendations from the system 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 the diagnostic hypotheses?)

This paper focuses on the modelling part of MEDICUS. For the creation of a Baysian network domain model, the modeller has to specify a lot of qualitative and quantitative information. On the qualitative level, the network implies dependence and independence relations, and it has to be verified that these implications are consistent with the domain model of the user. On the quantitative level, apriori and conditional distributions are needed for the network.

But even domain experts are usually hesitant to specify numerical relationships. In addition, we want MEDICUS to be usable also for mathematically untrained users. Therefore, MEDICUS has a *linguistic model editor* where the user can state her or his model assumptions in a simplified natural language. The system creates an initial Baysian network from the information stated in the linguistic model editor. In order to close the gap between the linguistically stated model and the additional qualitative and quantitative information needed for the Bayesian network, we need knowledge acquisition techniques that provide this information.

The next section provides a short overview of MEDICUS. (A more detailed description can be found in Folckers et al., 1996). After that, we will describe how qualitative and quantitative information will be acquired. The closing section will state some conclusions.

2. An Overview of MEDICUS

2.1 Design Principles

As stated, 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. We call our approach an *Intelligent Problem Solving Environment* (IPSE, Möbus, 1995): The learner acquires knowledge by actively *testing hypotheses*. This means that the learner creates solution proposals or models, tests hypotheses about their correctness, and the system analyses the proposals 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; 1991), Newell (1990), Anderson (1989; 1993), Gollwitzer (1990), and Heckhausen (1989). Briefly, it states that new knowledge is acquired as a result of problem solving and applying weak heuristics in response to impasses. In contrast, knowledge is optimized if applied successfully. In addition, 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.

MEDICUS is designed according to these criteria. For example, help information is or will be always available on demand. Planning a model is facilitated by the simplified-natural-language model editor which allows the learner to state her or his ideas in an informal way. The evaluation of models is supported qualitatively and quantitatively.

We chose to handle the uncertainty of knowledge by the Bayesian network approach. A Bayesian network (e.g., Neapolitan, 1990; Pearl, 1991) 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 "feaver" and "sore throat" are independent given knowledge about "influenza" and "infection of throat". This means that the information "feaver" 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.

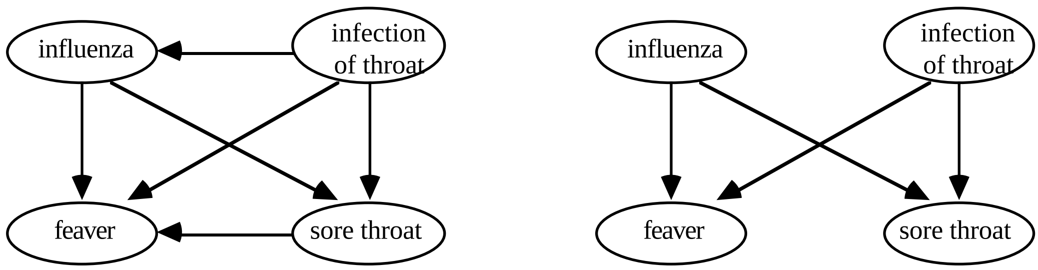

p(influenza, infection of throat, feaver, sore throat) =
p(feaver | sore throat, influenza, infection of throat) *
p(sore throat | influenza, infection of throat) *
p(influenza | infection of throat) * p(infection of throat)

p(influenza, infection of throat, feaver, sore throat) =
p(feaver | influenza, infection of throat) *
p(sore throat | influenza, infection of throat) *
p(influenza) * p(infection of throat)

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, and that qualitative reasoning as supported by Bayesian networks corresponds closely to human reasoning patterns (Henrion, 1987; Jungermann & Thüring, 1993; Waldmann & Holyoak, 1992; Waldmann et al., 1995). For a short discussion of alternative approaches (Dempster-Shafer Theory, Fuzzy Set Theory) for MEDICUS see Folckers et al. (1996).

2.2. The Implementation State

2.2.1 Supporting Model Construction

One of the main goals of MEDICUS is to assist the learner in developing a model of perceived causes, effects, and other relationships in a domain of interest 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, aposteriori distributions) which can be used for proposing recommendations, help, and modifications. At the same time it is necessary that the learner is able to state his ideas in an informal way which he is used to. Therefore, we developed a simplified-natural-language *linguistic model editor*. After stating his model in this editor, the system can generate an initial graph automatically.

Figure 2 shows an example with four sentences. Each sentence is placed in a sentence field. In order to create sentences, the learner 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 (Suppes, 1970) organised 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 model editor are marked by asterisks. They were selected as a result of discussions of a topic of environmental medicine with an expert, but we plan to extend this list. The sentences created by the learner are checked by a definite clause grammar. Besides syntactical correctness, semantic restrictions are checked. The learner receives feedback if the grammar finds errors.

The learner may ask the system to create a graph representation for the model specified (Figure 3). In creating the graph, nouns (= variable categories named by the user) are represented by nodes (propositional variables). Table 2 shows the propositions assigned to the variable categories. 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 variables in question. The learner may also ask for an explanation of the relationship between the sentences in the model editor, and the net. The explanation is based on the taxonomy of relation shown in Table 1.

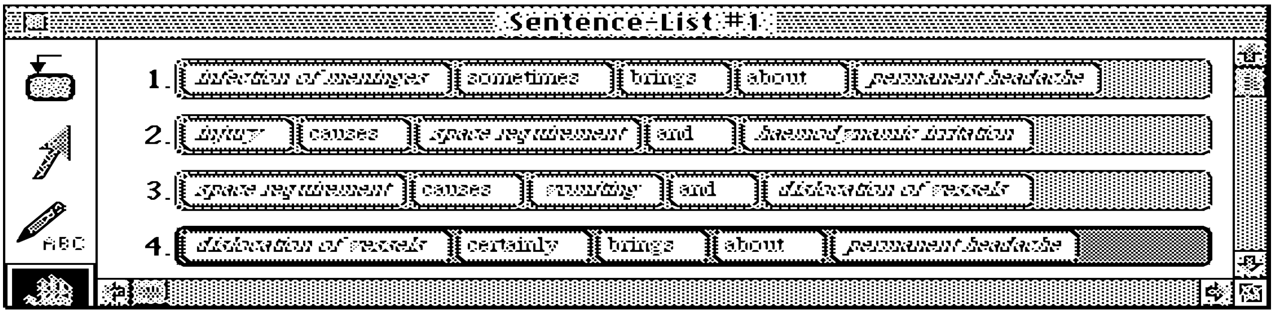

Figure 2: Four sentences created in the linguistic model editor

direction of influence		Directed and undirected relations: kind of influence — positive	negative	Representation in the graph:
	forward $tA \le tB$	A causes B* A brings about B* A triggers B* $p(B \mid A) > p(B)$	A counteracts B* A prevents B $p(B \mid A) < p(B)$	A B
	backward $tA \ge tB$	A follows B A is consequence of B ... $p(A \mid B) > p(A)$	A does not follow B ... $p(A \mid B) < p(A)$	B A
	undirected	A corresponds to B* A occurs with B*	A and B are mutually exclusive	Dialog
		Is-a and Part-of hierarchies:		Representation in the graph:
		A is example for B* A contains B* $p(B \mid A) = 1$ A is exemplified by B A is part of B $p(B \mid \neg A) = 0$		A B

Table 1: Taxonomy of relations used in the 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

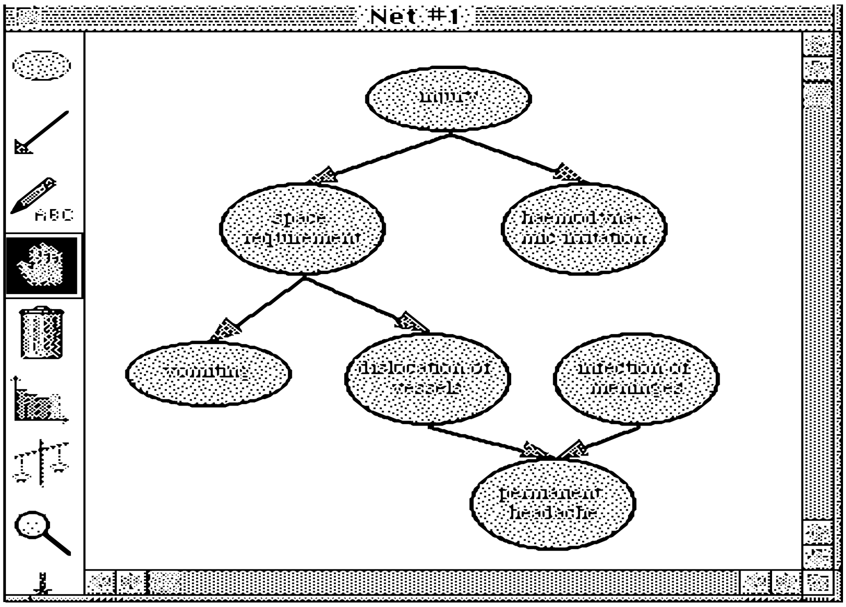

Figure 3: Graph representation generated for the sentences of Figure 2

The graph is an initial heuristic proposal which has to be refined by the learner qualitatively and quantitatively (see below). When the qualitative structure of the model is fixed, the modeller may quantify the net with apriori and conditional probabilities, enter evidences, and let the system generate posterior distributions. Like for example in ERGO and HUGIN, evidence propagation is implemented according to the Lauritzen & Spiegelhalter (1988) algorithm.

2.2.2 Training Diagnostic Reasoning

The current implementation state of MEDICUS generates qualitative recommendations. It lists the syndrome hypotheses most probable in the light of actual anamnestic evidence, and it recommends what 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.

Together with our cooperation partners, diagnostic support will be applied to problems of planning and interpreting clinical and environmental investigations ("Environmental Monitoring"). 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.

3 Acquiring Qualitative and Quantitative Knowledge for the Graph

3.1 Acquring Qualitative Knowledge by a Dialog

After the initial formulation of the model, it has to be analysed and revised on a qualitative level. In particular, it has to be verified that the dependencies and independencis implied by the graph correspond to the assertions stated by the modeller. As shown in Figure 1, in Bayesian networks independencies are expressed by missing links. For example, the graph in Figure 3 states that space requirement and haemodynamic irritation are independent, given injury (that is, p(space requirement | injury) = p(space requirement | injury, haemodynamic irritation)). This means that knowledge about a haemodynamic irritation is not relevant for the hypothesis "space requirement", if it is known that an injury took place or not. In contrast, if nothing is known about injury, information about a haemodynamic irritation is useful for the hypothesis "space requirement" (p(space requirement) ≠ p(space requirement | haemodynamic irritation)). Similarly, vomiting and dislocation of vessels are independent, given space requirement, but dislocation of vessels and infection of meninges are dependent, given permanent headache: If it is known that a patient suffers from permanent headache, then new evidence that weakens the hypothesis "dislocation of vessels" will strengthen the hypothesis "infection of meninges", and vice versa: Weakening one explanation for "permanent headache" strengthens the other one. Formally, conditional independence is described by the d-separation criterion (Pearl, 1991).

The knowledge of the modeller has to be acquired by the system in a way that is at the same time comfortable to the modeller and informative for generating independence assertions. Therefore, a knowledge acquisition facility is currently developed that can be used for model *construction* or 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 modeller specifies the initial anamnestic data and symptoms (left window in Figure 4: for example, "injury" and "dislocation of vessels"). Next, he specifies a diagnostic hypothesis (middle window in Figure 4, for example "space requirement"). Thirdly, he specifies what information he would look for next (right window in Figure 4: "haemodynamic irritation" and "vomiting" in this case). Independencies are constructed from this dialog in the following way: Information not considered relevant to the hypothesis by the modeller, given the anamnestic data and symptoms, is independent of the hypothesis. In Figure 4, "permanent headache" was *not* selected in the right window, so "permanent headache" and "space requirement" are considered independent, given "injury" and "dislocation of vessels": p(space requirement | injury, dislocation of vessels, permanent headache) = p(space requirement | injury, dislocation of vessels).

2. The modeller states the hypothesis that the graph is consistent with the information specified by her or him in the diagnostic dialog. The system analyses 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 modeller's graph. This may lead to the result that a) the graph and the in-/ dependencies are consistent, that b) edges have to be removed from the graph in

order to be consistent with the in-/ dependencies, that c) edges have to be added to the graph (this is the feedback given for the net in Figure 3 after the dialog of Figure 4 has taken place), or that d) edges have to be removed from and added to the graph as well.

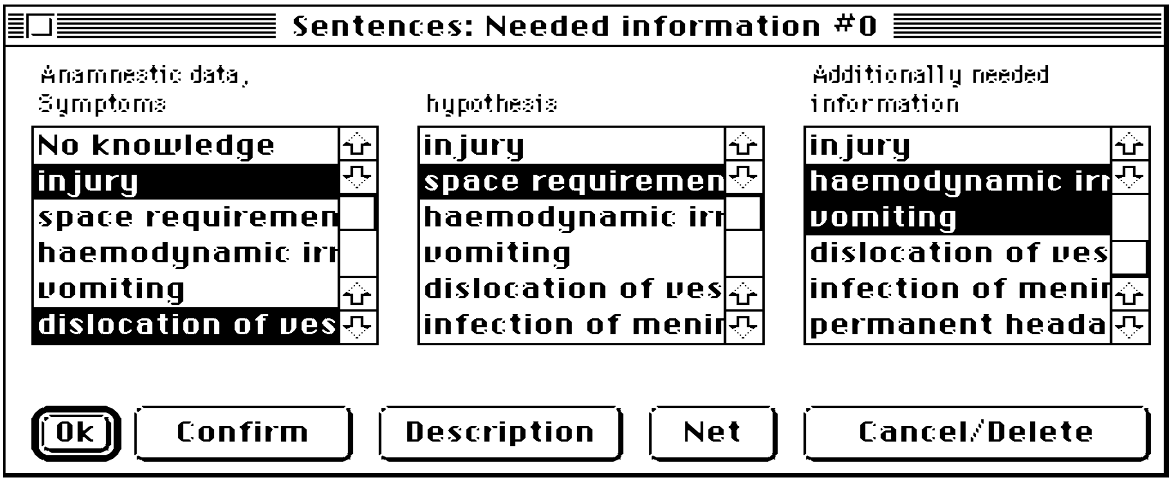

Figure 4: Diagnostic dialog for the acquisition of information about independencies

3. On further request, the modeller 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 an edge from "space requirement" to "haemodynamic irritation" because the modeller specified that "haemodynamic irritation" is informative for "space requirement" given "injury" and "dislocation of vessels".

3.2 Acquring Quantitative Knowledge from Fuzzy Relations

When the qualitative structure of the model is fixed, the modeller may quantify the net with apriori and conditional probabilities, enter evidences, and let the system generate posterior distributions. But as stated, even experts hesitate to specify numbers, so it would be appropriate to make use of the information contained in the "fuzzy" relations of the propositions stated in the linguistic model editor (like for example "sometimes brings about", see Figure 2). There is literature about the empirical investigation of the semantics of adverb phrases like "probably", "perhaps", "maybe", etc., and modal verb forms like "should", "will", "may", etc. (Kipper & Jameson, 1994; Teigen, 1988; Teigen & Brun, 1995; Wallsten et al., 1986; Budescu, Weinberg, & Wallsten, 1988; Wallsten & Budescu, 1995; Zimmer, 1983). Some of these studies try to assign membership functions to these linguistic forms by presenting different "wheel of fortune" configurations to subjects. For a list of linguistic terms, the subjects then have to indicate how well each term describes the "wheel of fortune" configuration presented. For example, a wheel of fortune with a winning area of 20% and a loosing area of 80% is better described by the statement "It is possible that I will win" than by the statement "It is very likely that I will win." We plan to extend this approach to relations. This means the following:

1. Application of the "wheel of fortune paradigm" to conditional events. This can be achieved by a wheel of fortune configuration as depicted in Figure 5: If spinning wheel A leads to the event "A+", then the wheel "B after A+" is spun, otherwise the wheel "B after A-" is spun. (For simplicity, we only consider binary variables here, although MEDICUS handles multivalued variables as well.)

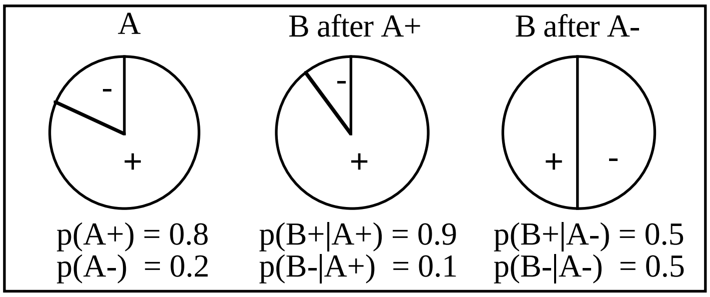

Figure 5: Wheel of fortune for relations between two binary variables A and B

2. Membership functions of these wheel-of-fortune configurations to a linguistic characterisation like for example "B is a typical consequence of A". These membership functions have to be obtained empirically. According to probabilistic concepts of causality (Suppes, 1970; Anderson, 1990), a positive influence of variable A on variable B can be expressed by p(B+ | A+) > p(B+ | A-), and a negative influence as p(B- | A+) > p(B- | A-). Thus it seems reasonable to hypothesise that fuzzy relations express a certain relationship between p(B+ | A+) and p(B+ | A-), or between p(B- | A+) and p(B- | A-), respectively. This relation can be expressed by the likelihood p(B+ | A+) / p(B+ | A-). Thus we will investigate whether the membership values can be expressed a function of the likelihood. Figure 6 illustrates this by presenting some fictituous membership functions for verbal phrases common in the literature of environmental medicine (Beyer & Eis, 1994, translated from German by the authors).

Thus the likelihood ratio that can be derived from *two* wheels of fortune ("B after A+" and "B after A-") is a bivariate analogue to the univariate situation with *one* wheel of fortune pursued in the studies mentioned above.

In order to interpret membership functions as probabilities, the following assumption has to be made (Kipper, 1995): The probability that a certain matter of facts F is expressed by the verbal phrase V, p(V | F), is proportional to the membership value of V for F. Using this assumption, probabilities p(verbal phrase | likelihood) can be obtained, like for example p("B+ is considered a cue for A+" | p(B+ | A+) / p(B+ | A-) = x). In addition, an apriori distribution of the likelihood, p(p(B+ | A+) / p(B+ | A-) = x), is needed. Then the desired probabilities p(likelihood | verbal phrase) = p([p(B+ | A+) / p(B+ | A-)] | "B+ is considered a cue for A+") can be obtained. The mode of this probability distribution is the likelihood ratio that best represents the verbal phrase in question. If there is more than one verbal phrase for a certain relationship, for example verbal descriptions provided by different experts (like "B+ is considered a cue for A+", "A+ may cause B+" and so on), then the aposteriori probabilities p([p(B+ | A+) / p(B+ | A-)] | "B+ is considered a cue for A+", "A+ may cause B+", ...) can be computed.

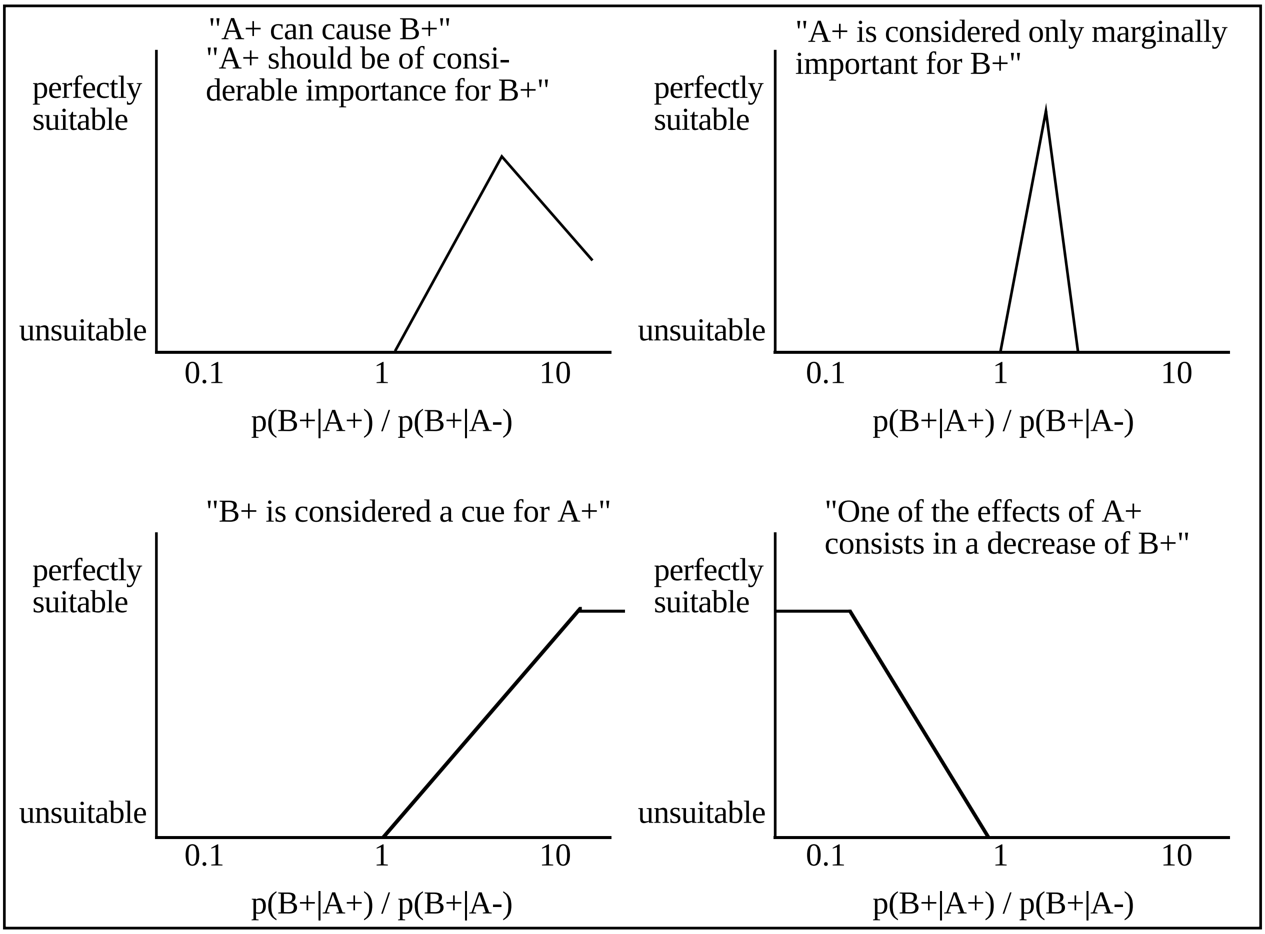

Figure 6: Fictituous membership functions for some verbal phrases found in the environmental medicinal literature

From the desired aposteriori probabilities p(likelihood | verbal phrase) the needed conditional distributions for the Bayesian network can be obtained. (For example, if p(B+ | A+) / p(B+ | A-) and p(B- | A+) / p(B- | A-) are known, p(B+ | A+), p(B+ | A-), p(B- | A+), and p(B- | A-) can be obtained.)

4. Conclusions and Further Work

One of our next steps will be a study to obtain likelihoods for verbal phrases stated in our linguistic model editor empirically. The likelihoods can then be used for determining the conditional probabilities needed for the network.

One of our long-term goals is to establish the system as a modelling and diagnostic reasoning tool within university and postqualification courses. This requires that the modelling component gets even more attractive and accepted especially for mathematically untrained persons. Therefore, with our cooperation partners we want to integrate diagnostic information sources in the modelling component as well. This means that the user will be able to state descriptions of diagnostic processes, case descriptions and histories, and heuristic representations common in medicine like decision trees, flowcharts, and so on. These information sources will be used to automatically construct a Bayesian network that can be used for diagnostic hypotheses testing by the modeller himself, and by other persons as well.

5. References

Anderson, J.R. (1989). A theory of the origins of human knowledge. *Artificial Intelligence*, 40, 313-351.

Anderson, J.R. (1990). *The Adaptive Character of Thought*, Hillsdale: Erlbaum.

Anderson, J.R. (1993). *Rules of Mind.* Hillsdale: Erlbaum.

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.

Barrows, H.S., Tamblyn, R.M. (1980). *Problem-Based Learning: An Approach to Medical Education*. Springer.

Beyer, A., Eis, D. (eds) (1994). *Praktische Umweltmedizin* [Practical Environmental Medicine], Berlin: Springer.

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.

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., Stanford, CA (Report HPP-84-48).

Elstein, A.S., Bordage, G. (1980). Psychology of Clinical Reasoning., in G.C. Stone, F. Cohen, N.E. Adler, and Associates (eds), *Health Psychology - A Handbook*. San Francisco: Jossey-Bass Publ., 333-367.

Elstein, A.L., Shulman, L.S., Sprafka, S.A. (1978). *Medical Problem Solving - An Analysis of Clinical Reasoning*, Cambridge: Harvard University Press.

Folckers, J., Möbus, C., Schröder, O., Thole, H.-J. (1996). *An Intelligent Problem Solving Environment for Designing Explanation Models and for Diagnostic Reasoning in Probabilistic Domains*, paper accepted for Intelligent Tutoring Systems ITS 96, Montreal, June 12 - 14, 1996.

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.

Heckerman, D.E. (1991). *Probabilistic Similarity Networks*, Cambridge: MIT Press.

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, H., Thüring, M. (1993). Causal Knowledge and the Expression of Uncertainty, in G. Strube, K.F. Wender (eds), *The Cognitive Psychology of Knowledge*. Amsterdam: Elsevier, 53-73.

Kipper, B. (1995). *Repräsentation und Verarbeitung propositionaler Einstellungen in natürlichsprachlichen Systemen*, University of Saarbrücken: Dissertation.

Kipper, B., Jameson, A. (1994) *Semantics and Pragmatics of Vage Probability Expressions*, University of Saarbrücken, Sonderforschungsbereich 314 Künstliche Intelligenz - Wissensbasierte Systeme, Bericht Nr. 105.

Lauritzen, S.L., Spiegelhalter, D.J. (1988). Local Computations with Probabilities on Graphical Structures and their Application to Expert Systems, *Journal of the Royal Statistical Society*, B50(2), 157-224.

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, C. (1995). Towards an Epistemology of Intelligent Problem Solving Environments: The Hypothesis Testing Approach, in J. Greer (ed), *Proc. of the World Conf. on Artificial Intelligence and Education AI-ED 95*, 138-145.

Neapolitan, R.E. (1990). *Probabilistic Reasoning in Expert Systems*, New York: Wiley.

Newell, A. (1990). *Unified theories of cognition*. Cambridge: Harward University Press.

Patel, V.L., Groen, G.J. (1986). Knowledge Based Solution Strategies in Medical Reasoning, *Cogn. Sci.*, 10, 91-116.

Pearl, J. (1991). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufman.

Shafer, G., Srivastava, R. (1990). The Bayesian and Belief-Function Formalisms: A General Perspective for Auditing, in G. Shafer, J. Pearl (eds), *Readings in Uncertain Reasoning*, San Mateo: Morgan Kaufman, 482-521.

Shortliffe, E.H. (1976). *Computer-Based Medical Consultations: MYCIN*, New York: North-Holland.

Suppes, P. (1970). *A Probabilistic Theory of Causality*, Amsterdam: North-Holland.

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.

Teigen, K.H. (1988). *The Language of Uncertainty*, Acta Psychologica, 68, 27-38.

Teigen, K.H., Brun, W.(1995). *Yes, But it is Uncertain: Direction and Communicative Intention of Verbal Probabilistic Terms*, Acta Psychologica, 88, 233-258.

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, 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.

Waldmann, M.R., Holyoak, K.J., Fratianne, A. (1995). *Causal Models and the Acquisition of Category Structure*. Journal of Experimental Psychology: General, 124, 2, 181-206.

Wallsten, T.S., Budescu, D.V., Rapoport, A., Zwick, R., Forsyth, B.(1986). *Measuring the Vague Meanings of Probability Terms*, Journal of Experimental Psychology: General, 115, 4, 348-365.

Wallsten, T.H., Budescu, D.V. (1995). *A Review of Human Linguistic Probability Processing: General Principles and Empirical Evidence*, The Knowledge Engineering Review, 10, 1, 43-62.

Weiss, S.M., Kulikowski, C.A., Amarel, S. (1978). A Model-Based Method for Computer-Aided Medical Decision Making, *Artificial Intelligence*, 11, 145-172.

Zimmer, A.C.(1983). *Verbal vs. Numerical Processing of Subjective Probabilities*, in R.W. Scholz (ed), Decision Making under Uncertainty, Elsevier.