Fuzzy Logic in Artificial Intelligence

Peter Klement **Wolfgang Slany***

*Christian Doppler Laboratory for Expert Systems
E184/2, TU Wien, A-1040 Vienna, Austria, Europe
Phone: +43–1–58801–6141 Fax: +43–1–5055304
URL: http://www.dbai.tuwien.ac.at:8080/staff/slany.html
E-Mail: wsi@vexpert.dbai.tuwien.ac.at

CD–Technical Report 94/67

Labor
Expertensysteme

**Christian Doppler Laboratory
for Expert Systems**

**Technische Universität Wien
Institut für Informationssysteme
Abteilung für Datenbanken und Expertensysteme**

Fuzzy Logic in Artificial Intelligence

Erich Peter Klement
Fuzzy Logic Laboratorium Linz
Johannes Kepler Universität
A-4040 Linz, Austria
E-mail: klement@flll.uni-linz.ac.at

Wolfgang Slany
Christian Doppler Laboratory for Expert Systems
Technische Universität Wien
A-1040 Vienna, Austria
E-mail: wsi@vexpert.dbai.tuwien.ac.at

June 23, 1994

Abstract

After a basic introduction of fuzzy logic, we discuss its role in artificial and computational intelligence. Then we present innovative applications of fuzzy logic, focusing on fuzzy expert systems, with one typical example explored in some detail. The article concludes with suggestions how artificial intelligence and fuzzy logic can benefit from each other.

I. Introduction

In 1948, Alan Turing wrote a paper [1] marking the begin of a new era, the era of the intelligent machine, which raised questions that still remain unanswered today. This era was heavily influenced by the appearance of the computer, a machine that allowed humans to automate their way of thinking.

However, human thinking is not exact. If you had to park your car *precisely* in one place, you would have extreme difficulties. To allow computers to really mimic the way humans think, the theories of fuzzy sets and fuzzy logic were created. They should be viewed as formal mathematical theories for the representation of uncertainty, which is essential for the management of real world systems, as it mimics the crucial ability of the human mind to summarize data and focus on decision relevant information.

Marvin Minsky, one of the founding fathers of artificial intelligence, once defined the latter as

> *"... the science of making machines do things that would require intelligence if done by men."*

Similarly, Lotfi A. Zadeh, who in 1965 wrote the founding paper on fuzzy set theory [2], once described the aim of this theory as being

> *"the construction of smarter machines."*

Zadeh recently coined the term MIQ (machine intelligence quotient) to refer to this particular aspect of the growing number of intelligent consumer products and industrial systems [3].

Proponents of the so-called 'strong' artificial intelligence believe that eventually, these machines will be as intelligent as we human beings are now. Thinking positively about technology, everything that is conceivable to be solved by artificial means will eventually be realized if it is interesting enough. Of course some intellectual processes have been shown to be emergent properties, such as 'consciousness'. The concept of emergent properties of complex systems was first observed by von Bertalanffy [4] in the 1920s in his study of complex biological systems. He noticed that complex assemblies of entities organized in particular ways can reveal unique properties not possessed by the individual entities alone. Emergent properties cease to exist if the whole is broken into components or if the components are organized in a different way. Additionally, emergent properties cannot be understood by the study of isolated components. Similar to the notion of a critical mass in physics, an emergent property will suddenly pop up when a sufficient amount of mass has been accumulated. Contrary to reductionistic approaches, these approaches normally assume a holistic view of the world, i.e. something complex can be more than simply the accumulation or 'sum of its parts'. Of course, as with the atomic bomb, which was in a certain sense the first artificial application of the physical effect described above, the ethical aspects have to be carefully considered. One has to be aware that any technology can be used for good or for evil. However, not the technology in itself is good or bad, but instead the humans that use it are so, since technology has so far been only a tool for human beings. In the case of intelligence, this might be not true anymore, since advanced intelligence may entail new ethical needs, but these new forms of intelligence have not yet reached a level where ethical aspects become prevalent.

II. Some elements of fuzzy logic

In this section, we briefly outline how fuzzy logic extends classical Boolean logic (or, equivalently, how fuzzy set theory generalizes Cantorian set theory).

Given a (crisp) universe of discourse X, a *fuzzy subset* A of X (see, e.g., [2, 5]) is characterized by its *membership function*

$$\mu_A : X \to [0,1],$$

where for $x \in X$ the number $\mu_A(x)$ is interpreted as the *degree of membership* of x in the fuzzy set A or, equivalently, as the *truth value* of the statement '*x is element of A*'.

The membership function of a fuzzy set is a natural generalization of the *characteristic function* of a (classical) subset A of X,

$$\mathbf{1}_A : X \to \{0,1\},$$

assigning to each element x in X the value 1 whenever x belongs to A, and the value 0 otherwise.

In order to generalize the set-theoretical operations like *intersection* and *union* (or the corresponding Boolean logical operations *conjunction* and *disjunction*, respectively), we need triangular norms and conorms [6, 7, 8]: A *triangular norm* (*t-norm*) is a binary operation on $[0,1]$, i.e., a function $T : [0,1]^2 \to [0,1]$, which is commutative, associative, monotone in both components, and satisfies the boundary condition

$$T(x,1) = x.$$

If T is a t-norm, then the dual *triangular conorm* (*t-conorm*) $S : [0,1]^2 \to [0,1]$ is defined by

$$S(x,y) = 1 - T(1-x, 1-y).$$

There are many, in fact infinitely many, t-norms and t-conorms, only few of which are used in applications. The most important t-norms, together with their dual t-conorms, are the following:

Minimum $T_{\mathbf{M}}$, *Maximum* $S_{\mathbf{M}}$

$$T_{\mathbf{M}}(x,y) = \min(x,y), \qquad S_{\mathbf{M}}(x,y) = \max(x,y),$$

Product $T_{\mathbf{P}}$, *Probabilistic Sum* $S_{\mathbf{P}}$

$$T_{\mathbf{P}}(x,y) = x \cdot y, \qquad S_{\mathbf{P}}(x,y) = x + y - x \cdot y,$$

Lukasiewicz t-norm $T_{\mathbf{L}}$, *Bounded Sum* $S_{\mathbf{L}}$

$$T_{\mathbf{L}}(x,y) = \max(x+y-1,0), \qquad S_{\mathbf{L}}(x,y) = \min(x+y,1).$$

Weakest t-norm $T_{\mathbf{W}}$, *Strongest t-conorm* $S_{\mathbf{W}}$

$$\begin{aligned} T_{\mathbf{W}}(x,y) &= \begin{cases} \min(x,y) & \text{if } \max(x,y) = 1, \\ 0 & \text{otherwise,} \end{cases} \\ S_{\mathbf{W}}(x,y) &= \begin{cases} \max(x,y) & \text{if } \min(x,y) = 0, \\ 1 & \text{otherwise.} \end{cases} \end{aligned}$$

Given a t-norm T, its dual t-conorm S, and fuzzy subsets A, B of the universe X, the membership functions of the intersection $A \cap B$, the union $A \cup B$, and the complement A^c are given by:

$$\begin{aligned} \mu_{A\cap B}(x) &= T(\mu_A(x), \mu_B(x)), \\ \mu_{A\cup B}(x) &= S(\mu_A(x), \mu_B(x)), \\ \mu_{A^c}(x) &= 1 - \mu_A(x). \end{aligned}$$

The values $\mu_{A\cap B}(x)$, $\mu_{A\cup B}(x)$, and $\mu_{A^c}(x)$ describe the truth values of the statements '*x is element of A AND x is element of B*', '*x is element of A OR x is element of B*', and '*x is NOT element of A*', respectively.

Given a t-norm T, its dual t-conorm S and propositions P, Q with truth values $||P||$ and $||Q||$, respectively, there are two main extensions of the Boolean implication $P \Rightarrow Q$:

S-implication

$$||P \Rightarrow_S Q|| = S(1 - ||P||, ||Q||),$$

R-implication

$$||P \Rightarrow_R Q|| = \sup\{\alpha \in [0,1] \mid T(\alpha, ||P||) \leq ||Q||\},$$

In Boolean logic, i.e., with truth values 0 and 1 only, S- and R-implication always are equivalent, which is no longer true in fuzzy logic: for the t-norm $T_{\mathbf{M}}$ the corresponding S-implication becomes the *Kleene-Dienes* implication and the R-implication is the *Gödel* implication [9] which are quite different. Note, however, that in the case of the Lukasiewicz t-norm $T_{\mathbf{L}}$ S- and R-implication coincide: they both yield the *Lukasiewicz* implication.

The fact that there is a wide range of possibilities for extending Boolean operations to fuzzy logic may look disturbing for beginners in the field. However, it reflects the richness of fuzzy logic, as it allows for a very sensitive fine tuning when modeling real world situations.

III. Fuzzy logic, artificial intelligence, and computational intelligence

It is important to note that the term fuzzy logic is used in two distinct senses. In its narrower sense, fuzzy logic is only one branch of fuzzy set theory. Fuzzy set theory was invented by Zadeh to be able to better represent such everyday notions as the set of 'tall persons'. Of course, this set is defined vaguely, and persons will more or less be a member of it, i.e. member to a certain degree. Fuzzy logic in this narrow sense deals in a natural way with the representation and inference from such vaguely formulated or uncertain knowledge, similarly to classical logic which deals with crisp knowledge where statements can only be either true or false (well, almost, at least if you do not count the findings of Kurt Gödel). In recent years, however, it has become increasingly common to employ the term fuzzy logic in a much broader sense, making the difference between the notions of fuzzy set theory and fuzzy logic vanish. To avoid confusion, we follow the trend to use fuzzy logic in its general sense.

James Bezdek, editor in chief of the *IEEE transactions on fuzzy systems*, defined fuzzy logic in a delightful essay [10] to be one part of 'computational intelligence', altogether with such research areas as neural networks, evolutionary computation, and genetic algorithms. Bezdek contrasts the ABC's on intelligence: artificial, biological and computational. In the strictest sense, computational intelligence "depends on numerical data supplied by manufacturers and [does] not rely on 'knowledge'." Artificial intelligence, on the other hand, uses what Bezdek calls 'knowledge tidbits'. Heuristically constructed artificial intelligence such as an expert system is an example. Practicing knowledge engineers and neural smiths know the distinction is at times not precise. Expert extraction of feature data for training a layered perceptron certainly falls in the area of artificial intelligence. Using these features to train the layered perceptron is primarily computational. Fuzzy inference engines crafted by experts fall into the definition of artificial intelligence. Algorithmic tuning of the engine with raw data, however, is computational intelligence.

Even though the boundary between computational intelligence and artificial intelligence is not distinct, we can, making certain assumptions, monitor the volume of research activity in each. Indeed, the separate identities of computational intelligence and artificial intelligence are confirmed by inspection of the recent volume of publishing and patent activity [11].

However, the term 'Computational Intelligence' itself is not undisputed, since it had already been widely used to mean artificial intelligence before it was redefined by Bezdek, see for example the journal 'Computational Intelligence', published since 1985, the conference 'Computational Intelligence' taking place annually since 1988, and numerous other publications and organizations using the term in this traditional sense.

In both cases, artificial intelligence as well as fuzzy logic, one tries in some sense to imitate life in its problem-solving capability. The ways how to achieve this goal are different in many respects, but there are also many common points where the two fields overlap: Robert Marks [11] counted 4811 entries on fuzzy logic in the INSPEC data base from 1989 to 1993, containing citations from over 4000 selected journals, books, conference proceedings and technical reports — "22% of them [were] cross categorized in the expert system category, and 12% with neural networks." Based on various 'bean countings', Marks concludes that the overlapping areas cover, depending on the way to count, from 14% to 33%.

It should not be left untold that there has been a lot of scientific antagonism between fuzzy logic and artificial intelligence, and, accordingly, skeptics on both sides exist and treat the other side with reservation, if not with open hostility. There are many reasons for this, e.g. some critics of fuzzy logic credit the word 'fuzzy' for being too controversial and misleading in itself, others maintain that anything that can be done with fuzzy logic and fuzzy set theory can be done equally well with classical logic and probability theory [12][1], and still others insist on denying fuzzy logic the status of a logic itself [13]. Of course these claims were refuted [14, 15, or see discussions in the archives of the news-groups mentioned later in this article]. Fuzzy logic in its narrow sense is simply a logic *of* fuzziness, not a logic which *itself* is fuzzy. Just as the laws of probability are not random, so the laws of fuzziness are not vague.

On the other hand, critics of artificial intelligence have observed that the sometimes over-ambitious predictions made in the past did not come true. Some even go as far as to deny that there has been even one successful expert system implemented that really became used. Others believe that the aim to create artificial intelligence is useless and impossible on philosophical grounds. However, such views are likely to become muted with the passage of time and a better understanding of the basic ideas underlying the theories of both artificial intelligence and fuzzy logic. We observe nevertheless that, nurtured by the current success of fuzzy logic in the real world, dangerously unrealistic predictions and claims appear again. Bart Kosko, a respected scholar in the field and author of a best-selling textbook on 'Neural Networks and Fuzzy Systems' [16] for instance predicts for the next few decades fuzzy logic based natural language understanding, machines that write interesting novels and screenplays in a selected style such as Hemingway's, or even sex robots with a humanlike repertoire of behavior [14]. Some researchers suggest however that as attempts are made to make fuzzy systems larger, they will encounter similar difficulties as conventional reasoning methodologies. Fuzzy logic is certainly not a philosopher's stone solving all problems that confront us today. But it has a considerable potential for practical applications. The management of uncertainty will be of growing importance. This uncertainty can have various reasons, ranging from uncertainty due to the lack of knowledge or evidence, due to an abundance of complexity and information, to uncertainty due to the fast and unpredictable development of scientific, political, social, and other structures nowadays.

[1] But Cheeseman, the author of [12], also rejects nonmonotonic reasoning, default logic, and Dempster Shafer's theory, arguing that probabilities are better suited to model the world, and that the other methods are at most harmless if not outright wrong.

IV. Applications of fuzzy logic

The applications of fuzzy technologies fall mainly into two categories: fuzzy control applications, which are often rather simple but very efficient fuzzy rule-based systems, such as autofocusing systems in cameras, washing machines, automobile transmissions, subway control, or even handwriting recognition. In these applications, fuzzy logic is used as a powerful knowledge representation technique that allows to hide unessential details and to handle uncertain data. However, their efficiency depends also heavily on the use of sensors and effectors, thus their success should really be explained by the interaction of these various parts. The second category consists of those much more complex systems that aim at supporting or even replacing a human expert. Such applications are exemplified by medical diagnosis systems, securities funds and portfolio selection systems, traffic control systems, fuzzy expert systems, and fuzzy scheduling systems. In this second category, there are still many problems that remain to be addressed, and there is an equally pressing need for a better understanding of how to deal with knowledge-based systems in which knowledge is both uncertain and imprecise.

Areas where fuzzy logic and artificial intelligence meet in current research include: fuzzy expert systems (e.g., for medical diagnosis or intelligent tutoring systems), theoretical investigations (e.g., combinations of fuzzy logic with modal logics and other forms of defeasible reasoning, i.e. based on questionable knowledge; this also includes investigations into fuzzy logic programming languages such as fuzzy extensions of PROLOG), machine learning (e.g., combinations of fuzzy logic with neural networks, genetic algorithms, associative memories, symbolic learning methods such as case based reasoning), robotics (involving motion control and planning capabilities, e.g. when flying a fully automated helicopter or driving a car on a freeway), pattern matching (e.g., face recognition), fuzzy deductive databases (e.g., to ease data retrieval in geographic information systems), or constraint satisfaction problem solving methods (applied for example in manufacturing process scheduling [17], or in bridge design).

V. Fuzzy expert systems

Let us take a closer look at fuzzy expert systems as the archetypical spin off coming from the combination of techniques from fuzzy logic and artificial intelligence. Classical expert systems are computer programs that emulate the reasoning of human experts or perform in an expert manner in a domain for which no human expert exists. This could be due to a dangerous working environment or simply because of a domain that is to large for one human being. These expert systems typically reason with uncertain and imprecise information, using various methods besides fuzzy logic to handle them. There are many sources of imprecision and uncertainty. The knowledge that the expert systems embody is often not exact, in the same way as a human's knowledge is imperfect. Given facts or user-supplied information are also often uncertain.

An expert system is typically made up of at least three parts: an inference engine, a knowledge base, and a working memory. The inference engine uses the domain knowledge together with acquired information about a problem to provide an expert solution. The knowledge base contains the expert domain knowledge for use in problem solving, very often in form of explicit facts and IF-THEN rules.

A fuzzy expert system, usually, is an expert system that uses a collection of fuzzy membership functions and rules to reason about data. The rules in a fuzzy expert system are typically of a form similar to the following:

IF heat is *low* AND pressure is *high* THEN valve is *closed*

where 'heat' and 'pressure' are (linguistic) input variables, i.e., names for known data values, 'valve' is a (linguistic) output variable, i.e., a name for a data value to be computed, *low* is one of the possible linguistic values of the variable 'heat' described by membership function of the corresponding fuzzy set, *high* is a linguistic value of the variable 'pressure', and *closed* is a linguistic value of the variable 'valve'. The antecedent (the rule's premise) describes to what degree the rule applies, while the conclusion (the rule's consequent) assigns a fuzzy set or, if defuzzification takes place, a crisp value to each of the output variables. Most tools for working with fuzzy expert systems allow for more than one conclusion per rule. The set of rules in a fuzzy expert system is known as the rulebase or knowledge base.

The general inference process proceeds in three (or four) steps.

1. In the fuzzification step, the linguistic terms defined through their associated fuzzy membership functions are matched with the actual values of the input variables, to determine the degree of truth for each rule's premise.

2. In the inference step, the truth values for the premises are propagated to the conclusion part of each rule. This results for each rule in one fuzzy subset that is assigned to an output variable. Usually, only *minimum* or *product* are used as inference methods. In *minimum* inferencing, the output membership function is clipped off at the height corresponding to the rule premise's computed degree of truth. In *product* inferencing, the output membership function is scaled by the rule premise's computed degree of truth.

3. In the composition step, all fuzzy subsets assigned to a given output variables are combined to form a single fuzzy subset for each output variable. Also in this case, two ways of composition dominate in most applications, namely, *max* and *bounded sum*. In *max* composition, the combined output fuzzy set is constructed by taking the pointwise maximum over all of the membership functions of the fuzzy sets assigned to the output variable by the inference rule. In *bounded sum* composition, the combined output fuzzy set is constructed by taking the pointwise sum (cut off at level 1 whenever it would be exceeded) of all the membership functions of the fuzzy sets assigned to the output variable by the inference rule.

4. The optional defuzzification step is used when it is useful and/or necessary to convert the output fuzzy set into a crisp value by choosing a crisp number which is representative for the fuzzy output set. This is almost always the case in technical control problems where the output must be crisp, less often in expert systems where output fuzzy sets are sometimes linguistically approximated. There are at least 30 different defuzzification methods. Two of the more common techniques are the *center of gravity* and *center of maxima* methods. In the *center of gravity* method, the crisp value of the output variable is computed by finding the variable value of the center of gravity of the membership function for the fuzzy value. In the *center of maxima* method, the

midpoint of the region where the fuzzy set assumes its maximum truth value is chosen as the representative crisp value for the output variable.

It should be mentioned that *min* and *product* are two out of an infinite number of possible multiple-valued generalizations of the logical operation 'and' (conjunction) in Boolean (two-valued) logic.

VI. A typical example

To cite one of the most prominent and successful fuzzy expert systems, we have to refer to a very long ranging project initiated as early as 1976 by Klaus-Peter Adlassnig and resulting in a system in use today. 'CADIAG-2', which is currently evolving to become 'CADIAG-3', is a medical diagnosis system based on fuzzy expert system technology ([3] contains a recent paper about this very large project which has resulted in an enormous amount of publications; [5] also contains a description of CADIAG-2). A typical rule of this system looks as follows (the rule has been slightly simplified for this example):

IF	fever	is	*frequent*	AND
	high fever	is	*frequent*	AND
	knee dropsy	is	*rare*	AND
	carditis	is	*very-rare*	AND
	articular pain	is	*almost-always*	AND
	erythema	is	*frequent*	AND
	previous tonsillitis	is	*very-frequent*	AND
	staphylokokkus	is	*never*	AND
	increased AST	is	*almost-always*	
THEN	rheumatic fever	is	*plausible*	

Here one can see again the two key concepts which play a central role in the application of fuzzy logic in expert systems. The first is that of a linguistic variable such as 'high fever', that is, a variable whose values are words or sentences in a natural or synthetic language such as *frequent* or *rare*. The other is that of a fuzzy IF-THEN rule in which the antecedent and consequent are propositions containing linguistic variables. The essential function served by linguistic variables is that of granulation of variables and their dependencies.

In effect, the use of linguistic variables and fuzzy IF-THEN rules results — through granulation — in soft data compression which exploits the tolerance for imprecision and uncertainty. Of course, the effective membership functions represented by words such as *very-rare* have also to be determined and must be known at inference time to the inference engine.

For a detailed account of what expert systems in general and fuzzy expert systems in particular are and how they work, we refer to [18, 19, 5].

VII. Can fuzzy logic help to reduce computational complexity?

In the final 'Future perspectives' section of his book [5], Hans-Jürgen Zimmermann, who is also editor in chief of the influential journal *Fuzzy Sets and Systems*, points out an important potential benefit that fuzzy logic could bring to the field of artificial intelligence. In particular, for many problems encountered in subfields of artificial intelligence, such as logic programming, deductive databases, or constraint satisfaction, complexity analysis results range from intractable to undecidable, i.e. these problems are computationally very hard to solve. It is very easy to encounter problems not solvable with current computing power. Zimmermann therefore raises the question whether fuzzy logic can help to solve such large and complex problems. At the time of the writing of his book, the answer was still 'no'.

We shortly restate here some common notions from complexity theory to analyze whether another answer can be found. Many of the complex problems most commonly encountered in the mentioned subfields of artificial intelligence have been proven to be so-called NP hard problems, referring to their combinatorially explosive need for computing resources. Especially the run-time behavior of algorithms that solve these hard problems is at least exponential in the size of the problem description. NP hardness formally implies that these problems are 'at least as hard' as so-called NP complete problems. According to Garey and Johnson [20], NP complete problems are known to be

> *"... 'just as hard' as a large number of other problems that are widely recognized as being difficult and that have been confounding the experts for years.*

> *... the knowledge that [a problem] is NP complete does provide valuable information about what lines of approach have the potential of being most productive. Certainly the search for an efficient, exact algorithm should be accorded low priority."*

The main result is that an exact and efficient solution for NP hard problems has eluded many researchers until now. These problems have therefore been termed intractable. It is however possible to relax one of two criteria, either exactness or efficiency, in which case the other criterion can be fulfilled in many cases. One suggestion could be to relax the problem somewhat in its unimportant characteristics, i.e. to model a simplified version that does yield an acceptable solution efficiently. This is often sufficient for real world optimization problems. Another suggestion is indicated by the second sentence in above quotation, which is worth some more investigation. In particular, it is interesting to know that NP complete problems can be solved in polynomial time by a nondeterministic computer. The scenario is often such that a solution has to be 'guessed', for instance by consulting an 'oracle', followed by calling a polynomial algorithm to check whether the guessed solution is correct. This would suggest that an algorithm that intelligently 'guesses' a complete instantiation and then checks whether it is a solution could be used to construct an algorithm that finds a 'reasonably good' solution for 'almost all' problems.

The following paragraph provides a little more background about the introduced notions. It is an open problem of complexity theory whether NP is equal to P, P being the problems solvable in polynomial time, i.e. the tractable problems. Indeed, many researchers even believe that this question is currently to be the most important open question in computer science as a whole. However, most researchers think that P and NP are different. This would mean that NP complete problems would remain, at least in the worst case, intractable, i.e. their execution time grows more than polynomially when the structural parameters grow linearly. In such a case, doubling the speed of the computer does not really help since only negligible larger problems will be solvable by that computer, which is usually by far not enough. The NP complete problems are characterized by the fact that they are NP problems and that *all* other NP problems can be polynomially reduced to these NP complete problems. This means that NP complete problems are at least as difficult as any other NP problem. To prove that a problem A is NP complete, it is sufficient to show that A belongs to NP and that one other problem B known to be NP complete can be polynomially transformed into A. A general search problem H belongs to the NP hard problems if and only if there exists a *polynomial time algorithm* for some decision problem C known to be NP complete, assuming that H could be used arbitrarily often for further computations at a computational cost of one unit-time interval by the polynomial time algorithm solving C (i.e., C must be polynomial time Turing-transformable into H).

Therefore, the answer to the question posed by Zimmermann has to remain negative on theoretical grounds. On the contrary, the computational complexity will often increase when a problem is fuzzified, since in general the search space will be enlarged.

However, knowledge modeled using fuzzy logic can easily approximate a problem in its important characteristics. In particular, problems encountered in the real world, such as in manufacturing, exhibit a tendency to obey certain intangible rules, thus allowing the approximation to be done successfully. The same would not be possible for random data, representing the worst case that must also be correctly handled by exact (but much slower) algorithms. Additionally, some heuristics such as described in [17] match particularly well with fuzzy knowledge representation schemes, and thus help to efficiently solve large and complex problems. In this connection, a 'solution' to a problem is often understood pragmatically as the result of a search for the best solution that can be found using all the available resources such as *available* computers, *available* time, and *available* algorithms. So, to provide a more positive answer to the question raised by Zimmermann, we conclude that fuzzy logic *can* help to reduce complexity *if* the problem to be solved can be adequately approximated, which is true for many real world problems.

VIII. Synergy effects

To emphasize again in what respect artificial intelligence and fuzzy logic can mutually benefit from each other, we want to point out that all complex systems and machines that where built so far required more than just one basic technology in order to be successful. In a large measure, techniques from artificial intelligence and from fuzzy logic are complementary rather than competitive. We believe that it is possible to fruitfully combine techniques from both fields in many areas. The resulting hybrid systems will be more and more important in the future. Following the line of reasoning given at the begin concerning emergent properties, the synergy effect resulting in this combination is necessary to achieve the ultimate goal of creating machines that act more and more intelligently for the benefit of mankind.

IX. Further readings ...

For readers interested in gaining a better understanding of one of the two fields, fuzzy logic and artificial intelligence, we would like to refer to some good introductory texts such as Winston's book on artificial intelligence [21], or, more recently, McNeill and Freiberger's book on fuzzy logic [14]. For those wanting to dig deeper or to answer more elaborate questions, we recommend to consult some of the following texts and media (the list could of course be much longer, but we limit ourselves to the most accessible items):

- The excellent 'Encyclopedia of Artificial Intelligence' edited by Shapiro [19] covers almost all possible subjects related to this field, including numerous articles on fuzzy logic.

- The 'Readings in Fuzzy Sets for Intelligent Systems' [22] to rapidly find the most influencing articles published in this field, as well as the 'Selected Papers by L. A. Zadeh' [23].

- The internet news-groups comp.ai and comp.ai.fuzzy, also accessible electronically via various mailing lists and blackboards, including their respective frequently-asked-questions (with answers) lists, which contain pointers to other electronic sources of information such as world-wide-web-servers, pointers to the most important conferences, major journals, scientific societies, research centers, major scientific projects, book-lists, as well as names of persons-to-know and companies related to the respective fields. These news-groups are also forums to discuss all topics related to the two fields, and are equipped with searchable archives extending over several years [15, 24].

For readers searching references covering primarily the intersection of artificial intelligence and fuzzy logic, we have compiled a list of some important textbooks [25, 18, 26, 27, 28] and conference proceedings [29, 3] in the bibliography.

References

[1] Alan M. Turing. Intelligent machinery. In D. C. Ince, editor, *Mechanical Intelligence*, Collected Works of A. M. Turing. North-Holland, 1992. Original paper appeared in B. Meltzer and D. Michie (Editors), *Machine Intelligence,* 5:3–23, 1969, Edinburgh University Press, but was actually written as early as in 1948.

[2] Lotfi A. Zadeh. Fuzzy sets. *Information and Control, New York: Academic Press.*, 8:338–353, 1965. Republished in [23].

[3] Erich Peter Klement and Wolfgang Slany, editors. *Fuzzy Logic in Artificial Intelligence. Proceedings of the 8th Austrian Artificial Intelligence Conference, FLAI'93, Linz, Austria, June 1993*, volume 695 of *Lecture Notes in Artificial Intelligence*. Springer Verlag Berlin Heidelberg, 1993. A conference report is available on the net: URL: ftp://mira.dbai.tuwien.ac.at/pub/slany/flai.txt.

[4] L. von Bertalanffy. The organism considered as a physical system. In L. von Bertalanffy, editor, *General system theory*. Braziller, New York, 1968. Republished from 1940.

[5] Hans-Jürgen Zimmermann. *Fuzzy Set Theory — and Its Applications*. Kluwer Academic Publishers, 2nd, revised edition, 1991.

[6] Erich P. Klement. Operations on fuzzy sets and fuzzy numbers related to triangular norms. In *Proceedings of the Eleventh International Symposium on Multiple-Valued Logic*, pages 218–225, Norman, 1981. IEEE, New York.

[7] B. Schweizer and A. Sklar. *Probabilistic Metric Spaces*. North-Holland, Amsterdam, 1983.

[8] D. Butnariu and E. P. Klement. *Triangular Norm-Based Measures and Games with Fuzzy Coalitions*. Kluwer, Dordrecht, 1993.

[9] Rudolf Kruse, Jörg Gebhardt, and Frank Klawonn. *Foundations of Fuzzy Systems*. John Wiley and Sons Ltd., Chichester, 1994.

[10] James C. Bezdek. On the relationship between neural networks, pattern recognition and intelligence. *The International Journal of Approximate Reasoning*, 6:85–107, 1992.

[11] Robert J. Marks II. Intelligence: Computational versus artificial. *IEEE Transactions on Neural Nets*, September 1993.

[12] Peter Cheeseman. In defense of probability. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 1002–1009, Los Angeles, Ca., August 1985. Morgan Kaufmann Publishers, Inc.

[13] Charles Elkan. The paradoxical success of fuzzy logic. In *Proceedings of the Eleventh National Conference on Artificial Intelligence, AAAI'93*, pages 698–703, Washington D.C., July 1993. URL: ftp://cs.ucsd.edu/pub/elkan/paradoxicalsuccess.ps.Z; Responses: URL: ftp://ftp.cs.cmu.edu/user/ai/areas/fuzzy/doc/elkan/response.txt.

[14] Daniel McNeill and Paul Freiberger. *Fuzzy Logic*. Simon & Schuster, 1993.

[15] Erik Horstkotte, Cliff Joslyn, and Mark Kantrowitz. comp.ai.fuzzy faq: Fuzzy logic and fuzzy expert systems 1/1 [monthly posting]. URL: ftp://ftp.cs.cmu.edu/afs/cs.cmu.edu/project/ai-repository/-ai/pubs/faqs/fuzzy/fuzzy.faq, January 1994.

[16] Bart Kosko. *Neural Networks and Fuzzy Systems*. Prentice-Hall, Englewood Cliffs, NJ, 1992.

[17] Wolfgang Slany. Scheduling as a fuzzy multiple criteria optimization problem. CD-Technical Report 94/62, Christian Doppler Laboratory for Expert Systems, Technical University of Vienna, 1994. Submitted to Fuzzy Sets and Systems. URL: ftp://mira.dbai.tuwien.ac.at/pub/slany/cd-tr9462.ps.Z.

[18] A. Kandel, editor. *Fuzzy Expert Systems*. CRC Press, Boca Raton, Ca., 1992.

[19] Stuart C. Shapiro, editor. *Encyclopedia of Artificial Intelligence.* John Wiley & Sons, Inc., 2nd enlarged and revised edition, 1992.

[20] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness.* Freeman and Co., 1979.

[21] Patrick Henry Winston. *Artificial Intelligence.* Addison-Wesley Publishing Company, Inc., 1977.

[22] Didier Dubois, Henri Prade, and Ronald R. Yager, editors. *Readings in Fuzzy Sets for Intelligent Systems.* Morgan Kaufmann, 1993.

[23] R. R. Yager, S. Ovchinnikov, R. M. Tong, and H. T. Nguyen, editors. *Fuzzy Sets and Applications: Selected Papers by L. A. Zadeh.* John Wiley & Sons, 1987.

[24] Mark Kantrowitz. comp.ai faqs: 6 parts [monthly posting]. URL: ftp://-ftp.cs.cmu.edu/afs/cs.cmu.edu/project/ai-repository/ai/pubs/faqs/ai/ai_[1-6].faq and other directories at the same site, January 1994.

[25] Didier Dubois and Henri Prade. *Possibility Theory: An Approach to Computerized Processing of Uncertainty.* Plenum Press, N.Y., 1988.

[26] R. Kruse, E. Schwecke, and J. Heinsohn. *Uncertainty and Vagueness in Knowledge Based Systems – Numerical Methods.* Springer Verlag, 1991.

[27] Constantin V. Negoiţă. *Expert Systems and Fuzzy Systems.* Benjamin/Cummings, 1985.

[28] Maria Zemankova-Leech and Abraham Kandel. *Fuzzy relational data bases – a key to expert systems.* Verlag TÜV Rheinland, 1984.

[29] D. Heckerman and A. Mamdani, editors. *Uncertainty in Artificial Intelligence (Proceedings of the 9th Conference).* Morgan Kaufmann Publishers, 1993.