ELSEVIER

Artificial Intelligence in Medicine 30 (2004) 1–26

Artificial Intelligence in Medicine

www.elsevier.com/locate/artmed

Knowledge acquisition in the fuzzy knowledge representation framework of a medical consultation system

Karl Boegl[a,b,*], Klaus-Peter Adlassnig[a,b], Yoichi Hayashi[c], Thomas E. Rothenfluh[d], Harald Leitich[a,b]

[a]*Section on Medical Expert and Knowledge-Based Systems, Department of Medical Computer Sciences, University of Vienna, Vienna, Austria*

[b]*Ludwig Boltzmann Institute for Expert Systems and Quality Management in Medicine, Vienna, Austria*

[c]*Department of Computer Science, Meiji University, Kawasaki, Japan*

[d]*Department of Psychology, University of Zurich, Zurich, Switzerland*

Received 19 March 2001; received in revised form 3 May 2001; accepted 11 January 2002

Abstract

This paper describes the fuzzy knowledge representation framework of the medical computer consultation system MedFrame/CADIAG-IV as well as the specific knowledge acquisition techniques that have been developed to support the definition of knowledge concepts and inference rules. As in its predecessor system CADIAG-II, fuzzy medical knowledge bases are used to model the uncertainty and the vagueness of medical concepts and fuzzy logic reasoning mechanisms provide the basic inference processes. The elicitation and acquisition of medical knowledge from domain experts has often been described as the most difficult and time-consuming task in knowledge-based system development in medicine. It comes as no surprise that this is even more so when unfamiliar representations like fuzzy membership functions are to be acquired. From previous projects we have learned that a user-centered approach is mandatory in complex and ill-defined knowledge domains such as internal medicine. This paper describes the knowledge acquisition framework that has been developed in order to make easier and more accessible the three main tasks of: (a) defining medical concepts; (b) providing appropriate interpretations for patient data; and (c) constructing inferential knowledge in a fuzzy knowledge representation framework. Special emphasis is laid on the motivations for some system design and data modeling decisions. The theoretical framework has been implemented in a software package, the Knowledge Base Builder Toolkit. The conception and the design of this system reflect the need for a user-centered, intuitive, and easy-to-handle tool. First results gained from pilot studies have shown that our approach can be successfully implemented in the context of a complex fuzzy theoretical

* Corresponding author. Tel.: +43-1-40400-6668; fax: +43-1-40400-6667.
E-mail address: karl.boegl@univie.ac.at (K. Boegl).

0933-3657/$ – see front matter © 2002 Elsevier B.V. All rights reserved.
PII: S0933-3657(02)00073-8

framework. As a result, this critical aspect of knowledge-based system development can be accomplished more easily.
© 2002 Elsevier B.V. All rights reserved.

Keywords: Fuzzy set theory; Knowledge acquisition; Knowledge representation; Knowledge-based systems; Medical consultation system

1. Introduction

In this paper, we describe the knowledge representation and the knowledge acquisition procedures that support medical experts to add, edit and update their knowledge in MedFrame/CADIAG-IV. First, we outline the representational framework by introducing some of the basic concepts and by explaining how simple and complex medical data and findings can be acquired and defined. Next, we explain the different steps and tools that are available to the user to establish fuzzy relationships between different medical entities. These two steps—the acquisition of all elements and their relationships—are the prerequisite for MedFrame/CADIAG-IV's inference processes which use real patient data to propose diagnostic and therapeutic hypotheses. Finally, we present a prototypical knowledge acquisition software, the Knowledge Base Builder Toolkit, that we have implemented to provide assistance for the domain experts and discuss some issues that require further research.

One of the crucial tasks in knowledge-based system development is to acquire domain knowledge that performs at least at the quality levels of true domain experts, does not impose additional representational burden or extra rigidity, allows for task-adequate, user-centered acquisition steps, and is adaptable to specialized uses. These requirements are especially true for the complex field of medical diagnosis and therapy planning. Due to the abundance, complexity, and uncertainty of medical knowledge, knowledge acquisition in this field is an especially difficult and time-consuming task [9,26].

Medical knowledge, especially the nature of the relationships between symptoms, physical signs, laboratory data, clinical findings, and diagnostic hypotheses can be characterized as a collection of empirical facts, statistical data, scientific cause–effect structures, and human experience. Uncertainty, with variations such as vagueness and imprecision, can be found at almost every step in medical reasoning. Certainty factors in computer systems like MYCIN [23], Bayesian inference models in ILIAD [27], Bayesian belief networks [21], or QMR's frequency weights and evoking strengths [18] are prominent examples of different approaches to capture some of that uncertainty. CADIAG-II, the predecessor of MedFrame/CADIAG-IV, was one of the first medical expert systems to successfully apply fuzzy set theory and fuzzy inference rules to a variety of medical fields such as rheumatology, gastroenterology and radiology [1,4,6,16].

The building of high-quality medical knowledge bases for medical consultation systems requires the experts to state their relevant knowledge as concise and logically correct as possible. Unfortunately, in many medical fields, there is often only few 'proven' knowledge to acquire: physicians usually feel uncomfortable to add their 'insights' and 'useful associations' as crisp rules, even when allowed to add some uncertainty to their inferences

[20]. MedFrame/CADIAG-IV tries to smoothen this translation by allowing for—and supporting the acquisition of—fuzziness in almost all knowledge-related steps (i.e. fuzzy representations, fuzzy associations, and fuzzy inferences). Experts are thus not forced to 'sharpen' or 'strengthen' their knowledge just to make it useful for computer use. Many types of uncertainty in data interpretations and the difficulty to create logically 'definitive' or 'true' inferences are accounted for in non-computerized situations by the use of imprecise language. The idea to use 'computations' even for fuzzy linguistic concepts was of course one of the motivations to establish the scientific study of fuzzy reasoning [29].

As already stated above, MedFrame/CADIAG-IV is a fuzzy medical consultation system in that it uses fuzzy methods for almost all knowledge processing tasks: it accepts fuzzy inputs, operates on fuzzy sets with fuzzy rules, and produces fuzzy sets as output. Of course, whenever appropriate or desired the system can defuzzify or approximate its statements to crisp values. For example, rank-ordered lists of confirmed, possible, and excluded diagnoses can be produced to help physicians to direct their next examination steps.

However, the goal of consultation systems such as MedFrame/CADIAG-IV is–in contrast to some early expert systems—not solely to come up with the 'best' diagnostic hypothesis for a given set of data and findings. Several studies in the past have shown that computerized systems are more useful and better accepted if the users receive support for their own diagnostic thinking styles and if they can structure the vast amount of their knowledge in representations that help them to understand complex connections and potential needs for refinement. The issue is then to assist the physician in the differential diagnostic process and to allow a smooth transition from 'established' medical knowledge to personal judgment and experience. This support can be achieved by relying on relatively simple, modifiable associations which indicate possible medical causes (diseases, syndromes) that may explain the patient's current data. With the help of a consultation system, these small knowledge components can be maintained and feasible connections to other knowledge components can be found, used, or 'debugged' in the context of real patient data.

In addition, a consultation system should not stop with a list of feasible conclusions, but should propose further examinations and tests that may help to confirm or exclude some of the hypotheses. It should also indicate pathological findings which are not yet accounted for. After reviewing the system's recommendations and explanations, clinicians can interactively refine their hypotheses until they reach acceptable decisions. In order to be accepted in a clinical context, not only the 'best' or 'common' hypotheses ought to be presented—it might be more realistic and beneficial to also emphasize rare or uncommon explanations which might otherwise be overlooked.

When these requirements are reviewed, it becomes clear that only a highly interconnected, dynamic network of medical entities can achieve such demanding tasks. It is impossible to fine-tune such a network to optimize its performance or to enforce its logical correctness. In contrast, the representational framework has to be expressive and flexible enough to account for all the possible connections along with easy means to express the inherent vagueness and the complexity of the many interconnections. Fuzzy representations implement this flexibility while still allowing for efficient and adaptive computations.

From a technical perspective, MedFrame/CADIAG-IV extends previous implementations in many ways. For example, it marks the transition from a system on a centralized host with a terminal-based interface and proprietary data representation schemes to a web-based, client–server solution with graphical user interfaces, object-oriented data models, and compatibility with many established standards. Further details about MedFrame/CADIAG-IV's design rationale, knowledge representation framework, knowledge-based implementation, and inference processes can be found in [15]. The conceptual models and the implementation details of MedFrame/CADIAG-IV's knowledge acquisition system and the Knowledge Base Builder Toolkit (which is described later) are described in [7].

In order to preserve the successful inference characteristics and large, specialized patient databases of predecessor systems, full backward compatibility was a special design requirement. Although MedFrame/CADIAG-IV is a completely new implementation that preserves the general mode of operation from CADIAG-II—with major changes at the levels of design, formalisms, representation, computations, and implementation—the most profound and most visible changes have been introduced to ease acquisition and maintainability of medical knowledge. Former versions required deep insights into the representational details of the system to be able to specify complex knowledge correctly. Detailed task and resource analyses of the knowledge acquisition and elicitation processes of previous users led to a specification of a simpler and yet more flexible acquisition process.

2. Basic concepts of fuzzy knowledge representation

One of the main formal characteristics of MedFrame/CADIAG-IV is the use of fuzzy set theory and fuzzy logic. Therefore, we introduce the concepts and notations of fuzzy set theory as used throughout this article.

2.1. Fuzzy sets

In medical science, it is rarely possible to give exact definitions or descriptions of medical concepts and relationships between concepts. For example, the assignment of laboratory test results to normal or pathological ranges is arbitrary in borderline cases and depends on the subjective estimation of the physician. Furthermore, precise descriptions of relationships between findings and diseases can rarely be given [2].

To express vagueness and imprecision of medical entities and relationships we employ the theory of fuzzy sets [28]. If U is any set, A is a *fuzzy subset* of U if there is a function μ_A (called membership function) such that

$$\mu_A : U \to [0, 1], \tag{1}$$

$$A = \{(x, \mu_A(x)) | x \in U\}. \tag{2}$$

A is a *fuzzy set* if there is a U such that A is a *fuzzy subset* of U. The set U is referred to as the base set or the universe of discourse. The *membership function* is a generalization of the characteristic function of ordinary sets, fuzzy sets thus being generalizations of ordinary sets. An element of a fuzzy set is specified by a pair $A(u)/u$, where $A(u)$ is the degree of membership of that element with that set. Finite fuzzy sets can thus be specified by the listing of such pairs $\{1/u_1, 0.8/u_2, 0.3/u_3, \ldots\}$. In addition, in CADIAG-II and MedFrame/CADIAG-IV elements of fuzzy sets can be specified as ν (‘unknown’) so that the membership function specifying a fuzzy set A of a set U is defined as $A : U \to [0, 1] \cup \nu$. The fuzzy power set of U, denoted by $\mathscr{F}(U)$, is defined as the set of all fuzzy sets of the set U.

2.2. Type-2 fuzzy sets

A type-2 fuzzy set $\tilde{A}$ of a set U is a fuzzy set whose degrees of membership are themselves fuzzy sets. This definition extends the original concept of a type-2 fuzzy set as given in [11], as this fuzzy set is not restricted to [0, 1] here. $\tilde{A}$ is defined by $\tilde{A} : U \to \mathscr{F}(V)$, where $\mathscr{F}(V)$ is the fuzzy power set of an ordinary set V.

2.3. Fuzzy relations

A *fuzzy relation* R between a set U and a set V is a fuzzy set of the Cartesian product $U \times V$ ($U \times V$ is the set of all ordered pairs (u, v), $u \in U$, $v \in V$). The membership function $R : U \times V \to [0, 1]$ assigns to every pair (u, v) a degree of membership.

As an illustration, if $U = \{$fever, dyspnea$\}$ and $V = \{$pulmonary embolism, pneumonia$\}$ then a fuzzy relation of ‘association’ of members of U and V might be expressed as ‘association’ = {0.1/(fever, pulmonary embolism), 0.95/(fever, pneumonia), 0.9/(dyspnea, pulmonary embolism), 0.8/(dyspnea/pneumonia)}.

In analogy to the definition of type-2 fuzzy sets, a type-2 fuzzy relation $\tilde{R}$ between a set U and a set V is a fuzzy set of Cartesian product $U \times V$, where the membership function $\tilde{R} : U \times V \to \mathscr{F}(W)$ assigns to every pair (u, v) a fuzzy set of the set W.

2.4. Linguistic variables

Zadeh introduced the concept of linguistic variables—variables whose values are linguistic terms rather than numerical values—to provide a means of approximate characterization of phenomena that are too complex or too ill-defined to be amenable to a description in conventional quantitative terms [29]. In medical diagnosis, which is strongly influenced by human perception and judgment, in many cases it is more appropriate to describe the underlying knowledge by means of linguistic variables than by quantitative descriptions.

A *linguistic variable* is characterized by the quintuple $\langle X, T(X), U, G, \mathrm{M}\rangle$ in which: (a) X is the name of the variable; (b) $T(X)$ is the *term set* of X, that is, the set of its linguistic values; (c) U a universe of discourse; (d) G a syntactic rule that generates the terms $T(X)$; and (e) M a semantic rule which associates with each linguistic value x its meaning, $M(x)$, where $M(x)$ denotes a fuzzy set of U. The syntactic rule G specifies the manner in which the linguistic values of $T(X)$ are generated. The meaning of a linguistic value x is characterized by a *compatibility function*, $C : U \to [0, 1]$, which associates with each u in U its compatibility with x.

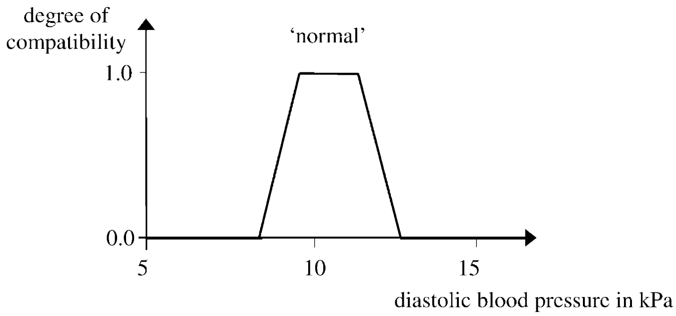

Fig. 1. The meaning of the linguistic term is characterized by a compatibility function which associates to every possible diastolic blood pressure a degree of compatibility with the concept 'normal diastolic blood pressure'.

As an example, the linguistic variable 'diastolic blood pressure' can be defined as $X =$ 'diastolic blood pressure', with $T(X) = \{\text{very low, low, normal, high, very high}\}$, and $U = [5, 17]$ kPa. The linguistic value 'normal' is represented by a compatibility (membership) function as depicted in Fig. 1.

3. Fuzzy knowledge representation in MedFrame/CADIAG-IV

An overall description of the knowledge representation formalism of MedFrame/CADIAG-IV can be found in detail in [15]. The following section gives a short introduction and a brief summary of the knowledge representation framework as necessary for understanding the knowledge acquisition tasks.

3.1. Representation of medical concepts

For all inferencing tasks, some kind of 'knowledge' is usually available to deal with the challenge of arriving at useful conclusions even though the task is usually ill-defined. In MedFrame/CADIAG-IV, the *medical concept* type represents the top-most abstraction level of medical knowledge items. The most important characteristic of a medical concept is that it is uniquely defined by a set of facets, that are based on SNOMED international module concepts [22], and a variable number of qualifiers. Thus, all entities can be defined and identified in a stringent, coherent, and semantically meaningful way.

Semantically, we distinguish between two subtypes of medical concepts which are the basic knowledge types within the knowledge representation framework.

3.1.1. Medical entities

Findings, diseases, and therapies are the basic building blocks for all possible statements about medical concepts. The definition of such knowledge constitutes the granularity of the system (i.e. the most atomic ingredients that can be reasoned about) and follows other approaches in the medical knowledge-based systems field.

3.1.2. Medical data

At another level than qualitative medical entities, medical data describe quantitative medical concepts such as measurements, results from physical examinations, and laboratory data (e.g. height, duration of morning stiffness, serum glucose levels).

The sets of findings F, diseases D, therapies T, the set of medical entities E ($E = F \cup D \cup T$), and the set of patients P are ordinary sets. When considering a *single* patient, the vagueness and imprecision of medical entities is taken into account by introducing fuzzy sets: (a) F^+ (specified by $F^+ : F \to [0, 1]$) denoting the patient's findings; (b) D^+ ($D^+ : D \to [0, 1]$) denoting the patient's diseases; and (c) T^+ ($T^+ : T \to [0, 1]$) denoting the therapies administered to the patient (which may also be fuzzy as for example in the case of a prescribed 'mild low-fat diet').

An example of this can be given by considering the set of findings F as defined in MedFrame/CADIAG-IV and the fuzzy set F^+ representing the findings of a single patient $F^+ = \{1.0/f_5 =$ 'hyperuricemia', $0.7/f_{31} =$ 'swelling of the ankle joint', $v/f_{72} =$ 'family history of gout', $\ldots\}$.

3.2. Data-to-entity conversion

As stated above, the fuzzy set F^+ denotes a single patient's findings and represents the respective medical concepts on a symbolic level. Since MedFrame/CADIAG-IV reasoning mechanisms operate at the level of symbolic concepts a data-to-entity conversion has to be employed to transform numeric medical data (e.g. body temperature: 39 °C) into symbolic concepts (e.g. 'body temperature normal', 'body temperature raised') (Fig. 2). For this purpose we define a linguistic variable that denotes the observed medical data. The term set of this linguistic variable represents the *interpretation categories* of the respective medical

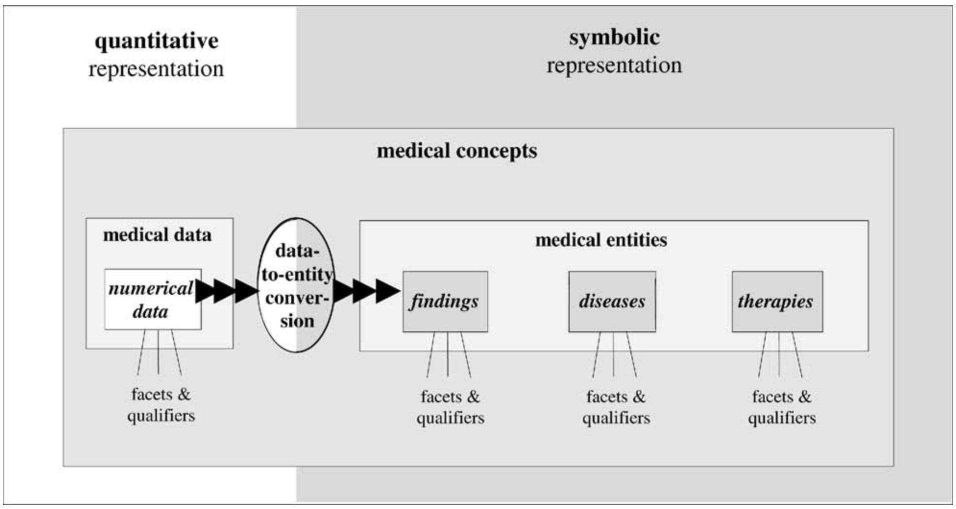

Fig. 2. Different types of medical concepts in MedFrame/CADIAG-IV. Quantitative medical data are converted into symbolic representations in a data-to-entity conversion step.

data or, in other words, the set of symbolic concepts (medical entities) a medical data can be converted into. The universe of discourse U is defined by the range of possible values of the medical data. The compatibility functions characterize the interpretation categories and assign to every value $u \in U$ of a medical data a degree of membership of u in the interpretation categories. At run-time, when actual patient data are used, these definitions will translate data values and assessments into symbolic, but fuzzified medical entities.

The interpretation of medical data in a clinical setting is not always straightforward. Whether or not a laboratory value is considered to be 'normal' might depend on the values of other measurements. Especially for patho-physiologically interdependent data (e.g. 'hematocrit level' and 'serum erythropoietin level') the interpretation is a multi-dimensional problem. Formally, these combinations require the use of type-2 fuzzy sets. In the case of two-dimensional problems, a practical solution to acquire the membership functions is to define data-to-entity conversions for discrete values of the second dimension and to compute the full fuzzy membership function through interpolations (for details see [8,17]).

A special variation of multi-dimensional data is the interpretation of fuzzy temporal trends, where time is considered to be a special kind of medical data. In many medical assessments, not only the absolute value of a medical parameter is important, but also the variation and course of the value over time (e.g. blood glucose tolerance tests). The computation of a temporal course, however, requires the computation of several compatibility functions over a period of time [24,25].

In many situations, the interpretation of actual patient data is only reasonable in special circumstances. MedFrame/CADIAG-IV allows the specification of fuzzy contexts that are used to qualify specific interpretations. For example, many medical entities and especially quantitative medical data give rise to different interpretation depending on age, sex, or special conditions such as pregnancy or preexisting diseases. Contexts can now be defined for all parameters individually. This freedom allows the medical expert to define different membership functions for different interpretations (e.g. given the context of pregnancy, low glucose levels may have other membership functions than in case of absence of gravidity). The selection of appropriate contexts and the computation of compatibility function for the selected context(s) are described in more detail in [8].

It has become a standard requirement of knowledge representations that they are transparent and meaningful to the people maintaining them. By separating out most of the difficult data-related decisions (e.g. 'does a temperature of 37.1° imply fever?') even at the more complex levels of multi-dimensional, time-dependent, or context-dependent interpretation, MedFrame/CADIAG-IV allows the decision maker to concentrate on the inferential reasoning process, knowing that appropriate attention is given to all possible interpretations of actual patient data.

3.3. Representation of relationships between medical concepts

At the core of MedFrame/CADIAG-IV's inference processes are reasoning mechanisms that deal with symbolic entities. These entities are connected by means of fuzzy rules. The basic inference process follows these relationships and recursively calculates fuzzy values

Table 1
The four types of relationships between antecedents and consequents

Positive association	$A \xleftarrow{F_p} C$	Frequency of occurrence of the antecedent with the consequent
	$A \xrightarrow{S_p} C$	Strength of confirmation of the antecedent for the consequent
Negative association	$A \xleftarrow{F_n} \neg C$	Frequency of occurrence of the antecedent with *not* the consequent
	$A \xrightarrow{S_n} \neg C$	Strength of exclusion of the antecedent for the consequent

The type of association determines the pair of applicable fuzzy relationships.

for connected entities. As already stated above, this type of rules employs an implication operator that demands symbolic concepts as antecedents and consequents. In general, we distinguish rules that denote *positive associations* between antecedents and consequents from rules that denote *negative associations* (a similar approach has been proposed in [12,13,23]). Both are characterized by a pair of fuzzy sets which in the case of a positive association are denoted as F_p the 'frequency of occurrence of the antecedent with the consequent' and S_p the 'strength of confirmation of the antecedent for the consequent', and in the case of a negative association as F_n 'frequency of occurrence of the antecedent with *not* the consequent' and S_n the 'strength of exclusion of the antecedent for the consequent' (Table 1). These fuzzy sets represent and qualify the vagueness and uncertainty of the relationships and are interpreted as fuzzy numbers. Actually, they are fuzzy numbers of the set $U = [0, 1]$ (this definition slightly extends the definition as given in [11]). For the definition of their characterizing membership functions we use the function $\Pi(x; \alpha, \beta, \gamma, \delta)$ for linear transition functions and $P(x; \alpha, \beta, \gamma, \delta)$ for non-linear transition functions (a graphical representation of Π is depicted in Fig. 3, examples of the non-linear type P are shown in Fig. 8).

$$\Pi(x; \alpha, \beta, \gamma, \delta) = \begin{cases} 0 & x < \alpha, \\ \dfrac{x-\alpha}{\beta-\alpha} & \alpha \le x < \beta \quad \wedge \alpha < \beta, \\ 1 & \beta \le x \le \gamma, \\ 1 - \dfrac{x-\gamma}{\delta-\gamma} & \gamma < x \le \delta \quad \wedge \gamma < \delta, \\ 0 & x > \delta. \end{cases}$$

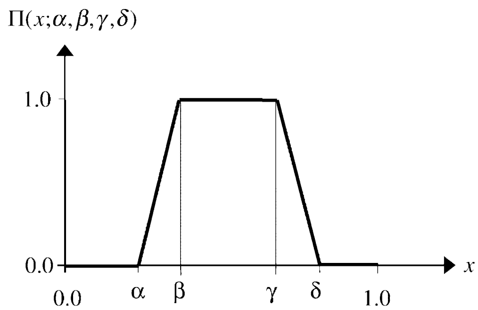

Fig. 3. Graphical representation of function Π.

$$P(x;\alpha,\beta,\gamma,\delta)=\begin{cases}0 & x<\alpha,\\ 2\left(\dfrac{x-\alpha}{\beta-\alpha}\right)^2 & \alpha\le x\le\dfrac{\beta-\alpha}{2}\quad\wedge\,\alpha<\beta,\\ 1-2\left(\dfrac{\beta-x}{\beta-\alpha}\right)^2 & \dfrac{\beta-\alpha}{2}<x<\beta\quad\wedge\,\alpha<\beta,\\ 1 & \beta\le x\le\gamma,\\ 1-2\left(\dfrac{x-\delta}{\delta-\gamma}\right)^2 & \gamma<x\le\dfrac{\delta-\gamma}{2}\quad\wedge\,\gamma<\delta,\\ 2\left(\dfrac{\delta-x}{\delta-\gamma}\right)^2 & \dfrac{\delta-\gamma}{2}<x\le\delta\quad\wedge\,\gamma<\delta,\\ 0 & x>\delta.\end{cases}$$

The overall *rule base* of MedFrame/CADIAG-IV, that is the total set of rules in the knowledge base, can be represented by two pairs of type-2 fuzzy relations which describe the relationships between the tuples $\langle e_i, e_j\rangle$ of medical entities (e_i, $e_j \in E$). For positive associations, this pair consists of the type-2 fuzzy relations $\tilde{R}_\mathrm{F}^\mathrm{p}$ 'frequencies of occurrence of the antecedents with the consequents'

$$\tilde{R}_\mathrm{F}^\mathrm{p} : E\times E \to \mathscr{F}([0,1]),$$

and $\tilde{R}_\mathrm{F}^\mathrm{p}$ 'strengths of confirmation of the antecedent for the consequent'

$$\tilde{R}_\mathrm{S}^\mathrm{p} : E\times E \to \mathscr{F}([0,1]).$$

For negative associations, this pair consists of the type-2 fuzzy relations $\tilde{R}_\mathrm{F}^\mathrm{n}$ 'frequencies of occurrence of the antecedents with *not* the consequents'

$$\tilde{R}_\mathrm{F}^\mathrm{n} : E\times E \to \mathscr{F}([0,1]),$$

and the $\tilde{R}_\mathrm{S}^\mathrm{n}$ 'strengths of exclusion of the antecedent for the consequent'

$$\tilde{R}_\mathrm{S}^\mathrm{n} : E\times E \to \mathscr{F}([0,1]).$$

The antecedents of rules are not restricted to single medical concepts but may be logical combinations of medical concepts and operators forming an operator tree. Two special knowledge representations, *disease profiles* and *explicit rules*, combine several medical entities in more complex ways. *Disease profiles* are intermediate representations that are approximations of the experts' mental models of diseases. They combine, in a table-like manner, all the defined medical entities (e.g. symptoms, physical signs, lab test results, clinical findings, examinations, syndromes, diseases, therapies) that are related to another entity (usually a disease or diagnostic hypothesis).

The most common and most descriptive use of fuzzy relationships is the connection of findings with diseases, but the same relationships can be used for any number and types of medical entities. Thus, medical entities that act as findings in a disease profile can themselves be complex fuzzy concepts that are composed of many combinations of other 'disease profiles'. The sum of all these fuzzy relationships constitute a network of linked concepts that defines the knowledge base, which is used by the inference processes. It is worthwhile mentioning that these disease profiles are not explicitly represented in MedFrame/CADIAG-IV's data model-they are intermediate representations of a concept that many physicians seem to be comfortable with and are generated from the 'normal' fuzzy relationships discussed above.

4. Acquisition of inferential knowledge

Given the complexity of the network of entities that is to be acquired and given the difficulty in assessing even a simple one-to-one relationship between a finding and a disease, a guided, stepwise knowledge-acquisition process has been established. Knowledge acquisition starts with the definition of medical concepts. The definition of medical data, which are a special subtype of medical concepts, additionally requires the definition of data-to-entity conversion rules to allow a processing of the represented information in the inference process. For this purpose, linguistic variables and adequate interpretation categories must be specified. Subsequently, the connection between medical entities can be established by fuzzy relationships.

4.1. Definition of the data-to-entity conversion rules

As stated above, the data-to-entity conversion mechanism transforms a patient's raw (numerical) medical data into symbolic concepts (medical entities). This step may be compared to a physician's assessment and interpretation of measured patient data. The definition of the medical entities that can be derived from medical data and the definition of the conversion rules require medical knowledge and are therefore a crucial part of the knowledge acquisition process.

When a medical datum is created, the knowledge acquisition tool formally employs a corresponding linguistic variable and asks for the total range of possible values of the medical datum. This range of possible values forms the universe of discourse U of the linguistic variable. Subsequently, the term set (interpretation categories) and the respective compatibility functions must be defined in three steps.

In a first step, an appropriate number of interpretation categories has to be chosen. In the example depicted in Fig. 4, three interpretation categories ('pathologic', 'suspect' and 'normal') were defined representing the findings 'pathologic chymotrypsin level in stool', 'suspect chymotrypsin level in stool' and 'normal chymotrypsin level in stool'. The more categories are established the more decisions and maintenance efforts are required to make them useful and the higher is the discriminative expressiveness.

In a second step, the selected interpretation categories can be defined as either *exclusive* or *inclusive* categories (Fig. 5). Exclusive (or complementary) categories are used when some medical data point could be used to differentiate between possible interpretations (e.g. the categorization of the 'total serum bilirubin level' can be useful to differentiate between hyperbilirubinemia caused by hemolysis, hepatitis or cholestasis). Inclusive categories are more useful when some data points imply gradual changes (e.g. both 'increased' and 'heavily increased' serum glucose levels are signs of diabetes, but they also allow for an improved diagnosis of the type and stage of diabetes).

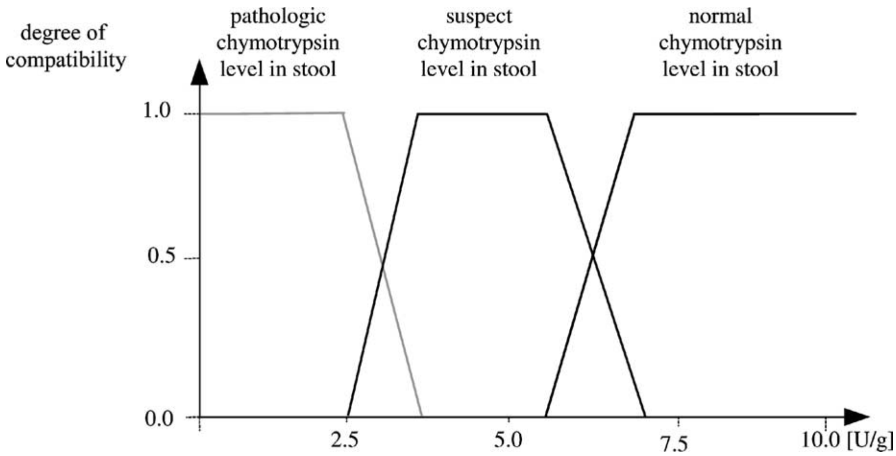

Fig. 4. The term set of the linguistic variable ‘chymotrypsin level in stool’ and the specification of the compatibility functions.

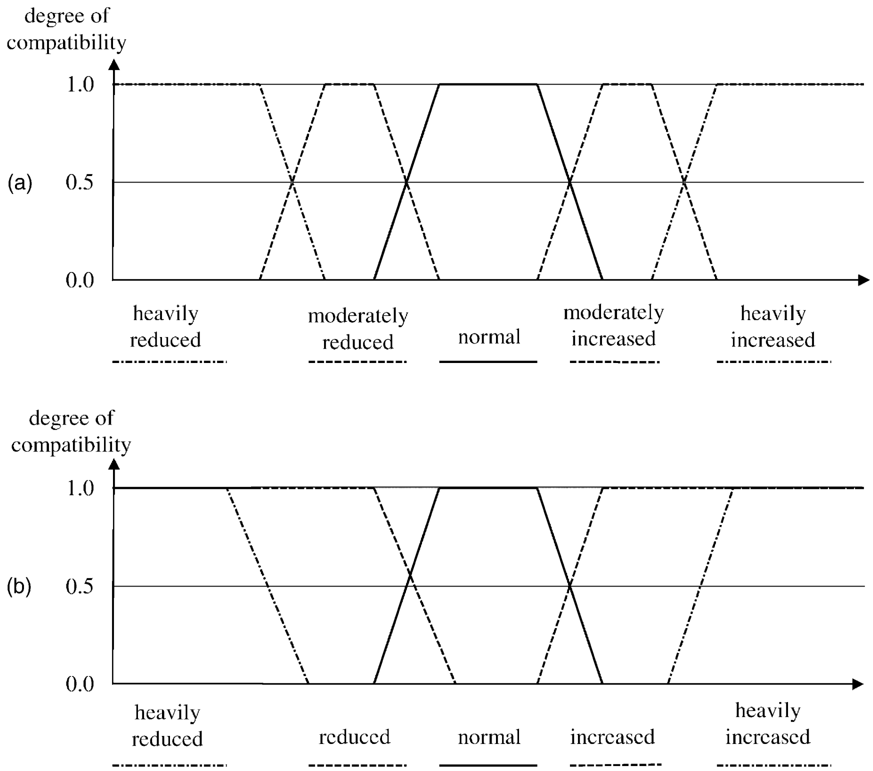

Fig. 5. Exclusive vs. inclusive interpretation categories for medical data.

As a last step, the compatibility functions characterizing the interpretation categories need to be defined. For a given medical data (x-axis) the degree of compatibility (y-axis) to the interpretation categories is assessed. A compatibility degree of ‘1’ denotes the range of data values for which an interpretation is always correct (in the case of exclusive categories it is also the only correct interpretation). The values of medical data to which no interpretation category are assigned as completely compatible (‘1’) describe the vagueness of the medical entities. These borderline values may be partially compatible with usually two adjacent medical interpretation categories.

Different fuzzy membership functions can be used to model the desired fuzziness characteristics outside of full compatibility. As a default, linear transition functions between two categories are proposed. The data-to-entity conversion process is completed when the whole defined range of possible data values is covered. It is, however, possible to just define parameter ranges that are ‘pathologic’ with respect to a certain class of diseases.

Partly, the medical knowledge that is necessary for the definition of the interpretation categories and the specification of the characterizing compatibility functions may be obtained from text books, reference tables, and patient databases, but in many cases depends—due to a lack of strict medical guidelines—on the subjective judgment of the medical expert.

The interpretation of medical data that are patho-physiologically interdependent and of data over time requires a type-2 fuzzy set representation of the interpretation categories. Accordingly, an interpretation category is characterized by a compatibility function $C : U \rightarrow \mathscr{F}(V)$. A practical solution to acquire the compatibility function is to define the fuzzy sets $V_i : V \rightarrow [0, 1]$ for a selected set $\bar{U} = \{u_1, \ldots, u_n\}$ of discrete values $u_i \in U, i \in N$, and define the compatibility function as $C : \bar{U}\mathscr{F}(V)$. If the set $\bar{U}$ is chosen adequately (i.e. $u_1 = \min U$, $u_n = \max U$, the granularity being fine enough) the compatibility function can be approximated for the whole range of $u \in U$ by linear interpolation. In the case of two interdependent medical data, the compatibility function assigns to the pair (u, v) of a medical datum $u \in U$ and a medical datum $v \in V$ the degree of compatibility with the interpretation category it specifies. In the case of temporal medical data, the compatibility function assigns to a temporal course $u(t)$ of a medical datum $u \in U$ ($t \in \mathrm{TI}$ denotes time) the degree of compatibility to a set of predefined temporal courses (which are special kinds of interpretation categories). Here, the fuzzy sets $U_i : U \rightarrow [0, 1]$ are defined at distinct time stamps $t_i \in \mathrm{TI}$, $i \in N$. As an illustration, the compatibility function of the fuzzy temporal trend ‘normal glucose tolerance test’ is depicted in Fig. 6. During knowledge acquisition, the user can select from a predefined set of fuzzy time-trends (e.g. constant, rising, falling, oscillating) with special parameters (onset-time, onset-value, direction) which cover many typical situations.

As already mentioned, compatibility functions employed in the data-to-entity conversion may also be defined in a context-dependent manner. In terms of the knowledge acquisition process, a default context, which is used whenever no specialized context is applicable, is always defined in a first step. Subsequently, a number of appropriate fuzzy contexts can be defined (or reused). For example, many medical entities and especially quantitative medical data give rise to different interpretation depending on age, sex, or special conditions such as pregnancy or preexisting diseases. Contexts can now be defined for all parameters individually. This freedom allows the medical expert to define different

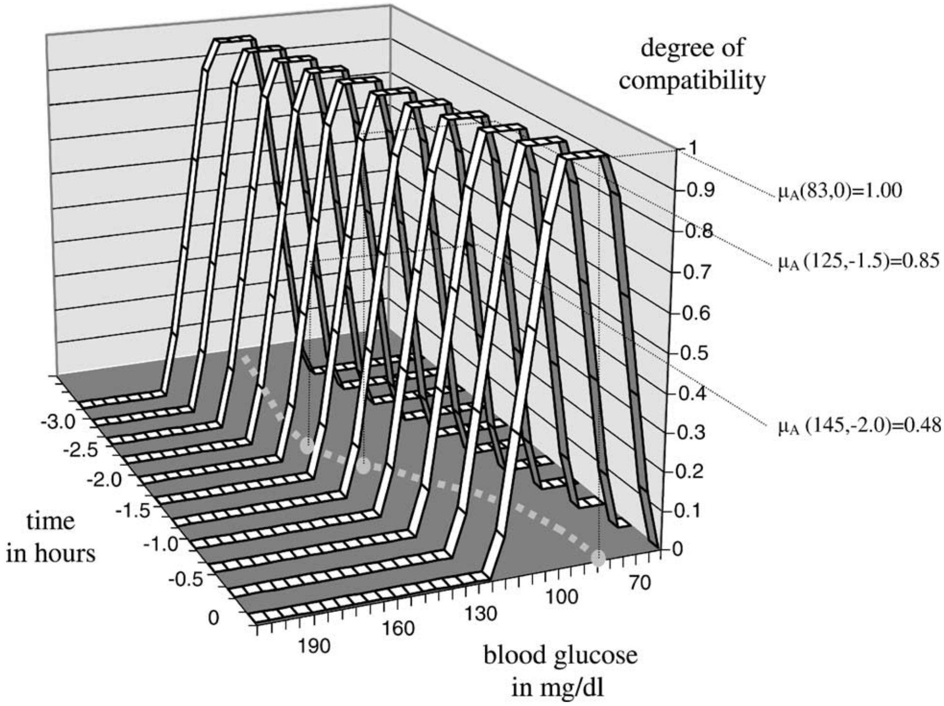

Fig. 6. Type-2 fuzzy set representation of the trend 'normal glucose tolerance test'. The degree of compatibility of a temporal course with the trend is calculated by assessing the degrees of compatibility to the temporal course at the distinct time stamps.

membership functions for different interpretations (e.g. given the context of pregnancy, normal glucose levels have different membership functions). The selection of appropriate contexts and the computation of compatibility function for the selected context(s) are described in more detail in [8].

4.2. Stepwise refinement of fuzzy medical relationships

The definition of the fuzzy relationships between medical entities is guided by a stepwise knowledge acquisition process. In a few steps, the connection between two entities can be established by medical experts starting from a simple association and ending up with fuzzy membership functions. These steps are optional refinements that are supported by a dedicated knowledge acquisition methodology. If an expert has acquired some familiarity with all the available forms of knowledge representation, direct interaction with the appropriate step is possible without a need to go through all previous steps. Although the interpretations in steps 1 and 2 could be represented using traditional logical formulas, they are also defined as fuzzy membership functions. Formally, all relationships express some degree of vagueness and uncertainty and are thus interpreted as fuzzy numbers. Medical entities are unrelated to each other unless an expert adds some knowledge about a specific relationship in the following steps.

4.3. Step 1: associations

To add knowledge about the relationship between two medical entities, a first step may consist in defining either a *positive*, a *neutral*, or a *negative* association between them. These associations are appropriate whenever causal relationships or at least empirical correlations are accepted as scientific facts. For example, a positive relationship between a finding and a disease implies that medical knowledge is available to always infer (confirm) the presence of a disease whenever the finding is present. To exclude the disease whenever the finding is present, a negative association would have been used. A neutral association adds knowledge that some kind of non-causal relationship exists, which could be a descriptive relationship between a patient's characteristic such as age, sex, or race and a certain disease.

Because positive or negative evidence is hard to find, many neutral associations are likely to appear in medical domains to account for experimental and casuistic knowledge.

Formally, a positive association is represented by the fuzzy sets F_{p} and S_{p} and characterized by membership functions

$$F_{\mathrm{p}}(u) = \begin{cases} 1 & u \neq 0, \\ 0 & u = 0, \end{cases}$$

and

$$S_{\mathrm{p}}(u) = \begin{cases} 1 & u \neq 0, \\ 0 & u = 0. \end{cases}$$

A negative association is specified by

$$F_{\mathrm{n}}(u) = \begin{cases} 1 & u \neq 0, \\ 0 & u = 0. \end{cases}$$

and

$$S_{\mathrm{n}}(u) = \begin{cases} 1 & u \neq 0, \\ 0 & u = 0. \end{cases}$$

By definition, a neutral association does not have a strength of confirmation or strength of exclusion and is therefore represented by F_{p} only

$$F_{\mathrm{p}}(u) = \begin{cases} 1 & u \neq 0, \\ 0 & u = 0. \end{cases}$$

As an example, the association between the finding 'increased serum glucose level' and the disease 'diabetes' is positive and accordingly specified as depicted in (Fig. 7).

Associations can be refined in the following steps to diminish the uncertainty or lack of established causal knowledge.

4.4. Step 2: predefined CADIAG relations

The basic relationship concepts of CADIAG-I and CADIAG-II, which were useful for acquiring medical knowledge from domain experts as well as for successful application of

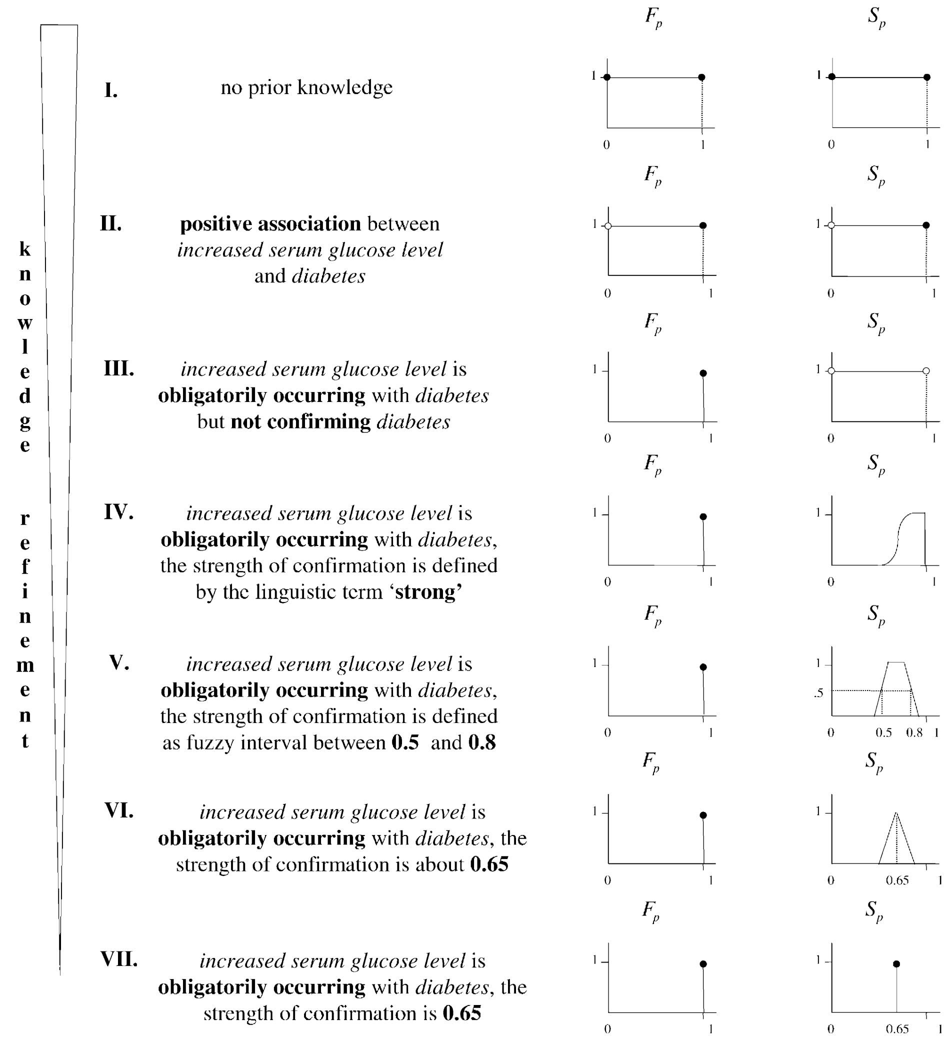

Fig. 7. Illustration of the stepwise refinement process of the relationship between two medical entities ('increased serum glucose level' and 'diabetes'). Full/empty circles denote full/no membership.

the fuzzy reasoning system, are 'frequency of occurrence' and 'strength of confirmation' [6]. They are used to differentiate positive and negative associations into several basic relationships (MedFrame/CADIAG-IV extends CADIAG-II by introducing explicit calculation of all combinations of negative evidence). These relations allow the expert to capture special kinds of relationships more easily (Table 2). Examples of these are given in Table 3.

Any relation is directed from an antecedent A to a consequent D (usually from a symptom, physical sign, lab test result, or finding to a disease) and characterized by the frequency of occurrence F_p and the strength of confirmation S_p. In a simple case, these

Table 2
The nine predefined CADIAG relations and their fuzzy-based representation

CADIAG relation		F_p	S_p	F_n	S_n
OC	Obligatorily occurring and confirming	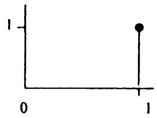	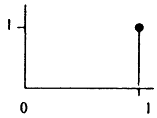	–	–
nOC	Not obligatorily occurring and confirming	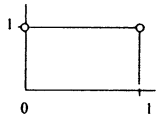	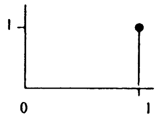	–	–
OnC	Obligatorily occurring and not confirming	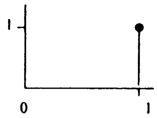	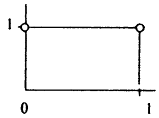	–	–
nOnC	Not obligatorily occurring and not confirming	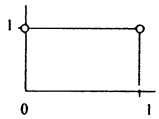	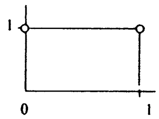	–	–
N	Neutral	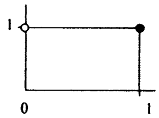	–	–	–
nOnE	Not obligatorily occurring with not the consequent and not excluding	–	–	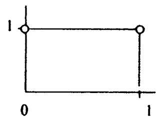	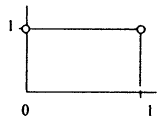
OnE	Obligatorily occurring with not the consequent and not excluding	–	–	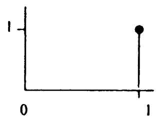	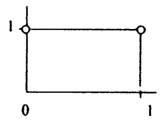
nOE	Not obligatorily occurring with not the consequent and excluding	–	–	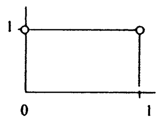	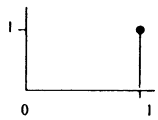
OE	Obligatorily occurring with not the consequent and excluding	–	–	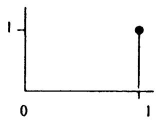	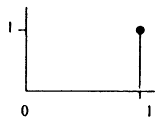

Table 3
The nine predefined CADIAG relations and case examples

CADIAG relation		Case example
OC	Obligatorily occurring and confirming	A failure of detection of galactose-1-phosphate-uridinediphosphate-galactose-transferase (antecedent) is *obligatorily occurring* with galactosaemia (consequent) and *confirms* galactosaemia by definition
nOC	Not obligatorily occurring and confirming	Intracellular urate cristalls in synovial fluid (antecedent) are *not obligatorily occurring* with gout (consequent) but a detection of intracellular urate cristalls *confirms* gout by definition
OnC	Obligatorily occurring and not confirming	An increased serum glucose level (antecedent) is *obligatorily occurring* with diabetes (consequent) but does *not confirm* diabetes
nOnC	Not obligatorily occurring and not confirming	Rheumatoid factors (antecedent) are *not obligatorily occurring* with rheumatoid arthritis (consequent) and do *not confirm* rheumatoid arthritis
N	Neutral	Gout (consequent) is observed twenty times more frequently in male patients (antecedent). However, the patient's sex has no strength of confirmation for gout and is therefore *neutral*
nOnE	Not obligatorily occurring with not the consequent and not excluding	A failure of detection of malignant cells in biopsy material of the stomach (antecedent) is *not excluding* gastric cancer (consequent). A detection of malignant cells in biopsy material of the stomach is *not obligatorily occurring* in other malignant diseases
OnE	Obligatorily occurring with not the consequent and not excluding	Physical well-being (antecedent) does *not exclude* illness (consequent) but is *obligatorily occurring* with health (= not illness)
nOE	Not obligatorily occurring with not the consequent and excluding	A failure of detection of malignant cells in a resected stomach *excludes* gastric cancer. A detection of malignant cells in a resected stomach is *not obligatorily occurring* in other malignant diseases
OE	Obligatorily occurring with not the consequent and excluding	Vital signs (antecedent) *exlude* death (consequent) and are *obligatorily occurring* with life (= not death)

parameters can be crisp judgments (e.g. obligatorily present *and* confirming, OC) which could be used like standard (predicate) logical relations. Of course, all logical combinations including negations are possible which results in four different relations (e.g. obligatorily present and not confirming, OnC).

Additionally, the same relations have to be defined for the absence of the consequent (negation of D, $\neg D$), because a low or zero value of strength of confirmation is semantically different from an exclusion. Thus, the frequency of occurrence F_n (obligatorily for $\neg D$) and the strength of exclusion S_n (excluding or not) are defined as well.

The special case of a 'neutral' relationship has been defined to include only F_p, without qualifying S_p. In summary, the user has nine possible relations by specifying F_p and S_p, or F_n and S_n, respectively.

As the finding 'increased serum glucose level' is obligatorily occurring with the disease 'diabetes' but does not confirm the disease 'diabetes', we use the predefined CADIAG relation OnC for the further specification of the relationship. The membership functions are defined as

$$F_p(u) = \begin{cases} 1 & u = 0, \\ 0 & u \neq 0. \end{cases}$$

and

$$S_p(u) = \begin{cases} 0 & u = 1, \\ 1 & u \neq 0 \quad \wedge u \neq 1, \\ 0 & u = 0. \end{cases}$$

So far, evidence and counter-evidence for a medical entity are represented to be either 'present' or 'absent' given some other medical entity. In order to qualify the degree of relationship, fuzzy memberships are used.

4.5. Step 3: fuzzy membership definition by linguistic terms

It is certainly possible to adopt frequentistic or probabilistic interpretations of the relationships F_p, S_p, F_n, and S_n, and in fact, many experts use just that kind of knowledge to 'fuzzify' their relationships. The knowledge acquisition framework therefore supports another approach to deal with uncertainty and vagueness by providing a set of linguistic terms to define the membership functions F_p, S_p, F_n and S_n. The term sets of the linguistic variables characterize the gradation of the strengths of the relationships. We offer a set of seven predefined linguistic terms for the definition of F_p and F_n, or S_p and S_n, respectively (characterized by the compatibility functions defined in Tables 4 and 5, Figs. 8 and 9, respectively).

Table 4
Linguistic terms used for the definition of F_p and F_n

Linguistic terms	Compatibility functions
Almost always	$M(x) = P(x; 0.88, 0.94, 1.00, 1.00)$
Very often	$M(x) = P(x; 0.75, 0.88, 1.00, 1.00)$
Often	$M(x) = P(x; 0.50, 0.75, 1.00, 1.00)$
Medium	$M(x) = P(x; 0.25, 0.50, 0.50, 0.75)$
Seldom	$M(x) = P(x; 0.00, 0.00, 0.25, 0.50)$
Very seldom	$M(x) = P(x; 0.00, 0.00, 0.13, 0.25)$
Almost never	$M(x) = P(x; 0.00, 0.00, 0.06, 0.13)$

Table 5
Linguistic terms used for the definition of S_p and S_n

Linguistic terms	Compatibility functions
Almost definitely	$M(x) = P(x; 0.88, 0.94, 1.00, 1.00)$
Very strong	$M(x) = P(x; 0.75, 0.88, 1.00, 1.00)$
Strong	$M(x) = P(x; 0.50, 0.75, 1.00, 1.00)$
Medium	$M(x) = P(x; 0.25, 0.50, 0.50, 0.75)$
Weak	$M(x) = P(x; 0.00, 0.00, 0.25, 0.50)$
Very weak	$M(x) = P(x; 0.00, 0.00, 0.13, 0.25)$
Almost no	$M(x) = P(x; 0.00, 0.00, 0.06, 0.13)$

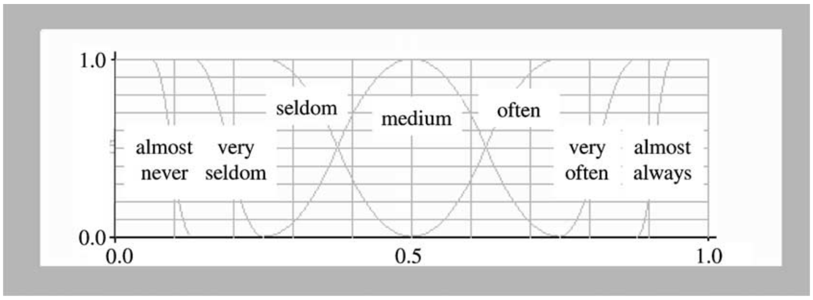

Fig. 8. Compatibility functions of the linguistic frequency of occurrence with D and frequency of occurrence with $\neg D$ variables.

In our example, we use the linguistic term 'strong' for the definition of the S_p. F_p has already been defined as obligatorily occurring in the previous step. The membership functions are defined as

$$F_p(u) = \begin{cases} 1 & u = 1, \\ 0 & u \neq 1, \end{cases}$$

and

$$S_p(u) = P(u; 0.5, 0.75, 1, 1).$$

4.6. Step 4: manipulating membership functions

Alternatively, or as a further refinement to the use of linguistic terms, the fuzzy membership functions of S_p and F_p, or S_n and F_n respectively, can be manipulated directly. Doing so it is possible to either define: (a) fuzzy intervals; (b) fuzzy values; or (c) crisp values (which are all subtypes of fuzzy numbers).

As an example, the expert may wish to define the strength of confirmation of an 'increased serum glucose level' for 'diabetes' as being somewhere in the fuzzy interval [0.45, 0.85] and, therefore, an interval representation is chosen

$$F_p(u) = \begin{cases} 1 & u = 1, \\ 0 & u \neq 1, \end{cases}$$

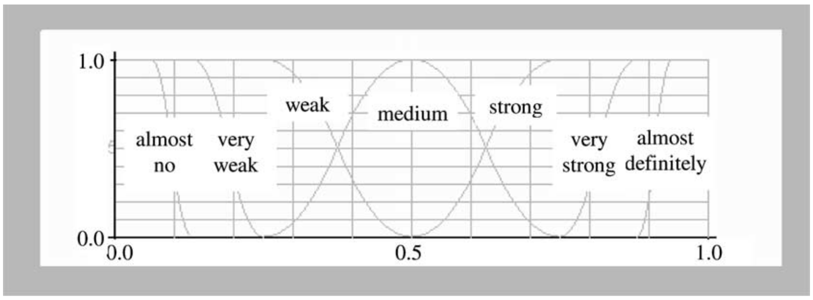

Fig. 9. Compatibility functions of the linguistic strength of confirmation and strength of exclusion variables.

and

$$S_{\mathrm{p}}(u) = \Pi(u; 0.45, 0.55, 0.75, 0.85).$$

If the expert is more confident of his judgment, S_{p} might be defined as being approximately 0.65 and the corresponding membership functions as

$$F_{\mathrm{p}}(u) = \begin{cases} 1 & u = 1, \\ 0 & u \neq 1, \end{cases}$$

and

$$S_{\mathrm{p}}(u) = \Pi(u; 0.55, 0.65, 0.65, 0.75).$$

Finally, if the exact values of the strengths of the relationships can be given they can be specified by crisp values which here can be seen as a special case of a fuzzy number. Accordingly, S_{p} can be characterized by a fuzzy membership function

$$F_{\mathrm{p}}(u) = \begin{cases} 1 & u = 1, \\ 0 & u \neq 1, \end{cases}$$

and

$$S_{\mathrm{p}}(u) = \begin{cases} 1 & u = 0.65, \\ 0 & u \neq 0.65. \end{cases}$$

Formally, even these values remain defined as fuzzy values.

4.7. Semiautomatic calculation of fuzzy relationships

The difficulty to judge and assess fuzzy compatibility functions, even if supported by a stepwise refinement procedure as outlined above, calls for additional support. Medical inferences are often guided by further knowledge from base rates and probabilities (e.g. incidence and prevalence) which are scientifically established for various types of patient populations and medical conditions.

Furthermore, in-depth knowledge of the local patient population, well-researched patient samples or hypothetical cases can be used during knowledge acquisition to serve as ‘gold standards’ or at least as robustness indicators to evaluate changes in the knowledge base.

We therefore offer a technique that allows a semiautomatic acquisition of fuzzy relationships [3]. The medical knowledge acquired by using this technique is statistical in nature and may, in some cases, differ from the judgmental knowledge of an expert. Thus, we refer to it as a semiautomatic knowledge acquisition technique assuming that the experts critically review the achieved results.

The semiautomatic knowledge acquisition technique is based on the calculation of frequencies of co-occurrence of medical concepts in a patient data base. With respect to the different types of knowledge representation, we distinguish between two different computational models. If patient data is stored in a conventional database the vagueness and fuzziness of medical concepts is not taken into account and a medical concept is either present in the patient's medical record or not. If, on the contrary, data are obtained from a fuzzy database (as the MedFrame/CADIAG-IV database), medical concepts may not only be present or absent, but also be present to a certain degree. In that case, the calculation is performed by using Sigma-Counts [5].

The proposed semiautomatic knowledge acquisition technique was first tested in CADIAG-II [3]. For this purpose, a set of batch programs that were implemented on an IBM 4341 model 2 were applied on large patient databases within the hospital information system of the Vienna General Hospital. As the knowledge representation formalism of CADIAG-II has been extended in MedFrame/CADIAG-IV we had to adapt the computational model.

5. The Knowledge Base Builder Toolkit

We have developed a prototype of a knowledge acquisition software, the Knowledge Base Builder Toolkit, that adopts the conceptual knowledge acquisition model as described above. The software has been programmed in the Java programming language (using JDK 1.1) and was designed as an Internet-based client in the MedFrame/CADIAG-IV client–server environment (Fig. 10).

The Knowledge Base Builder Toolkit has been implemented to allow a user-centered acquisition of knowledge and to make easier and more accessible the three main knowledge acquisition tasks of: (a) defining medical concepts; (b) providing appropriate interpretations for patient data; and (c) constructing inferential knowledge.

For this purpose, the Knowledge Base Builder Toolkit comprises of a Thesaurus that facilitates the definition, administration, and retrieval of all instances of medical concepts and thus supports the building of the definitional hierarchy of the controlled MedFrame/CADIAG-IV vocabulary.

A Data-to-Entity Conversion Rule Builder supports the process of defining the conversion rules with several assistants and a Rule Builder Assistant facilitates the definition, syntax checking, and maintenance of rules which usually are composed of user-defined medical entities and a system-defined set of operands (e.g. Boolean and fuzzy operators).

The stepwise refinement of fuzzy medical relationships is supported by a corresponding sequence of dialogs. If an expert has acquired some familiarity with all the available features, direct interaction with the appropriate step is possible without a need to go through all previous steps. In the user interface, textual definitions (i.e. function type, values, bounds, and ranges) as well as the corresponding graphical representations (i.e. the membership functions) can be manipulated to define or adapt fuzzy members.

Finally, a Profile Editor allows the domain expert to build disease profiles and to immediately see the consequences of any changes with respect to a predefined case base. The Profile Editor employs a spreadsheet metaphor to allow an intuitive performance of this task. A new disease profile may either be established from scratch or be based upon one of two predefined knowledge resources, either another already defined disease profile, or a complete knowledge base module.

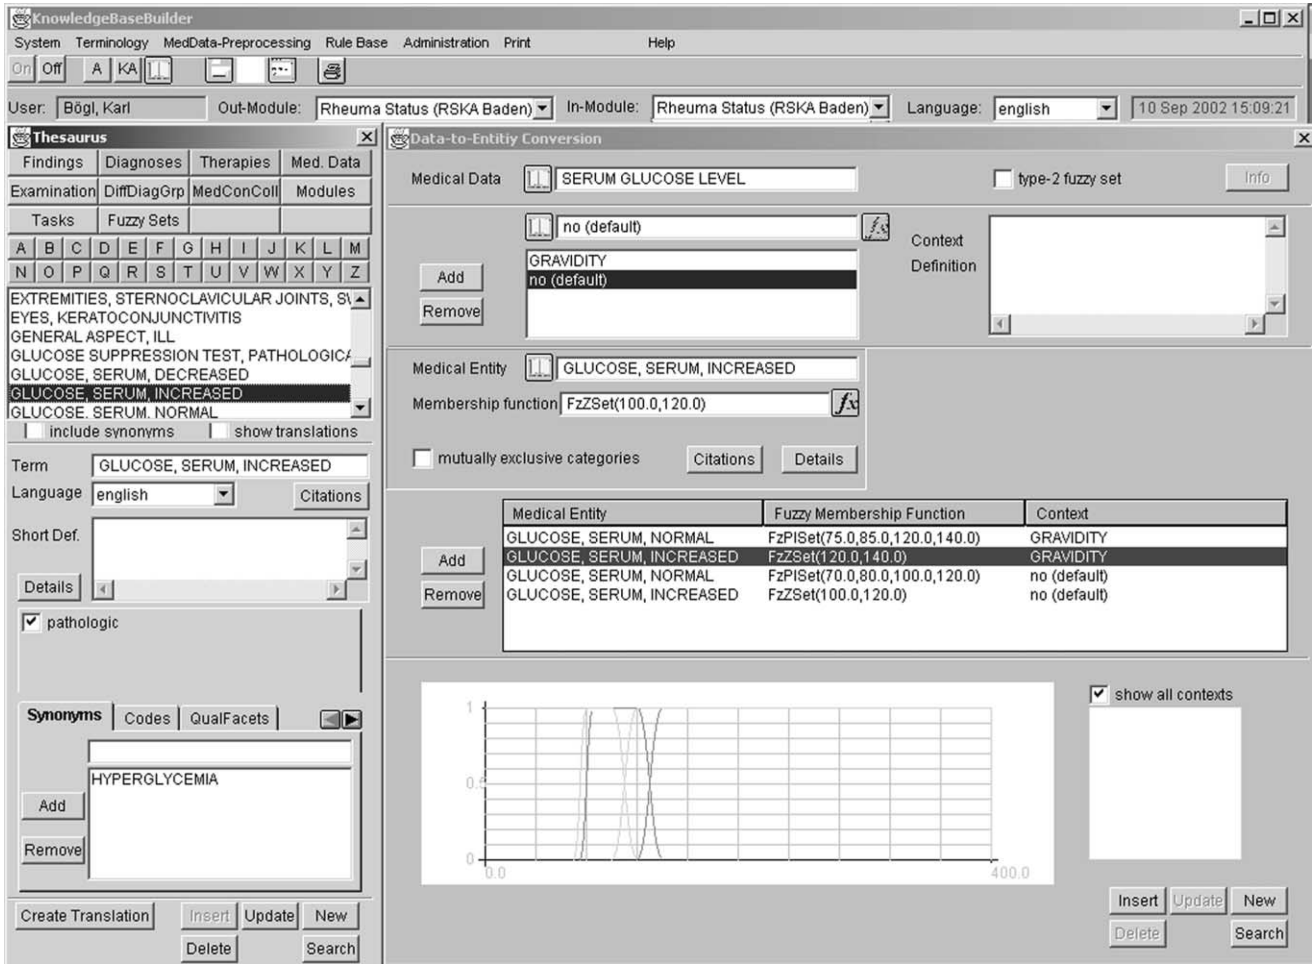

Fig. 10. Case example illustrating the definition of context-dependent data-to-entity rules using the Knowledge Base Builder Toolkit.

6. Discussion

It is premature to evaluate the Knowledge Base Builder Toolkit software because the crucial test—successful everyday use in a clinical setting—has not yet been performed. However, early experiences with the porting of one of the CADIAG-II knowledge bases (rheumatology) and with the construction of several new knowledge bases are promising.

The most ambitious project so far has been the construction of a hepatologic knowledge base, that comprises 120 different liver diseases and more than 900 findings, lab tests and clinical findings. It turned out that the modified knowledge representations and acquisition procedures were manageable and did not confuse experts or knowledge engineers once they were accustomed to the basic concepts. However, given the fact that reasoning with negative evidence is not easily understood even at the formal (logical, mathematical, philosophical) level, it should come as no surprise that difficulties in using the full spectrum for negative evidence representations have been reported. A second application has been skeletal radiology with special emphasis on radiological manifestations of rheumatologic diseases [14].

The balance between providing more flexibility and trying to reduce judgmental efforts seems still difficult to maintain. While our framework tries to minimize complexity with predefined acquisition steps and supportive tools, it remains difficult to assess the impact of decisions on the rest of the system. Additional features in our model, such as the ability to classify and constrain knowledge pieces to particular uses may alleviate 'collision' effects. For example, most definitions are restricted to a single medical subfield or confined to users of a particular hospital department where local standards may overrule generally accepted knowledge. However, already relatively simple cases of multi-criteria decision making (e.g. patho-physiologically interacting entities, time-dependent entities, disease profiles) are difficult to define.

By introducing intuitive methods for a stepwise definition of fuzzy medical entities and fuzzy relationships we allow the medical experts to establish medical knowledge bases with minimum additional support from a knowledge engineer. Results gained with the Knowledge Base Builder Toolkit show that an acquisition of domain knowledge can be achieved and that the process provides now a transparent—and easier to maintain—interface to the knowledge base. However, the construction of useful medical knowledge bases remains a time-consuming and demanding task. The methods and tools provided in our system are designed to make the knowledge acquisition task easier, but they still require familiarity of the domain experts with the methodology as described above. However, the explicit definition of the required steps and the transparent supporting representations and computer interfaces certainly enhance the cooperation and mutual understanding between experts and knowledge engineers.

Recently knowledge acquisition from very large databases, including medical databases and DNA sequence databases, called 'rule extraction' or 'rule generation' has been attracting much interest in the artificial intelligence and artificial neural network community [19]. Very recently, d'Avila Garcez et al. have published an approach to symbolic knowledge acquisition by rule extraction [10]. This paper is dealing with a fundamental knowledge base refinement approach. It is very likely that approaches like these will substantially ease the knowledge acquisition bottleneck in the future.

The fuzzy knowledge representation framework of MedFrame/CADIAG-IV allows medical knowledge acquisition at varying levels of precision and certainty. Thus, the expert may define relationships between medical entities very vaguely or unambiguously. The fuzzy inference mechanism of MedFrame/CADIAG-IV is capable of dealing with both highly precise as well as incomplete, uncertain, and vague knowledge. Thus, even knowledge that would not have been incorporated into non-fuzzy knowledge-based systems may contribute to the system's output.

In summary, the conception and the design of our system reflect the need for a user-centered, intuitive, and easy-to-handle tool for establishing and maintaining fuzzy medical knowledge bases. First results with an early prototype have shown that this critical phase of knowledge-based system development, which is of special importance in the complex mathematical context of applied fuzzy set theory, can now be passed more easily.

Acknowledgements

We gratefully acknowledge the financial support provided by ARC Seibersdorf Research GmbH, Biomedical Engineering, Seibersdorf, Austria, the corporate research partner of the Ludwig Boltzmann Institute for Expert Systems and Quality Management in Medicine, Vienna, Austria.

References

[1] Adlassnig K-P. A fuzzy logical model of computer-assisted medical diagnosis. Methods Inform Med 1980;19:141–8.

[2] Adlassnig K-P. Fuzzy set theory in medical diagnosis. IEEE Trans Syst Man Cybernet 1986;16:260–5.

[3] Adlassnig K-P, Kolarz G. Representation and semiautomatic acquisition of medical knowledge in CADIAG-1 and CADIAG-2. Comput Biomed Res 1986;19:63–79.

[4] Adlassnig K-P, Akhavan-Heidari M. CADIAG-2/GALL: an experimental diagnostic expert system for gallbladder and biliary tract diseases. Artif Intell Med 1989;1:71–7.

[5] Adlassnig K-P, Kolarz G, Scheithauer W, Grabner H. Approach to a hospital-based application of a medical expert system. Med Inform 1986;11:205–23.

[6] Adlassnig K-P, Kolarz G, Scheithauer W, Effenberger H, Grabner G. CADIAG: approaches to computer-assisted medical diagnosis. Comput Biol Med 1985;15:315–35.

[7] Boegl K. Design and implementation of a web-based knowledge acquisition toolkit for medical expert consultation systems. Doctoral thesis, Technical University of Vienna, Austria; 1997.

[8] Boegl K, Leitich H, Kolousek G, Rothenfluh TE, Adlassnig K-P. Clinical data interpretation in MedFrame/CADIAG-IV using fuzzy sets. Biomed Eng Appl Basis Commun 1996;8(6):488–95.

[9] Boose JH. Knowledge acquisition. In: Shapiro SS, editor. Encyclopedia of artificial intelligence, vol. 1, 2nd ed. New York: Wiley; 1992.

[10] d'Avila Garcez AS, Broda K, Gabbay DM. Symbolic knowledge extraction from trained neural networks: a sound approach. Artif Intell 2001;125:155–207.

[11] Dubois D, Prade H. Fuzzy sets and systems. New York: Academic Press; 1980.

[12] Hayashi Y, Nakai M. Efficient method for multi-dimensional fuzzy reasoning and its application to fault diagnosis. In: Proceedings of the International Workshop on Artificial Intelligence for Industrial Applications, Hitachi, Japan, 1988. p. 27–32.

[13] Hayashi Y, Tazaki E, Yoshida K, Dey P. Medical diagnosis using simplified multi-dimensional fuzzy reasoning. In: Proceedings of IEEE 1988 International Conference on Systems, Man and Cybernetics, Beijing, China, 1988. p. 58–62.

[14] Kainberger F, Gahleitner A, Steingruber I, Bögl K, Kettenbach J, Kolarz G, et al. Delphi-expert—a knowledge acquisition tool for expert systems in diagnostic radiology. In: Lemke HU, Vannier MW, Inamura K, editors. Computer assisted radiology and surgery: CAR'97. Amsterdam: Elsevier, 1996, p. 569–74.
[15] Kolousek G, The system architecture of an integrated medical consultation system and its implementation based on fuzzy technology. Doctoral thesis, Technical University of Vienna, Austria, 1997.
[16] Leitich H, Kiener HP, Kolarz G, Schuh Ch, Graninger W, Adlassnig K-P. A prospective evaluation of the medical consultation system CADIAG-II/RHEUMA in a rheumatological outpatient clinic. Methods Inform Med 2001;40:213–20.
[17] Leitich H, Boegl K, Kolousek G, Rothenfluh TE, Adlassnig K-P. A fuzzy model of data interpretation for the medical expert system MedFrame/CADIAG-IV. In: Trappl R, editor. Proceedings of Cybernetics and Systems'96, Austrian Society of Cybernetic Studies, Schottengasse 3, A-1010 Vienna, 1996. p. 300–3.
[18] Miller RA, Masarie FE. Quick Medical Reference (QMR): an evolving, microcomputer-based diagnostic decision support program for general internal medicine. In: Proceedings of the Annual Symposium on Computer Applications in Medical Care. Washington (DC): IEEE Press, 1989. p. 947–8.
[19] Mitra S, Hayashi Y. Neuro-fuzzy rule generation: survey in soft computing framework. IEEE Trans Neural Networks 2000;11(3):748–68.
[20] Musen MA. An overview of knowledge acquisition. In: David JM, Krivine JP, Simmons R, editors. Second generation expert systems. Berlin: Springer; 1993. p. 415–38.
[21] Pearl J. Probabilistic reasoning in intelligent systems: networks of plausible inference. Menlo Park: Morgan Kaufmann; 1988.
[22] Rothwell DJ. SNOMED-based knowledge representation. Methods Inform Med 1995;34:209–13.
[23] Shortliffe EH. Computer-based medical consultations: MYCIN. New York: Elsevier; 1976.
[24] Steimann F. Diagnostic monitoring of clinical time series. Doctoral thesis, Technical University of Vienna, Austria, 1995.
[25] Steimann F. The interpretation of time-varying data with DIAMON-1. Artif Intell Med 1996;8:343–57.
[26] van Bemmel J, Musen MA. Handbook of medical informatics. Berlin: Springer; 1998.
[27] Warner HR, Sorenson D, Bouhaddou O. Knowledge engineering in health informatics. New York: Springer Verlag; 1997.
[28] Zadeh LA. Fuzzy sets. Inform Control 1965;8:338–53.
[29] Zadeh LA. Linguistic variables, approximate reasoning and dispositions. Med Inform 1983;8:173–86.