In *Fuzzy systems in computer science*, Kruse R, Gebhardt J, Palm R, eds, 155-169, Vieweg, Braunschweig/Wiesbaden 1994.

Fuzzy Systems in AI

Christian Freksa

Abstract

This paper reviews motivations for introducing fuzzy sets and fuzzy logic to knowledge representation in artificial intelligence. First we consider some areas of successful application of conventional approaches to system analysis. We then discuss limitations of these approaches and the reasons behind these limitations.

We introduce different levels of representation for complex systems and discuss issues of granularity and fuzziness in connection with these representation levels. We make a distinction between decomposable and integrated complex systems and discuss the relevance of this distinction for knowledge representation and reasoning. We also distinguish fuzzy relations between quantities of different granularity within one domain from fuzzy relations between two different domains and discuss the need of considering both in artificial intelligence.

We distinguish methods for describing natural, artificial, and abstract systems and contrast the modeling of system function with the modeling of system behavior in connection with the representation of fuzziness. The paper briefly discusses recent criticism of the fuzzy system approach and concludes with a prospect on soft computing in AI.

1 Why do we Need Fuzzy Sets and Fuzzy Logic in AI?

The notion of a fuzzy set [Zadeh 1965] and the development of *fuzzy set theory* and *fuzzy logic* were motivated by the severe difficulties to adequately characterize complex systems by conventional approaches of system analysis. “Adequate” means, for example, that insignificant variations on the component level of a system should not add up to significant changes on the system level. This criterion is an absolute requirement for understanding complex systems in terms of their components.

Conventional approaches represent complex systems in a reductionist manner by specifying well-defined components and their individual interactions. We will investigate the question why these approaches are of limited use in artificial intelligence and cognitive science.

The success of conventional approaches in well-defined systems

Until fairly recently, most researchers in artificial intelligence believed, AI systems could benefit from the same virtues that made much of modern science and technology a success: precise measurements, complete knowledge of the domain, and rigorous tools for dealing with them. Before discussing this belief I will briefly investigate under which conditions dealing with precise, crisp, and complete measurements has been successful.

A fundamental prerequisite for succeeding with high-precision characterizations of complex systems is the availability of basic entities and relations suitable for capturing everything that is relevant in the system under consideration. Under this condition, we can decompose complex systems into less complex subsystems and/or basic components and we can explain their role in terms of the basic components and their interactions.

Obviously, we can correctly describe complex systems in terms of their components when (1) the components conform with their definitions, (2) the interactions conform with their specifications, and (3) no other interactions interfere. These conditions certainly hold in *abstract systems* that are specified according to the three conditions given above. Examples are complex games defined by simple rules like the board games chess and go.

The three conditions also can be assumed to hold in *technical systems* whose components can be studied in isolation and whose interactions can be restricted to controlled local exchanges. Examples are closed chemical, mechanical, and electronic systems. Complex computers and computer networks constitute excellent examples that the bottom-up approach to characterizing systems on the basis of crisp notions can be highly successful.

However, the more complex these systems get, the more obvious becomes the need for developing high-level views and languages to better capture the essential aspects including actions at the relevant system level. Thus, the granularity of a description can make a big difference even if fine and coarse descriptions of crisp systems can be considered equivalent, in a formal sense.

The problem with ill-defined systems

The great success in representing complex systems in terms of their simple components that could be achieved in closed *technical systems* may be responsible for a blind belief that the same approach could be applied to other complex systems equally well. Let us consider a class of non-technical complex systems whose representation in terms of components has caused great difficulties: systems of economics or climate systems may serve as examples. Why are such systems different?

In case of an economic system, the three above requirements seem fulfilled, at first glance: (1) monetary units define the economic value of any item in basic terms, (2) mathematical rules precisely derive the effects of economic transactions, and (3) no other interactions besides transactions determine the monetary values on the component level. In fact, the strict obedience of the mathematical laws by economic transactions make it possible to precisely analyze in quantitative terms why someone became rich, became poor, went bankrupt, etc.

The problem is, that it is not very interesting to have models that only can be used for *post hoc* analysis. In the case of climate systems even the post hoc analysis does not carry very far. We would like to build models to make predictions! Why is it possible to make predictions on the component level for board games but not for economic systems?

The role of system complexity

The fact is that practically speaking it is not even possible for board games like chess or go to make reliable predictions on the component level of description. The complexity of the chess game consisting of only 32 pieces on an eight by eight checkers board already is too big for analyzing all legal moves. Nevertheless, it is possible to build useful higher-level descriptions of chess constellations from the basic entities and use these descriptions to generate reasonable predictions.

In systems of economy, in contrast, this approach has not proven practicable; condition (3) is severely violated: individual transactions in economy systems can not be predicted on the basis of local transactions; they are determined by complex global patterns and their dynamics involving psychological and other factors which cannot be captured in terms of the elementary configurations.

A similar situation is given in the case of climate systems and weather prediction: our knowledge about the preconditions in terms of fundamental facts will never be complete enough and our knowledge about the physical laws on the local level does not suffice for making useful predictions about global or local weather conditions. Thus, system complexity is only one aspect that must be considered; the availability of knowledge and the structure of that knowledge also play an import role for the representation we can use.

How does fuzziness come into the picture?

If we compare systems that are well-defined (from the bottom up) with systems that we know primarily from their global properties, we find different mapping relations between the low-level and the high-level notions. The high-level notions built up from low-level primitives typically are found in a *crisp* relation to the synthesized complex notions while the low-level notions postulated from the high-level concepts are found in a *fuzzy* relation.

For example, a taxonomy of plants and/or animals based on low-level primitives will yield a crisp classification of creatures as found in biology textbooks. On the other hand, the identification of low-level features for the definition of the high-level everyday notion “living animal” yields fuzzy relations between the high-level notion and the low-level features, as it is impossible to precisely capture the everyday notion universally by composition of low-level features. In this sense, “living animal” is a fuzzy concept, when related to low-level primitives – while on the high level on which we typically use the notion it would not be considered fuzzy at all.

2 The level of representation

Systems about whose properties we learn from global behavior can be described meaningfully on a global level, for example in the case of an economic system we might know “when the interest rate goes up, the money flow decreases”. Such a rule implicitly can take into account complex interaction regularities which cannot be captured on the component level. Although the total money flow results from individual monetary transactions, the rule does not hold for each individual transaction. Therefore it is difficult to give precise definitions of the global notions in terms of local transactions; it is much easier to identify the global effect as the net effect of local transactions and to control economy on a global level (e.g. by manipulating the interest rate) than through local transactions.

In the case of weather prediction we encounter a similar situation: on a coarse level (which may be relevant for agriculture, for example), predictions may be quite reliable, while they may be useless on the level of description on which local measurements of weather indicators (like rainfall per square meter) are made.

As a consequence, by representing complex systems not on the level of the most primitive notions but on a coarser level, these systems may become tractable. Numerous complex interactions on the local level simply can be ignored! Lotfi Zadeh recognized the importance of abstraction from low-level properties at an early stage of the artificial intelligence enterprise. He characterized intelligent systems by their ability to summarize complex descriptions by abstracting from details. In the case of everyday non-synthetic systems this requires taking into account the fuzzy relations between the high-level and the low-level features.

Integrated complex systems

There is a class of complex systems in which fuzzy relations play a particularly important role. This class consists of systems whose components cannot be studied in isolation or whose components cannot be studied in all relevant conditions. Most *natural complex systems* belong to this class. We can observe global behavior under varied condition patterns; from these we infer local influences.

Examples are biological systems which we describe in terms of presumed local functions and observed effects. Neither the description of the global effects nor the description of the local functions are suitable to capture all possible situations and to crisply represent the system. The reason is, that we are bound to use concepts which have a meaning outside of these systems (like “living animal”) since elementary local components are not available and there is no way to guarantee that these concepts precisely match the components of the described system.

Fuzzy relations between different domains

Fuzzy pattern recognition, fuzzy control, and most other successful applications of fuzzy set theory have focused on the fuzzy relation between the fine and coarse levels of representation of the artifacts involved. But from the inception of fuzzy set theory, its inventor Zadeh also suggested to represent fuzzy relations between *real* entities like physical objects and *mental* entities like concepts which are manifested in natural language expressions. In artificial intelligence, this type of relations is of particular interest.

Classical artificial intelligence had been focusing its attention on the representation structures within the medium computer. Little attention had been paid to ontological questions and the actual representation problem, i.e. the mapping between what is represented and what is representing [c.f. Palmer 1978] and to the problem of building representation structures from existing knowledge. Many AI-researchers did not consider this issue a problem; they believed, concepts could be used like nuts and bolts to screw together intelligent systems and the role of each concept would be clearly defined.

But when AI-systems grew up and moved outside their purely synthetic laboratory environments it became evident that there was a serious matching problem between natural concepts derived from the use and the behavior of systems and artificial concepts synthesized from low-level components. It became clear that we could not simply view natural concepts as imperfect entities whose objectives would be much better served by artificial substitutes.

Instead, expert system research and research in cognitive science investigated structures and properties of human knowledge in order to exploit its potential and to understand more of its function. In this process, fuzzy sets have played an important role in characterizing the relation between human concepts and natural or artificial entities and in mimicking their interactions (c.f. [Zimmermann 1992], [Dubois et al. 1993], [Kruse et al. 1994]). In this way, many of the properties of human thought and natural language, specifically with regard to their modification and their combination, could be simulated.

3 Natural systems vs. artifacts

As argued in the previous sections, the development of fuzzy set theory and fuzzy logic were influenced by properties of natural human concepts and their relation to entities in the real world. Specifically, Zadeh suggested to use fuzzy sets to represent notions like *tall* and *beautiful* in natural language phrases like *The tall dwarf is more beautiful than the small giant.* However, the present success in the application of fuzzy set theory is not so much in artificial intelligence – e.g. in the representation of natural language expressions or human concepts – but rather in control engineering, e.g. due to the improvement of purely artificial systems like household appliances, photo equipment, trains, and helicopters [Munakata, Jani 1994]. Why have fuzzy sets not caught on in artificial intelligence to the same degree as in control engineering?

The utility of fuzzy sets in artificial systems is not only due to reduced complexity as a consequence of the coarser representation of systems, but also due to properties of human concepts which support the engineering process of these systems. Fuzzy sets serve as a knowledge transfer vehicle, as a way of getting the judgment of engineers into complex systems. They are particularly suited for this task since they allow for reasonable representation relations even if the representation system is not yet completely understood. In this sense, human concepts are involved in the synthetic products of the engineers. But does this mean that fuzzy sets represent human concepts?

The characteristic function which defines a fuzzy set *characterizes* the relationship between real world entities and specific concepts (or labels typically associated with concepts); however, it does not model or explain how this fuzzy relationship comes about. For certain domains or types of tasks, a characterization of the fuzzy relationship between a concept and given instances in 'reality' is sufficient (for some tasks even this relationship is not required) - but there is an important class of problems which requires going beyond the characteristic function.

Shallow vs. deep modeling

I will argue that classical fuzzy sets represent human concepts on a rather shallow level, on the level of denoting entities within a well-defined framework, a framework in which the relevant dimensions are known, but precise values within these dimensions are missing. Human concepts, however, are not merely underdetermined physical values. They form a strong system on their own which have a meaning and make sense independent of real-world instances even though during the knowledge acquisition process they may have been derived from such instances.

In Figure 1, I present a classification of different epistemic qualities of knowledge associated with methods appropriate for processing them. On the basis of this classification I will suggest ways of extending the fuzzy set philosophy towards representing human concepts on a deeper level. The goal is to adapt notions developed in fuzzy set theory and fuzzy logic to make them better suitable for processing human knowledge and knowledge representation in artificial intelligence.

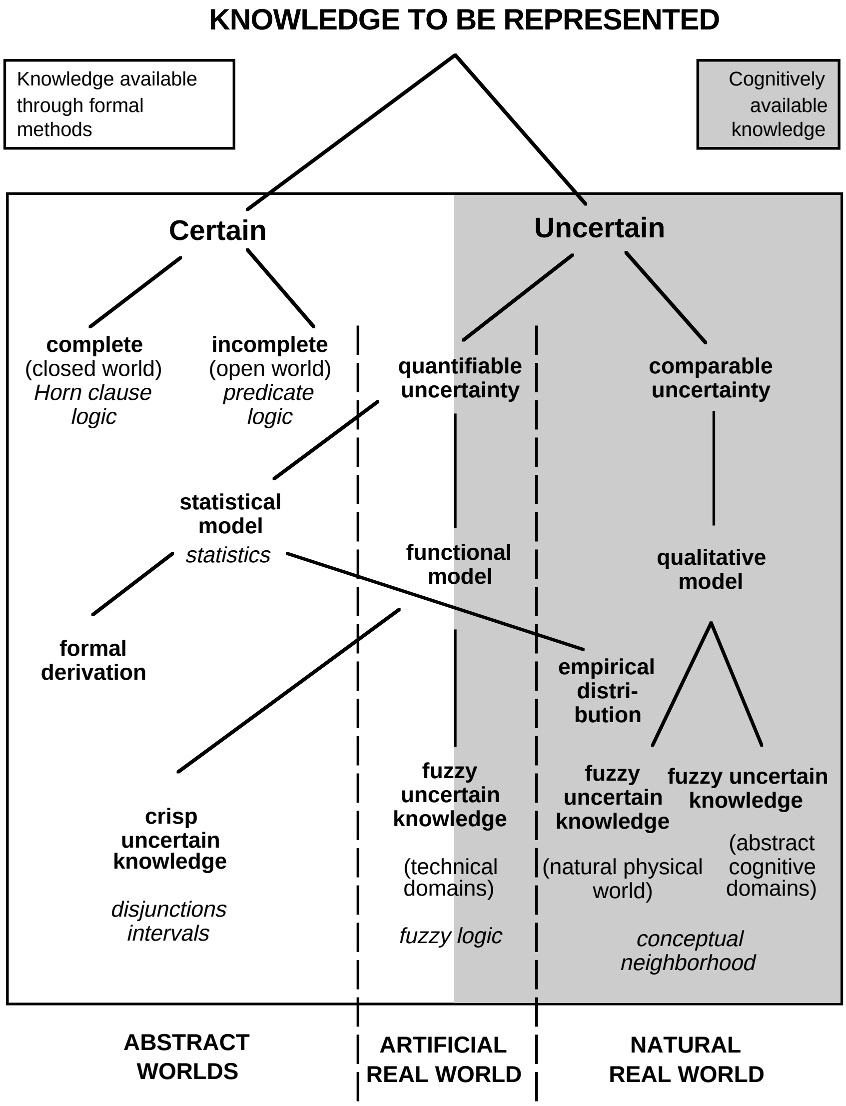

Figure 1: Classification of knowledge types on the basis of accessibility criteria

A classification of knowledge types

The classification of knowledge types presented in Figure 1 distinguishes different types of situations in which knowledge can be acquired. I will first present the distinctions made in this classification and subsequently I will associate types of approaches which appear suitable for dealing with the respective situations and knowledge.

The top-level distinction in this classification is made between *certain* knowledge and *uncertain* knowledge. Certain knowledge is only available in *abstract* domains where facts and rules can be postulated to be true. Knowledge obtained from *concrete* domains are subjected to uncertainty, due to limitations in modeling and knowledge acquisition. Certain knowledge may be further classified into *complete* knowledge in which the non-existence of a true fact is equivalent to the existence of the negated fact *(closed world assumption)* and *incomplete knowledge* in which this assumption is not generally valid.

Uncertain knowledge can be further classified into knowledge with *quantifiable* uncertainty, i.e. knowledge with which we can associate an absolute degree of uncertainty, and *comparable* uncertainty, i.e. knowledge that we merely can rank according to relative degrees of uncertainty.

Uncertainty may be due to statistical effects, due to incomplete knowledge, or due to fuzziness; accordingly uncertainty can be quantified on the basis of statistical models or on the basis of functional models. Statistical models can be classified according to the source of the probability distribution: are the probability distributions derived from a *formal* model on the basis of assumptions about the processes involved (e.g. each face of a die is equally likely to be thrown) or from empirical observations.

In the class of functional models of uncertainty we can distinguish between those yielding *crisp* sets of possible values and those yielding *fuzzy* sets. The crisp sets may be either intervals (in the case of continuous domains) or disjunctions (in the case of discrete domains). Fuzzy sets can be obtained when the functional dependencies are clearly enough defined to allow for a quantitative characterization of the elasticity of constraints. This is the case in *synthetic* domains in which fuzzy membership values can be derived from the system specifications. For an interesting illustration of the advantages of studying synthetic systems rather than natural systems, see Braitenberg's experiments in synthetic psychology [Braitenberg 1984].

In the domain of natural language or other natural domains, quantitative membership values generally are not available; but frequently membership values can be compared like in “I would rather say ‘John is tall’ than ‘John is very tall’”. Such comparisons lead to qualitative models of uncertainty. We can further distinguish between fuzzy uncertain knowledge about the natural physical world where we may have empirical sensory evidence for comparative uncertainty judgments and fuzzy uncertain knowledge about abstract cognitive domains where it is not possible to identify objective correlates to the comparative judgments.

In summary, this classification yields three different major types of knowledge processing situations: 1) The situation of abstract domains in which we work on the basis of formal specifications; here we either deal with complete certainty or with well-controlled uncertainty. 2) Real-world situations which are well-controlled in the sense that we know all relevant factors and the possible extreme situations; this is the case in many technical domains whose function *in principle* could be completely described on the basic level by exhaustive analysis; such artificial domains are ideal for the application of classical fuzzy logic. 3) In artificial intelligence, we are frequently confronted with situations in the natural world in which the available knowledge does not justify a quantitative fuzzy set approach; in particular, we must deal with open worlds in which representations in terms of numerical membership functions do not make sense; instead, qualitative knowledge may be available. For this type of situation, more adequate representational approaches are needed.

I would like to suggest that fuzzy reasoning may become as successful as fuzzy control and will be useful for natural language processing and other AI applications when it is integrated with deep conceptual representations of the domain and of the language describing the domain. In particular, the ‘horizontal’ dimension connecting cooperating and competing concepts must be exploited in addition to the ‘vertical’ dimension which connects concepts with their definitions. This horizontal dimension is required for determining which concepts are to be used in the first place. Depending on this selection, fuzziness may or may not play a role for the given task.

A possibility for extending the scope of fuzzy reasoning on the basis of the notion of *conceptual neighborhood* is currently being explored by the author. The notion of conceptual neighborhood was introduced in the context of qualitative temporal and spatial reasoning [Freksa 1992]. It addresses the logical structure of the represented domain under the operations that can take place in the domain. Specifically, two relations are conceptual neighbors when there is an operation in the represented domain which can transform one of the two relations into the other. Conceptual neighborhood can be used to represent horizontal relationships between concepts and to form coarse concepts from fine concepts and vice versa.

4 Criticism of the fuzzy systems approach

In a prize-winning article, Charles Elkan expressed his surprise about the success of fuzzy logic [Elkan 1993]. Specifically, Elkan attempts to show formally that fuzzy logic collapses to classical two-valued logic and he argues that it is not adequate for reasoning about uncertain evidence in expert systems from an empirical perspective. Elkan’s overall judgment is that fuzzy logic is fundamentally wrong and will cause serious problems in more challenging applications.

Like Zeno's famous ancient paradoxes [c.f. Vlastos 1967], Elkan’s paradox appears rather convincing at first glance, when the reader submits himself to the formal framework used by the author. However, a more careful analysis of Elkan’s argument reveals that – like Zeno – he presents a restricted view. He does this by forcing a notion of equivalence valid for two-valued logic on the analysis of fuzzy logic. In this way, Elkan addresses only special cases in which two-valued and fuzzy logic in fact are equivalent [cf. Shastri, to appear]. It is easy to provide numerical counterexamples to Elkan’s equivalence assertion using standard notions of fuzzy sets for intermediate truth values.

Limitations of formal analysis

I will take the occasion of this attack on fuzzy systems on the basis of purely formal arguments to draw your attention to an important issue of knowledge representation systems which *in principle* cannot be resolved by formal analysis. Formal analysis helps us to understand systems which are entirely formal. However, in representing knowledge about the real world, one part of the system is the body of knowledge to be represented, another part is the representing formal structure, and a third part establishes the relations between the body of knowledge and the formal structure [c.f. Palmer 1978].

Only the second part, the formal structure, can be rigorously analyzed formally. The first part, the body of knowledge is not accessible with formal tools directly; human perception and/or intuition present the knowledge to be represented by the formalism. The representation of the knowledge can be only as good as our understanding of the structure of the knowledge itself! For example, if we take a natural language statement like "John is tall" merely as a different way of writing the predicate logic statement *tall (John) is true*, then we will never be able to reach aspects of the original statement which are not coded in the predicate logic 'equivalent'. We easily can become victims of the same kind of fallacy Charles Elkan was subjected to when he viewed the world of fuzzy logic through the glasses of two-valued logic.

In reasoning about the real world, making formally sound inferences is only one aspect. Equally important – but much more difficult than widely believed – is to adequately formalize real world knowledge in the first place. The "paradoxical success" of the fuzzy logic approach in restricted domains may be considered as an indication that clearer perception or sharper intuition about the relation between the domain knowledge and the domain states have been involved in the knowledge formalization process. Of course, having found a new representation structure, we must develop appropriate reasoning methods to go along with.

It may well be impossible to find methods which will both fit the more adequate representation structure for real world phenomena and satisfy the classical criteria of formal analysis such as logical equivalence. Nevertheless, the resulting inferences may be more useful than formally correct inferences on the basis of less adequate knowledge structures. The notion of representational adequacy is not yet sufficiently understood.

5 Soft computing

Fuzziness is one of several aspects of our knowledge about the real world which must be taken into account in knowledge representation and processing. In general, we must deal with imprecision, uncertainty, and partial truth. The human mind can be viewed as a working realization of a system which rather successfully deals with all of these aspects simultaneously. In contrast to conventional (hard) computing approaches, systems that are tolerant of these aspects of everyday knowledge are united by the label *soft computing.*

The guiding principle of soft computing is: Exploit the tolerance for imprecision, uncertainty, and partial truth to achieve tractability, robustness, and low cost solutions [Zadeh 1994]. The basic ideas underlying soft computing have links to many early influences of fuzzy set theory, including Zadeh's original publication on fuzzy sets [Zadeh 1965], his paper on the analysis of complex systems and decision processes [Zadeh 1973], and his paper on possibility theory and soft data analysis [Zadeh 1981].

Besides fuzzy logic, probabilistic approaches for reasoning under uncertainty and related models for belief maintenance and revision play an important role. In artificial intelligence, Pearl's probabilistic reasoning in Bayesian networks, Nilsson's probabilistic logic, the certainty factor model used in the MYCIN expert system for medical diagnosis, Dempster-Shafer's theory of evidence have attracted much attention (c.f. [López de Mántaras 1990], [Kruse et al. 1991]). In the mid 1980s, neural network theory also joined into the soft computing effort.

Combining different approaches to soft computing

It has become evident during the past ten or twenty years, that no single approach to the study of cognitive or intelligent processes will succeed in understanding the interactions of cognitive agents with complex environments and no single approach to representing complex knowledge will fulfill all our requirements. Successful AI approaches must take into account effectiveness, efficiency, timeliness, robustness, adequacy, and cost of the solutions. Classical requirements like provability of correctness and completeness of the solution can be expected as little from computer systems reasoning about complex situations as from humans in the same situation.

After an era of increasing specialization in almost all areas of research and technology, we have now entered an era in which the interaction of approaches is of particular importance and concern. This is true for numerous areas, but interdisciplinary efforts like cognitive science and multi-approach efforts like the Berkeley Initiative in Soft Computing (BISC) might serve as examples. Such efforts require a considerable amount of re-orientation, as we have to recognize that the former competitors must become partners.

Although fuzzy set theory and fuzzy logic have faced strong opposition from conventionally oriented theoreticians in artificial intelligence and logic during the past 30 years, the rapidly growing number of successful applications developed mainly in Japan have shifted the focus of interest from local formalistic concerns to global system considerations. As in the case of the Fifth Generation Computing Project in the early 80s it required the Japanese challenge before European and American opposition was matched by a growing interest in the industry and an increased willingness by theoreticians to understand the principles of soft computing.

In Europe, we now find an increasing interest in the theory and applications of soft computing techniques in artificial intelligence. This is evident from the growing number of fuzzy logic workshops and soft computing contributions to artificial intelligence conferences, from the establishment of special interest groups in fuzzy logic and soft computing (e.g. within the German computer science society) and from the growing number of tutorials offered both by the industry and by academic institutions.

Three papers on specific topics of fuzzy reasoning

In the remainder of this chapter you find three articles dealing with soft computing for artificial intelligence.

The article by Sascha Dierkes, Bernd Reusch, and Karl-Heinz Temme presents a tool for supporting the representation of fuzzy knowledge and for fuzzy reasoning in an experts system shell.

The article by Jörg Gebhardt uses the possibilistic interpretation of fuzzy sets in the context of model-based reasoning. The approach described in the paper allows for evidential reasoning in multidimensional numerical hypothesis spaces under imprecision and uncertainty.

Jochen Heinsohn presents a language to extend the taxonomic knowledge representation approach of terminological logics by a probabilistic knowledge representation component. In this way uncertain knowledge can be included in the reasoning process.

Acknowledgments

I thank Jörg Gebhardt, Jochen Heinsohn, and Ramon López de Mántaras for valuable comments on an earlier version of this paper.

References

[Braitenberg 1984] V. Braitenberg *Vehicles.* MIT-Press, Cambridge 1984.

[Dubois et al. 1993] D. Dubois, H. Prade, R.R. Yager (eds.) *Readings in Fuzzy sets for intelligent systems.* Morgan Kaufmann Publishers, San Mateo 1993.

[Elkan 1993] C. Elkan: The paradoxical success of fuzzy logic. *Proc. AAAI-93*. 698-703, AAAI Press/MIT Press, Menlo Park 1993.

[Freksa 1992] C. Freksa: Temporal reasoning based on semi-intervals, *Artificial Intelligence* 54 (1992) 199-227.

[Kruse et al. 1991] R. Kruse, E. Schwecke, J. Heinsohn: *Uncertainty and vagueness in knowledge based systems: numerical methods*. Series Artificial Intelligence, Springer, Heidelberg 1991.

[Kruse et al. 1994] R. Kruse, J. Gebhardt, F. Klawonn: *Foundations of Fuzzy Systems*. John Wiley and Sons, Chichester, 1994.

[López de Mántaras 1990] R. López de Mántaras: *Approximate reasoning models*. Ellis Horwood, Chichester 1990.

[Munakata, Jani 1994] T. Munakata, Y. Jani: Fuzzy Systems: An overview. *Communications of the ACM* vol. 37, 3, 69-76, 1994.

[Palmer 1978] S.E. Palmer: Fundamental aspects of cognitive representation. In: Rosch, E., Lloyd, B. (eds.), *Cognition and categorization*, Erlbaum, Hillsdale 1978.

[Shastri, to appear] L. Shastri (ed.): Symposium on Fuzzy Logic, *IEEE Expert* (to appear).

[Vlastos 1967] G. Vlastos: Zeno of Elea. In: P. Edwards (ed.) *The Encyclopedia of Philosophy* vol. 8, 369-379. New York: Macmillan 1967.

[Zadeh 1965] L.A. Zadeh: Fuzzy sets. *Information and Control* 8, 338-353, 1965.

[Zadeh 1973] L.A. Zadeh: Outline of a new approach to the analysis of complex systems and decision processes. *IEEE Trans. SMC* 3, 1, 28-44, 1973.

[Zadeh 1981] L.A. Zadeh: Possibility theory and soft data analysis. *Mathematical Frontiers of the Social and Policy Sciences,* L. Cobb and R.M. Thrall (eds.), 69-129. Westview Press, Boulder 1981.

[Zadeh 1994] L.A. Zadeh: Fuzzy logic, neural networks, and soft computing. *Communications of the ACM* vol. 37, 3, 77-84, 1994.

[Zimmermann 1992] H.-J. Zimmermann. *Fuzzy set theory and its applications*. Kluwer, Boston 1992.