

Representation and Acquisition of Knowledge for a Fuzzy Medical Consultation System

Thomas E. Rothenfluh¹, Karl Bögl², and Klaus-Peter Adlassnig²

¹ Department of Psychology, University of Zurich, Zürichbergstr. 43, CH-8044 Zurich, Switzerland

² Department of Medical Computer Sciences, University of Vienna Medical School, Spitalgasse 23, A-1090 Wien, Austria

Email: rothen@methpsy.unizh.ch, {k.b,p.a}@trulli.imc.akh-wien.ac.at

Abstract. This paper describes the knowledge representations that are used in MedFrame/CADIAG-IV, a medical computer consultation system. Similar to its predecessor system, CADIAG-2 [1], 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 engines. Because the elicitation and acquisition of the required knowledge is a difficult and time-consuming task (even more so when unfamiliar representations like fuzzy membership functions are to be acquired), MedFrame/CADIAG-IV has been designed to provide better support for knowledge engineers and domain experts to define fuzzy knowledge concepts as well as fuzzy inference rules. Knowledge acquisition procedures and computer tools have been implemented in order to make the three main tasks of (a) defining medical concepts, (b) providing appropriate interpretations for patient data, and (c) constructing inferential knowledge easier and more accessible. The paper discusses the motivations and rationale for some system design and data modeling decisions and explains how the main tasks are supported both by special representations and by a stepwise knowledge acquisition process.

Keywords. Fuzzy medical consultation systems, knowledge acquisition

1 Introduction

One of the crucial tasks in today's software engineering process 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 idiosyncratic, complex and specialized uses. These requirements are especially true for the complex field of medical diagnosis and therapy planning in a wide domain such as internal medicine which is the major application domain of MedFrame/CADIAG-IV.

Medical knowledge, especially the precise nature of the relationships between laboratory data, medical signs, findings, and diagnostic hypotheses can be characterized as a (sometimes huge, sometimes sparse) 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 [16], Bayesian inference models in ILIAD [19] belief networks [e.g., 8] or QMR's frequency weights and evoking strengths [7; 14] are prominent examples of different approaches to capture some of that uncertainty. CADIAG-2, a 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 [2].

Complex classification systems (e.g., UMLS, SNOMED, ICD-10) have been developed to ease the representational task, that is, to allow for efficient and precise communication about medical knowledge. Many efforts have been made to use such classifications as a cornerstone for the knowledge acquisition task, that is, for the step where human knowledge has to be "translated" into computer-usable representations. This complicated and error-prone step can either be achieved by directly translating the experts' natural language into computer representations [e.g., GALEN, 15] or by supporting the acquisition steps with guidelines [e.g., to cope with SNOMED, see 5], methods and tools that are closely modeled after the experts' explanation or reasoning schemes.

Most computer-supported diagnostic tools strive to implement "best knowledge practices" and thus require the expert to state the relevant knowledge as concise and logically correct as possible. Unfortunately, in fields like internal medicine, there is only few "proven" knowledge to acquire: physicians usually feel uncomfortable to add their "insights" and "useful associations" as hard rules, even when allowed to add some uncertainty to their inferences. MedFrame/CADIAG-IV tries to smoothen that 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 applicable for computer use. Many types of uncertainty in data interpretations and the difficulty to define logically "definitive" or "true" inferences are accounted for in non-computerized situations with the use of imprecise language. The idea to use "computations" even for fuzzy linguistic concepts was of course one of the motivations [20] to establish the scientific study of fuzzy reasoning. Although medical scientific literature is full with fuzzy statements and qualifications, there is usually no simple translation from experts' linguistic statements of uncertainty into computer representations—additional support is needed to allow physicians to define and refine their expert knowledge.

MedFrame/CADIAG-IV is a *fuzzy* medical consultation system in that it uses fuzzy representations for almost all representations of knowledge: it accepts (or

even fuzzifies) its inputs, operates on fuzzy sets with fuzzy rules, and produces fuzzy sets as output. Whenever appropriate or desired, the system can defuzzify or approximate its statements to crisp values or defined sets with the help of user-definable thresholds and specific, domain-dependent algorithms. For example, rank-ordered lists of confirmed, possible, and excluded diagnoses can be produced in order 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. Clinicians find computerized systems more useful if they receive support for their own, personal 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 the 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 and 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 otherwise might 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.

MedFrame/CADIAG-IV extends previous implementations in many ways. For example, it marks the transition from a system on a centralized host with a text-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-base implementation, and inference processes can be found in Kolousek [9]. The conceptual models and the implementational details of MedFrame/CADIAG-IV's knowledge acquisition system and the Knowledge Base Builder Toolkit (KBBT) are described in [3].

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-2 – 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 in order 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 processes [10].

In this paper, we describe the knowledge representation and the knowledge acquisition procedures that support 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 symptoms can be acquired and defined. Next, we explain the different steps and tools that are available to the user to define fuzzy relationships between different medical entities. These two steps – the definition of all elements and their relationships – are the prerequisite for MedFrame/CADIAG-IV's inference processes (not described in this paper) which use real patient data to propose diagnostic and therapeutic hypotheses. Finally, we discuss some issues that require further research.

2 Basic Fuzzy Representations

MedFrame/CADIAG-IV uses *fuzzy sets* that define the degree of membership for an element in a set [for definitions of fuzzy terminology see 6]. *Fuzzy numbers* are used to specify fuzzy membership functions $\mu(x; \alpha, \beta, \gamma, \delta)$ and are constrained to values in $[0, 1]$. Additionally, a special value representing 'unknown' is provided. This value is a way of distinguishing an element that is known to have some, but (yet) unknown, membership relation to a set from other elements whose membership degree has not yet been assessed at all.

Membership functions can be represented (and acquired, as will be explained in a following section) through numerical or graphical representations (0). Transition functions between endpoints (between α and β or between γ and δ) can accept linear or other shapes.

Type 2 fuzzy sets are fuzzy sets whose degrees of memberships are themselves fuzzy sets. They are used whenever combinations of fuzzy elements are used (see

the examples in the section on data-to-entity conversions). Formally, the result of such combinations are fuzzy power sets.

Fuzzy relations between two (ordinary) sets are defined as the fuzzy set of the Cartesian product between the elements of those sets. Every element in this set is characterized by a membership function. Since fuzzy numbers are used to define the basic sets themselves (see section '4 Stepwise Acquisition of Inferential Knowledge'), fuzzy relations in MedFrame/CADIAG-IV's knowledge base are essentially a combination of type 2 fuzzy sets.




Fig.1. Graphical representation of the fuzzy membership function $\mu(x; \alpha, \beta, \gamma, \delta)$.

3 Knowledge Acquisition for Fuzzy Concepts

For all diagnostic 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, we distinguish between two basic knowledge types that are defined in the knowledge representation framework [9]:

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. With the help of a facet and qualifier system, which is similar and compatible with medical classification systems such as SNOMED, all entities can be defined and identified in a stringent, coherent, and semantically meaningful way.

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). Since MedFrame/CADIAG-IV reasoning mechanisms operate at the level of symbolic concepts (i.e., medical entities), a data-to-entity conversion has to be employed to transform quantitative medical data into medical entities.

3.1 Data-to-Entity Conversion

The transformation of medical data into meaningful interpretation categories (medical entities) can be compared to the definition of a linguistic variable in other fuzzy systems. This definition has to be established in the knowledge acquisition phase for all meaningful data values of a medical parameter. At run-time, when actual patient data is used, these definitions will translate data values (e.g., white blood cell count) and assessments (e.g., morning stiffness last more than 15 minutes but less than 30 minutes) into symbolic, but fuzzy medical entities.

The representation of medical data includes information about lower/upper bounds of acceptable values and convertible measurement units. Then, in the first step of a data-to-entity conversion, the defined range of possible values needs to be partitioned into an appropriate number of categories. These categories usually define some "normal" category and any number of "abnormal" or "pathologic" categories. Depending on the nature of the medical data, just one category, simple dichotomies, or one- and two-sided ordinal categories are used. For example, a patient's temperature reading could be classified just into one categories "fever" (implicitly, this also defines "not fever") or into finer-grained categories such as "low temperature", "normal temperature", "high temperature", and "fever." The number of appropriate categories depends on the required level of expressiveness. The more categories are established the more decisions and maintenance efforts are required to make them useful.

In a second step, the selected categories can be defined as *exclusive* or *inclusive* categories. 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 for diabetes, but they also allow for an improved diagnosis of stages of diabetes).

In a last step, the compatibility functions for the interpretation categories need to be defined. In MedFrame/CADIAG-IV, a data-to-entity conversion rule builder supports this process with several assistants [3]. For a given medical data value (represented on the x-axis in 0), the degree of compatibility (y-axis) to the selected 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). Different fuzzy membership functions can be used to model the desired fuzziness characteristics (α , δ , and transition functions in 0) 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.




Fig.2. Compatibility functions for medical data values of “total serum glucose level”. Part (a) represents exclusive interpretation categories, whereas part (b) the case represents where the category “heavily increased” includes the interpretation “increased” [adapted from 10].

3.2 Multi-Dimensional Data

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’, for details see [4], 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.

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 data is important, but also the data's variation and course 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 [12; 17]. During knowledge acquisition, the user can select from a predefined set of fuzzy time-trends (constant, rising, falling, oscillating) with special parameters (onset-time, onset-value, direction) which cover many typical situations.

3.3 Context-Dependent Data

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.

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, an unlimited 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 gravidity 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 a gravidity context, low glucose levels have other membership functions as normal or high glucose levels). The selection of appropriate contexts and the computation of compatibility function for the selected context(s) are described in more detail in [4].

It has become a standard requirement of knowledge representations that they are transparent and meaningful to the people maintaining them and that all "representational tricks" to use side-effects will eventually lead to the system's death. By separating out most of the difficult data-related decisions (e.g. "does a temperature of 37.1 degrees 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.

4 Stepwise Acquisition of Inferential Knowledge

At the core of MedFrame/CADIAG-IV's inference processes are reasoning mechanisms that deal with symbolic medical entities. These entities are connected by means of relations and the basic inferencing follows these relations and

recursively calculates (fuzzy) values for connected entities. Two special knowledge representations, *disease profiles* and *explicit rules*, combine several medical entities in more complex ways.

Rules are standard representations found in many other medical consultation systems and are mainly used in MedFrame/CADIAG-IV to express complex diagnostic rules. Simple *definition rules* allow the expression of medical knowledge based on logical and/or combinations of defined medical entities. For example, the confirmation of the finding ‘symmetrical arthritis’ requires the assessment of a quotient of symmetrically involved joints divided by the total number of involved joints. More elaborate rules are able to express common diagnostic concepts like “If three out of the following nine findings are present, then conclude X” with the help of predefined and user-definable (mainly mathematical) operators and functions. However, rules are not restricted to logical combinations (AND, OR, NOT, etc.) but can be formulated as fuzzy rules as well. Even well established medical criteria, such as the ACR/ARA criteria for rheumatic diseases, which are formulated as a set of rules, can benefit from fuzzification [11]. A *rule builder* has been implemented in MedFrame/CADIAG-IV’s KBBT [3] that 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., arithmetic, Boolean, fuzzy, and magnitude operators).

Another implicit type of “rules” work in the background of MedFrame/CADIAG-IV: the definitional hierarchy of the *controlled medical vocabulary*. Based on other multi-axial medical classification systems like SNOMED, a facet system is used that allows the definition of networks of qualified terms. For example, a medical concept like “*acute bacterial inflammation of the right lower lobe of the lung*” is defined along topographical and morphological axes. These facet systems (modules in SNOMED) have hierarchical structures, whose paths allow logical inferences like abstractions and generalizations. In similar ways, synonyms and language variants – which can be specified to provide better adaptability and customizability of the vocabulary – allow implicit inferences and reasoning steps which are not expressed in rules.

In extension to explicitly formulated rules, *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, findings, examinations, syndromes, diseases, therapies) that are related to another entity (usually a disease or diagnostic hypothesis). The relation between two entities is a *fuzzy membership relation* and is defined with the help of a stepwise refinement process which is outlined in detail below. The most common and most descriptive use of fuzzy relationships is the connection of symptoms with diseases, but the same relations can be used for any number and types of medical entities. Thus, medical entities that act as symptoms in a disease profile can themselves be complex fuzzy concepts that are composed of many combinations

of other “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 relations discussed below.

4.1 Stepwise Refinement of Fuzzy Relations

Given the complexity of the network of entities that is to be defined and given the difficulty in assessing even a simple one-to-one relation between a finding and a diagnosis, a *guided, stepwise knowledge-acquisition process* has been established. In a few steps, the connection between two entities can be established by medical experts starting from a simple association and ending with complex fuzzy membership functions. These steps are optional refinements that are supported by a corresponding sequence of dialogs in the system’s graphical user interface. 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. 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 relations 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 relation in the following steps.

4.2 Step 1: Associations

To add knowledge about the relation 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 relations or at least empirical correlations are accepted as scientific facts. For example, a positive relation between a symptom and a disease implies that medical knowledge is available to always infer (confirm) the presence of a disease whenever the symptom is present. To exclude the disease whenever the symptom is present, a negative association would have been used.

A neutral association adds weak knowledge that some kind of relation exists which is different from definitive exclusion and thus a necessary association. For example, if a symptom is present, the disease *may* be present, but the symptom is *not* required to be present. Neutral associations are also used to connect purely descriptive, non-inferential knowledge to other entities (and thus can be retrieved from those entities for explanation purposes).

Because definitive positive or negative evidence is hard to find, these associations can be refined in the following steps to reflect the uncertainty or lack of

established causal knowledge in order to account for valuable experiential and casuistic knowledge.

4.3 Step 2: Relations

The basic relation concepts of CADIAG-2, which were useful for acquiring medical knowledge from domain experts as well as for successful application of the fuzzy reasoning system, are “*frequency of occurrence*” and “*strength of confirmation*.” They are used to differentiate positive and negative associations into nine basic relations (MedFrame/CADIAG-IV slightly extends CADIAG-2 by introducing explicit calculation of all combinations of negative evidence).

Any relation is directed from an antecedent A to a consequent D (usually from a 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 parameters can be crisp judgements (e.g. obligatory present AND confirming) 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. obligatory present AND NOT confirming).

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 (obligatory for $\neg D$) and the strength of confirmation S_n (excluding or not) are defined as well. To be less formal, S_n is more intuitively thought of as “strength of exclusion” and F_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 & S_p , or F_n & S_n , respectively.

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.

4.4 Step 3: Fuzzy Memberships with Linguistic Terms

It is certainly possible to adopt frequentistic or probabilistic interpretations of the relationships introduced above – and in fact, many experts use just that kind of knowledge to “fuzzify” their relations. MedFrame/CADIAG-IV tries to support another approach to define uncertainty and vagueness by providing different explicit ways to define and adapt fuzzy membership functions by specializing on the concepts introduced in step 2.

Based on the relations to be refined, intelligent assistants propose appropriate tools (note that not all relations can be refined). For example, if F_p or F_n were established in the previous step not to be 1 (obligatory occurring) they can be further refined with linguistic terms. For frequency of occurrence they denote the

concepts “almost never”, “very seldom”, “seldom”, “medium”, “often”, “very often”, and “almost always.” Similarly, for S_p and S_n , the linguistic terms “almost no”, “very weak”, “weak”, “medium”, “strong”, “very strong”, and “almost definitely” can be used to modify the strength of confirmation. Associated with these terms are predefined fuzzy membership functions.




Fig. 3. Illustration of the stepwise refinement process of the relationship between two medical entities, INCREASED SERUM GLUCOSE LEVEL and DIABETES. F_p , F_n , S_p , and S_n are defined in the text

4.5 Step 4: Manipulating Membership Functions

Alternatively, or as a further refinement to the use of linguistic terms, the fuzzy membership functions can be manipulated directly. 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 graph) can be manipulated to define or adapt fuzzy intervals or fuzzy values.

The summarizing example in 0 does stop with this refinement. However, if enough knowledge becomes available to provide exact values for frequency of occurrence or for the strength of confirmation, the fuzzy functions can be converted to values. Formally, even these values remain defined as fuzzy values.

4.6 Semi-Automatic Acquisition and Evaluation

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 or 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. MedFrame/CADIAG-IV's profile editor, for example, allows a user to immediately see the consequences of any changes to values for frequency of occurrence or strength of confirmation in a disease profile with respect to a predefined case-base.

As a possible solution, fuzzy membership functions for F_P , F_N , S_P , and S_N can be calculated based on reference patient databases. The calculation of the membership function can either be based on frequencies (in the case of crisp entities, e.g., if a finding is present or absent) or on Sigma-Counts (as in MedFrame/CADIAG-IV's case of fuzzy entities, e.g., a finding is normal with $[[[x, 5, 6, 12, 13]]]$). The detailed procedures and formulas for those calculations are outlined elsewhere [10]. It is important to note here that this semi-automatic acquisition of membership functions is not a panacea. Different assumptions about the influence of prevalences (i.e., the frequency with which symptoms or diseases are present in the patient population) and about the nature of the patient samples (e.g., what is the interpretation of patients not having the disease) influence the validity of the calculations. Thus, these statistically derived values will always need critical review by experts; a view that is shared by others who are trying to use cases to infer fuzzy rules [13].

5 Discussion

It is premature to evaluate MedFrame/CADIAG-IV's knowledge acquisition process 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 CADIAG-2 knowledge bases (rheumatology), with the construction of a new knowledge base for hepatology (with 120 different diseases and about 400 other medical data and entities), and with a radiologic application are promising.

Of course, because the enhanced inference model is backward compatible with CADIAG-2, performance measurements with knowledge bases and patient data from that system are merely a means of checking the new implementation. 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.

So far, we have only in-house experience with newly acquired knowledge bases with respect to usage and acceptance of the new features. 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 (e.g., F_N , S_N) are reported. Further investigations into methodical aspects such as reasoning with negative evidence, consistency checking, and semi-automatic construction of fuzzy relations from patient data are under way.

For example, the use of *linguistic variables* is not without problems. Although the respective membership functions can be empirically validated and their robustness can be checked with simulation studies, the interpretation is still subjective and based on accepted conventions. This is usually not a problem within a specialized domain of expertise where some agreement about these terms has been established long before computer systems tried to model decisions. If, however, multiple domains are aggregated in such a system – as is the case for MedFrame/CADIAG-IV's ambitious goal to integrate many subfields of internal medicine – the interpretation of those terms might become more problematic. Furthermore, any modification in the set of a linguistic variable itself (e.g., adding some new fuzzy quantifier) may require adaptations of previously entered knowledge [13]. By providing even finer-grained access to the membership functions, MedFrame/CADIAG-IV is not restricted to the use of linguistic terms – but at the price of imposing additional decision tasks on the domain expert.

The balance between providing more flexibility and trying to reduce judgmental effort seems still difficult to maintain. While MedFrame/CADIAG-IV 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 MedFrame/CADIAG-IV, 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. Furthermore, although many of the desired reasoning properties “emerge” once all inference mechanisms are executed, the local representation may be very different from what domain experts think it should be. To ease the comprehension of the effects which a decision in the knowledge acquisition might have, empirical studies are needed that lead to aggregation and reasoning functions which are closer modeled after human decision making processes [e.g., 21]. This will not only help to improve the system detailed design but it will also lessen the decisional burden on the domain experts because the required precision and granularity can be assessed.

We tried to improve MedFrame/CADIAG-IV with respect to its support for the knowledge acquisition process. Knowledge acquisition is one of the tasks (sometime called a bottleneck) that are crucial for the success of any system that embodies domain knowledge at expert levels – successful computational and implementations features will not suffice [see also 18, for a recent overview of the field]. Preliminary results show that the acquisition of domain knowledge can be achieved quite easily 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.

References

1. Adlassnig, K.-P. (1980). A fuzzy logical model of computer-assisted medical diagnosis. *Methods of Information in Medicine*, 19, 141-148.
2. Adlassnig, K.-P. & Kolarz, G. (1986). Representation and semiautomatic acquisition of medical knowledge in CADIAG-1 and CADIAG-2. *Computers and Biomedical Research*, 19, 63-79.
3. Bögl, K. (1997). *Design and implementation of a web-based knowledge acquisition toolkit for medical expert consultation systems*. Ph.D. Thesis, Technical University, Vienna, Austria.
4. Bögl, K., Leitich, H., Kolousek, G., Rothenfluh, T. E. & Adlassnig, K.-P. (1996). Clinical data interpretation in MedFrame/Cadiag-4 using fuzzy sets. *Biomedical Engineering: Applications, Basis & Communications*, 8(6), 488-495.
5. Cimino, J. J. & Clayton, P. D. (1994). *Coping with changing controlled vocabularies*. Proceedings of the Annual Symposium on Computer Application in Medical Care (SCAMC'94), Washington, DC.

6. Dubois, D. & Prade, H. (1980). *Fuzzy sets and systems*. New York: Academic Press.
7. Giuse, D. A., Giuse, N. B. & Miller, R. A. (1990). Towards computer-assisted maintenance of medical knowledge bases. *Artificial Intelligence in Medicine*, 2, 21-33.
8. Henrion, M., Breese, J. S. & Horvitz, E. J. (1991). Decision analysis and expert systems. *AI Magazine*, 12(4), 64-91.
9. Kolousek, G. (1997). *The systems architecture of an integrated medical consultation system and its implementation based on fuzzy technology*. Ph.D. Thesis, Technical University, Vienna, Austria.
10. Leitich, H. (1995). *Anforderungen an eine Wissenserwerbssystem für das medizinische Expertensystem Cadiag-IV*. Master Thesis, Technical University, Vienna, Austria.
11. Leitich, H., Adlassnig, K.-P. & Kolarz, G. (1996). Development and evaluation of fuzzy criteria for the diagnosis of rheumatoid arthritis. *Methods of Information in Medicine*, 35, 334-342.
12. Leitich, H., Bögl, K., Kolousek, G., Rothenfluh, T. E. & Adlassnig, K.-P. (1996). *A fuzzy model of data interpretation for the medical expert system MedFrame/CADIAG-4*. Cybernetics and Systems '96, Vienna, Austria.
13. Martinez-Béjar, R., Shiraz, H. & Compton, P. (1998). *Using ripple down rules-based systems for acquiring fuzzy domain knowledge*. Paper presented at the Knowledge Acquisition Workshop (KAW'98), Banff.
14. Miller, R. A. & Masarie, F. E. (1989). *Quick Medical Reference (QMR): An evolving, microcomputer-based diagnostic decision support program for general internal medicine*. Proceedings of the Annual Symposium on Computer Applications in Medical Care (SCAMC'89), Washington, DC.
15. Rector, A., Rossi, A., Consorti, M. & Zasnstra, P. (1998). Practical development of re-usable terminologies: GALEN-IN-USE and the GALEN Organisation. *International Journal of Medical Informatics*, 48(1-3), 71-84.
16. Shortliffe, E. H. (1976). *Computer-Based Medical Consultations: MYCIN*. New York: Elsevier.
17. Steimann, F. (1996). The interpretation of time-varying data with DIAMON-1. *Artificial Intelligence in Medicine*, 8, 343-357.
18. van Bommel, J. & Musen, M. A. (Eds.). (1998). *Handbook of Medical Informatics*. Berlin: Springer.
19. Warner, H. R., Haug, P., Bouhaddou, O., Lincoln, M., Warner, H., Sorenson, D., Williamson, J. W. & Fan, C. (1988). *ILLIAD as an expert consultant to teach differential diagnosis*. Proceedings of the Annual Symposium on Computer Applications in Medical Care (SCAMC'88), Philadelphia.
20. Zadeh, L. A. (1983). Linguistic variables, approximate reasoning and dispositions. *Medical Informatics*, 8, 173-186.
21. Zimmermann, H.-J. (1997). Operators in models of decision making. In D. Dubois, H. Prade & R. R. Yager (Eds.), *Fuzzy information engineering* (pp. 471-496). New York: Wiley.

Piotr S. Szczepaniak
Paulo J.G. Lisboa
Janusz Kacprzyk (Eds.)

Fuzzy Systems in Medicine

With 299 Figures
and 64 Tables

Physica-Verlag
A Springer-Verlag Company

Piotr S. Szczepaniak, PhD, DSc
Associate Professor
Institute of Computer Science
Technical University of Łódź
ul. Sterlinga 16/18
90-217 Łódź
Poland
office@ics.p.lodz.pl

Prof. Paulo J.G. Lisboa
School of Computing and Mathematical Sciences
Liverpool John Moores University
Liverpool L3 3AF
United Kingdom
p.j.lisboa@livjm.ac.uk

Prof. Janusz Kacprzyk
Systems Research Institute
Polish Academy of Sciences
ul. Newelska 6
01-447 Warsaw
Poland
kacprzyk@ibspan.waw.pl

ISBN 3-7908-1263-3 Physica-Verlag Heidelberg New York

Cataloging-in-Publication Data applied for
Die Deutsche Bibliothek – CIP-Einheitsaufnahme
Fuzzy systems in medicine: with 64 tables / Piotr S. Szczepaniak; Paulo J.G. Lisboa; Janusz Kacprzyk (eds.) – Heidelberg; New York: Physica-Verl., 2000
(Studies in fuzziness and soft computing; Vol. 41)
ISBN 3-7908-1263-3

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Physica-Verlag. Violations are liable for prosecution under the German Copyright Law.

© Physica-Verlag Heidelberg 2000
Printed in Germany

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Hardcover-Design: Erich Kirchner, Heidelberg

SPIN 10748155 88/2202-5 4 3 2 1 0 – Printed on acid-free paper