The Paradoxical Success of Fuzzy Logic*

Charles Elkan
Department of Computer Science and Engineering
University of California, San Diego
La Jolla, California 92093-0114

Abstract

This paper investigates the question of which aspects of fuzzy logic are essential to its practical usefulness. We show that as a formal system, a standard version of fuzzy logic collapses mathematically to two-valued logic, while empirically, fuzzy logic is not adequate for reasoning about uncertain evidence in expert systems. Nevertheless, applications of fuzzy logic in heuristic control have been highly successful. We argue that the inconsistencies of fuzzy logic have not been harmful in practice because current fuzzy controllers are far simpler than other knowledge-based systems. In the future, the technical limitations of fuzzy logic can be expected to become important in practice, and work on fuzzy controllers will also encounter several problems of scale already known for other knowledge-based systems.

1 Introduction

Fuzzy logic methods have been applied successfully in many real-world systems, but the coherence of the foundations of fuzzy logic remains under attack. Taken together, these two facts constitute a paradox, which this paper attempts to resolve. More concretely, the aim of this paper is to identify which aspects of fuzzy logic render it so useful in practice and which aspects are inessential. Our conclusions are based on a new mathematical result, on a survey of the literature on the use of fuzzy logic in heuristic control, and on our own practical experience developing two large-scale expert systems.

This paper is organized as follows. First, Section 2 proves and discusses the theorem mentioned above, which is that only two truth values are possible inside a standard system of fuzzy logic. In an attempt to understand how fuzzy logic can be useful despite this paradox, Sections 3 and 4 examine the main practical uses of fuzzy logic, in expert systems and heuristic control. Our tentative conclusion is that successful applications of fuzzy logic are successful because of factors other than the use of fuzzy logic. Finally, Section 5 shows how current work on fuzzy control is encountering dilemmas that are already well-known from work in other areas of artificial intelligence, and Section 6 provides some overall conclusions.

2 A paradox in fuzzy logic

As is natural in a research area as active as fuzzy logic, theoreticians have investigated many different formal systems, and applications have also used a variety of systems. Nevertheless, the basic intuitions are relatively constant. At its simplest, fuzzy logic is a generalization of standard propositional logic from two truth values *false* and *true* to degrees of truth between 0 and 1.

Formally, let A denote an assertion. In fuzzy logic, A is assigned a numerical value $t(A)$, called the degree of truth of A, such that $0 \leq t(A) \leq 1$. For a sentence composed from simple assertions and logical connectives "and" ($\wedge$), "or" ($\vee$), and "not" ($\neg A$ or $\overline{A}$), degree of truth is defined as follows:

Definition 1:

$$\begin{aligned} t(A \wedge B) &= \min\{t(A), t(B)\} \\ t(A \vee B) &= \max\{t(A), t(B)\} \\ t(\neg A) &= 1 - t(A) \\ t(A) = t(B) &\text{ if } A \text{ and } B \text{ are} \\ &\text{logically equivalent.} \quad \blacksquare \end{aligned}$$

In the last case of this definition, let "logically equivalent" mean equivalent according to the rules of classical two-valued propositional calculus. The use of alternative definitions of logical equivalence is discussed at the end of this section.

Fuzzy logic is intended to allow an indefinite variety of numerical truth values. The result proved here is that only two different truth values are in fact possible in the formal system of Definition 1.

Theorem 1: For any two assertions A and B, either $t(B) = t(A)$ or $t(B) = 1 - t(A)$.

Proof: Let A and B be arbitrary assertions. Consider the two sentences $\overline{A \wedge \overline{B}}$ and $B \vee (\overline{A} \wedge \overline{B})$. These are logically equivalent, so

$$t(\overline{A \wedge \overline{B}}) = t(B \vee (\overline{A} \wedge \overline{B}).$$

*This work was supported in part by the National Science Foundation under Award No. IRI-9110813.

Now

$$\begin{aligned} t(\overline{A \wedge \overline{B}}) &= 1 - \min\{t(A), 1 - t(B)\} \\ &= 1 + \max\{-t(A), -1 + t(B)\} \\ &= \max\{1 - t(A), t(B)\} \end{aligned}$$

and

$$t(B \vee (\overline{A} \wedge \overline{B})) = \max\{t(B), \min\{1 - t(A), 1 - t(B)\}\}.$$

The numerical expressions above are different if

$$t(B) < 1 - t(B) < 1 - t(A),$$

that is if $t(B) < 1 - t(B)$ and $t(A) < t(B)$, which happens if $t(A) < t(B) < 0.5$. So it cannot be true that $t(A) < t(B) < 0.5$.

Now note that the sentences $\overline{A \wedge \overline{B}}$ and $B \vee (\overline{A} \wedge \overline{B})$ are both re-expressions of the material implication $A \rightarrow B$. One by one, consider the seven other material implication sentences involving A and B

$$\begin{aligned} \overline{A} &\rightarrow B \\ A &\rightarrow \overline{B} \\ \overline{A} &\rightarrow \overline{B} \\ B &\rightarrow A \\ \overline{B} &\rightarrow A \\ B &\rightarrow \overline{A} \\ \overline{B} &\rightarrow \overline{A}. \end{aligned}$$

By the same reasoning as before, none of the following can be true:

$$\begin{aligned} 1 - t(A) < t(B) &< 0.5 \\ t(A) < 1 - t(B) &< 0.5 \\ 1 - t(A) < 1 - t(B) &< 0.5 \\ t(B) < t(A) &< 0.5 \\ 1 - t(B) < t(A) &< 0.5 \\ t(B) < 1 - t(A) &< 0.5 \\ 1 - t(B) < 1 - t(A) &< 0.5. \end{aligned}$$

Now let $x = \min\{t(A), 1 - t(A)\}$ and let $y = \min\{t(B), 1 - t(B)\}$. Clearly $x \leq 0.5$ and $y \leq 0.5$ so if $x \neq y$, then one of the eight inequalities derived must be satisfied. Thus $t(B) = t(A)$ or $t(B) = 1 - t(A)$. ∎

It is important to be clear as to what exactly is proved above, and what is not proved. The first point to note is that nothing in the statement or proof of the theorem depends on any particular definition of the meaning of the implication connective, either in two-valued logic or in fuzzy logic.Theorem 1 could be stated and proved without introducing the symbol $\rightarrow$, since $A \rightarrow B$ is used just as a syntactic abbreviation for $B \vee (\overline{A} \wedge \overline{B})$.

The second point to note is that the theorem also applies to any more general formal system that includes the four postulates listed in Definition 1. Any extension of fuzzy logic to accommodate first-order sentences, for example, collapses to two truth values if it admits the propositional fuzzy logic of Definition 1 as a special case. The theorem also applies to fuzzy set theory, because Definition 1 can be understood as axiomatizing

and complements.

On the other hand, the theorem does not necessarily apply to any version of fuzzy logic that modifies or rejects any of the four postulates of Definition 1. It is however possible to carry through the proof of the theorem in many variant systems of fuzzy logic. In particular, the theorem remains true when negation is modelled by any operator in the Sugeno class [Sugeno, 1977], and when disjunction or conjunction are modelled by operators in the Yager classes [Yager, 1980].[1]

Of course, the last postulate of Definition 1 is the most controversial one, and the postulate that one naturally first wants to modify in order to preserve a continuum of degrees of truth. Unfortunately, it is not clear which subset of classical tautologies and equivalences should be, or can be, required to hold in a system of fuzzy logic. What all formal fuzzy logics have in common is that they reject at least one classical tautology, namely the law of excluded middle (the assertion $\overline{A} \vee A$). Intuitionistic logic [van Dalen, 1983] also rejects this law, but rejects in addition De Morgan's laws, which are entailed by the first three postulates of Definition 1. One could hope that fuzzy logic is therefore a formal system whose tautologies are a subset of the classical tautologies, and a superset of the intuitionistic tautologies. However, Theorem 1 can still be proved even if logical equivalence is restricted to mean intuitionistic equivalence.[2] It is an open question how to choose a notion of logical equivalence that simultaneously (i) remains philosophically justifiable, (ii) allows useful inferences in practice, and (iii) removes the opportunity to prove results similar to Theorem 1.

3 Fuzzy logic in expert systems

Any logical system or calculus for reasoning such as fuzzy logic must be motivated by its applicability to phenomena that we want to reason about. The operations of the calculus must model the behaviour of the ideas in certain classes. One way to defend a calculus is to show that it succeeds in interesting applications, which has certainly been done for fuzzy logic. However, if we are to have confidence that the successful application of the calculus is reproducible, we must be persuaded that the calculus correctly models the interaction of all phenomena in a well-characterized general class.

The basic motivation for fuzzy logic is clear: many ideas resemble traditional assertions, but they are not

[1]The postulates of standard fuzzy logic have been used quite widely, but it happens that even the same author sometimes adopts them and sometimes does not. For example (following [Gaines, 1983]) Bart Kosko explicitly uses all four postulates to resolve Russell's paradox of the barber who shaves all men except those who shave themselves [Kosko, 1990], but in later work he uses addition and multiplication instead of maximum and minimum [public lecture at UCSD, 1991].

[2]The Gödel translations [van Dalen, 1983; p. 172] of classically equivalent sentences are intuitionistically equivalent. For any sentence, the first three postulates of Definition 1 make its degree of truth and the degree of truth of its Gödel translation equal. Thus the proof given for Theorem 1 can be carried over directly.

naturally either true or false. Rather, uncertainty of some sort is attached to them. Fuzzy logic is an attempt to capture valid patterns of reasoning about uncertainty. The notion is now well accepted that there exist many different types of uncertainty, vagueness, and ignorance [Smets, 1991]. However, there is still debate as to what types of uncertainty are captured by fuzzy logic.[3] Many papers have discussed at a high level of mathematical abstraction the question of whether fuzzy logic provides suitable laws of thought for reasoning about probabilistic uncertainty. Our conclusion from practical experience in the construction of expert systems is that fuzzy logic is not uniformly suitable for reasoning about uncertain evidence. A simple example shows what the difficulty is.

Suppose the universe of discourse is a collection of melons, and there are two predicates *red* and *watermelon*, where *red* and *green* refer to the colour of the flesh of a melon. For some not very well-known melon x, suppose that $t(red(x)) = 0.5$ and $t(watermelon(x)) = 0.8$, meaning that the evidence that x is red inside has strength 0.5 and the evidence that x is a watermelon has strength 0.8. According to the rules of fuzzy logic, $t(red(x) \wedge watermelon(x)) = 0.5$. This is not reasonable, because watermelons are normally red inside. Redness and being a watermelon are mutually reinforcing facts, so intuitively, x is a red watermelon with certainty greater than 0.5.

The deep issue here is that the degree of uncertainty of a conjunction is not in general determined uniquely by the degree of uncertainty of the assertions entering into the conjunction. There does not exist a function f such that the rule $t(A \wedge B) = f(t(A), t(B))$ is always valid, when t represents the degree of certainty of fragments of evidence. The certainty of $A \wedge B$ depends on the content of the assertions A and B as well as on their numerical certainty. This fact is recognized implicitly in probabilistic reasoning, since probability theory does not assign unique probability values to conjunctions. What probability theory says is that

$$\begin{aligned} 1 - \Big(1 - Pr(A) + 1 - Pr(B)\Big) &\leq Pr(A \cap B) \\ &\leq \min\{Pr(A), Pr(B)\}. \end{aligned}$$

The actual probability value depends on further aspects of the situation that have not been stated. For example, if the two assertions A and B are independent, then the probability of their conjunction is $Pr(A) \cdot Pr(B)$.

Although probability theory is more flexible than fuzzy logic, the red watermelon example shows that it is not a universally adequate system of laws of thought for reasoning about all types of uncertainty either. If $t(red(x)) = 0.5$ and $t(watermelon(x)) = 0.8$, then it is natural to want $t(red(x) \wedge watermelon(x)) > 0.5$, which probability theory cannot permit.

[3]Misunderstanding on these issues has reached the non-technical press: see articles based on [Kosko, 1990] in *Business Week* (New York, May 21, 1990), the *Financial Times* (London, June 5, 1990), the *Economist* (London, June 9, 1990), *Popular Science* (New York, June 1990), and elsewhere.

The difficulties identified here with fuzzy logic and probability theory as formalisms for reasoning about uncertainty do occur in practice. We have recently designed, implemented, and deployed at IBM two large-scale expert systems [Hekmatpour and Elkan, 1993; Hekmatpour and Elkan, 1992]. One system, CHATKB, solves problems encountered by engineers while using VLSI design tools. The other system, WESDA, diagnoses faults in machines that polish semiconductor wafers. The knowledge possessed by each system consists of a library of cases and a deep domain theory which is represented as a decision tree where each node corresponds to a fact about the state of the tool being diagnosed. Relevant cases are attached to each leaf of the decision tree. Roughly, the children of each node represent evidence in favour of the parent node, or potential causes of the parent node. CHATKB or WESDA retrieves an old case to solve a new problem by choosing a path through its decision tree. A path from the root to a leaf is chosen by combining *a priori* child node likelihoods with evidence acquired through questioning the user. We have found that this process of combining evidence is a type of reasoning about uncertainty that cannot be modelled adequately by the axioms of fuzzy logic, or by those of probability theory.

Methods for reasoning about uncertain evidence are an active research area in artificial intelligence, and the conclusions reached in this section are not new. Our practical experience does, however, independently confirm previous arguments about the inadequacy of systems for reasoning about uncertainty that propagate numerical factors according only to which connectives appear in assertions [Pearl, 1988].

4 Fuzzy logic in heuristic control

Heuristic control is the area of application in which fuzzy logic has been the most successful. There is a wide consensus that the techniques of traditional mathematical control theory are often inadequate. The reasons for this include the reliance of traditional methods on linear models of systems to be controlled, their propensity to produce "bang-bang" control regimes, and their focus on worst-case convergence and stability rather than typical-case efficiency. Heuristic control techniques give up mathematical simplicity and performance guarantees in exchange for increased realism and better performance in practice. A heuristic controller using fuzzy logic is shown to have less overshoot and quicker settling in [Burkhardt and Bonissone, 1992] for example.

The first demonstrations that fuzzy logic could be used in building heuristic controllers were published in the 1970s [Zadeh, 1973; Mamdani and Assilian, 1975]. Work using fuzzy logic in heuristic control continued through the 1980s, and recently there has been an explosion of industrial interest in this area; for surveys see [Yamakawa and Hirota, 1989] and [Lee, 1990]. One reason why fuzzy controllers have attracted so much interest recently is that they can be implemented by embedded specialized microprocessors [Yamakawa, 1989].

Despite the intense industrial interest (and, in Japan, consumer interest) in fuzzy logic, the technology continnational Joint Conference on Artificial Intelligence (IJCAI'91, Sydney, Australia) Takeo Kanade gave an invited talk on computer vision in which he described at length Matsushita's camcorder image stabilizing system [Uomori *et al.*, 1990], without mentioning that it uses fuzzy logic.

Almost all currently deployed heuristic controllers using fuzzy logic are similar in five important aspects. A good description of a prototypical example of this standard architecture appears in [Sugeno *et al.*, 1989].

- First, the knowledge base of a typical fuzzy controller consists of under 100 rules; often under 20 rules are used. Fuzzy controllers are orders of magnitude smaller than systems built using traditional artificial intelligence formalisms: the knowledge base of CHATKB, for example, occupies many megabytes.
- Second, the knowledge entering into fuzzy controllers is structurally shallow, both statically and dynamically. It is not the case that some rules produce conclusions which are then used as premises in other rules. Statically, rules are organized in a flat list, and dynamically, there is no run-time chaining of inferences.
- Third, the knowledge recorded in a fuzzy controller typically reflects immediate correlations between the inputs and outputs of the system to be controlled, as opposed to a deep, causal model of the system. The premises of rules refer to sensor observations and rule conclusions refer to actuator settings.[4]
- The fourth important feature that deployed fuzzy controllers share is that the numerical parameters of their rules and of their qualitative input and output modules are tuned in a learning process. Many different learning algorithms have been used for this purpose, and neural network learning mechanisms have been especially successful [Keller and Tahani, 1992; Yager, 1992]. What the algorithms used for tuning fuzzy controllers themselves have in common is that they are gradient-descent "hill-climbing" algorithms that learn by local optimization [Burkhardt and Bonissone, 1992].
- Last but not least, by definition fuzzy controllers use the operators of fuzzy logic. Typically "minimum" and "maximum" are used, as are explicit possibility distributions (usually trapezoidal), and some fuzzy implication operator.

[4] Rule premises refer to qualitative ("linguistic" in the terminology of fuzzy logic) sensor observations and rule conclusions refer to qualitative actuator settings, whereas outputs and inputs of sensors and actuators are typically real-valued. This means that two controller components usually exist which map between numerical values and qualitative values. In fuzzy logic terminology, these components are said to defuzzify outputs and implement membership functions respectively. Their behaviour is not itself describable using fuzzy logic, and typically they are implemented procedurally.

features of fuzzy controllers identified above are essential to their success. It appears that the first four shared properties are vital to practical success, because they make the celebrated credit assignment problem solvable, while the use of fuzzy logic is not essential.

In a nutshell, the credit assignment problem is to discover how to modify part of a complex system in order to improve it, given only an evaluation of its overall performance. In general, solving the credit assignment problem is impossible: the task is tantamount to generating many bits of information (a change to the internals of a complex system) from just a few bits of information (the input/output performance of the system). However, the first four shared features of fuzzy controllers make the credit assignment problem solvable for them.

First, since it consists of only a small number of rules, the knowledge base of a fuzzy controller is a small system to modify. Second, the short paths between the inputs and outputıs of a fuzzy controller mean that the effect of any change in the controller is localized, so it is easier to discover a change that has a desired effect without having other undesired consequences. Third, the iterative way in which fuzzy controllers are refined allows a large number of observations of input/output performance to be used for system improvement. Fourth, the continuous nature of the many parameters of a fuzzy controller allows small quantities of performance information to be used to make small system changes.

Thus, what makes fuzzy controllers useful in practice is the combination of a rule-based formalism with numerical factors qualifying rules and the premises entering into rules. The principal advantage of rule-based formalisms is that knowledge can be acquired from experts or from experience incrementally: individual rules and premises can be refined independently, or at least more independently than items of knowledge in other formalisms. Numerical factors have two main advantages. They allow a heuristic control system to interface smoothly with the continuous outside world, and they allow it to be tuned gradually: small changes in numerical factor values cause small changes in behaviour.

None of these features contributing to the success of systems based on fuzzy logic is unique to fuzzy logic. It seems that most current applications of fuzzy logic could use other numerical rule-based formalisms instead, if a learning algorithm was used to tune numerical values for those formalisms, as is customary when using fuzzy logic.

Several knowledge representation formalisms that are rule-based and numerical have been proposed besides fuzzy logic. For example, well-developed systems are presented in [Sandewall, 1989] and [Collins and Michalski, 1989; Dontas and Zemankova, 1990]. To the extent that numerical qualification factors can be tuned in these formalisms, we expect that they would be equally useful for constructing heuristic controllers. Indeed, at least one has already been so used [Sammut and Michie, 1991].

5 Recapitulating mainstream AI

Several research groups are attempting to scale up systems based on fuzzy logic, and to lift the architectural limitations of current fuzzy controllers. For example, a methodology for designing block-structured controllers with guaranteed stability properties is studied in [Tanaka and Sugeno, 1992], and methodological problems in constructing models of complex systems based on deep knowledge are considered in [Pedrycz, 1991]. Controllers with intermediate variables, thus with chaining of inferences, are investigated in [von Altrock *et al.*, 1992].

However, the designers of larger systems based on fuzzy logic are encountering all the problems of scale already identified in traditional knowledge-based systems. It appears that the history of research in fuzzy logic is recapitulating the history of research in other areas of artificial intelligence. This section discusses the knowledge engineering dilemmas faced by developers of fuzzy controllers, and then points to dealing with state information as another issue arising in research on fuzzy controllers that has also arisen previously.

The rules in the knowledge bases of current fuzzy controllers are obtained directly by interviewing experts. Indeed, the original motivation for using fuzzy logic in building heuristic controllers was that fuzzy logic is designed to capture human statements involving vague quantifiers such as "considerable." More recently, a consensus has developed that research must focus on obtaining "procedures for fuzzy controller design based on fuzzy models of the process" [Driankov and Eklund, 1991]. Mainstream work on knowledge engineering, however, has already transcended the dichotomy between rule-based and model-based reasoning.

Expert systems whose knowledge consists of *if-then* rules have at least two disadvantages. First, maintenance of a rule base becomes complex and time-consuming as the size of a system increases [Newquist, 1988]. Second, rule-based systems tend to be brittle: if an item of knowledge is missing from a rule, the system may fail to find a solution, or worse, may draw an incorrect conclusion [Abbott, 1988]. The main disadvantage of model-based approaches, on the other hand, is that it is very difficult to construct sufficiently detailed and accurate models of complex systems. Moreover, models constructed tend to be highly application-specific and not generalizable [Bourne *et al.*, 1991].

Many recent expert systems, therefore, including CHATKB and WESDA, are neither rule-based nor model-based in the standard way. For these systems, the aim of the knowledge engineering process is not simply to acquire knowledge from human experts, whether this knowledge is correlational as in present fuzzy controllers, or deep as in model-based expert systems. Rather, the aim is to develop a theory of the situated performance of the experts. Concretely, under this view of knowledge engineering, knowledge bases are constructed to model the beliefs and practices of experts and not any "objective" truth about underlying physical processes. An important benefit of this approach is that the organization of an expert's beliefs provides an implicit organization of knowledge about the external process with which the knowledge-based system is intended to interact.

The more sophisticated view of knowledge engineering just outlined is clearly relevant to research on constructing fuzzy controllers more intricate than current ones. For a second example of relevant previous artificial intelligence work, consider controllers that can carry state information from one moment to the next. These are mentioned as a topic for future research in [von Altrock *et al.*, 1992]. Symbolic AI formalisms for representing systems whose behaviour depends on their history have been available since the 1960s [McCarthy and Hayes, 1969]. Neural networks with similar properties (called recurrent networks) have been available for several years [Elman, 1990], and have already been used in control applications [Karim and Rivera, 1992]. It remains to be seen whether research from a fuzzy logic perspective will provide new solutions to the fundamental issues of artificial intelligence.

6 Conclusions

Applications of fuzzy logic in heuristic control have been highly successful, despite the collapse of fuzzy logic as a formal system to two-valued logic, and despite the inadequacy of fuzzy logic for reasoning about uncertainty in expert systems. The inconsistencies of fuzzy logic have not been harmful in practice because current fuzzy controllers are far simpler than other knowledge-based systems. First, long chains of inference are not performed in controllers based on fuzzy logic, so there is no opportunity for inconsistency between paths of reasoning that should be equivalent to manifest itself. Second, the knowledge recorded in a fuzzy controller is not a consistent causal model of the process being controlled, but rather an assemblage of visible correlations between sensor observations and actuator settings. Since this knowledge is not itself consistent and probabilistic, the probabilistic inadequacy of fuzzy logic is not an issue. Moreover, the ability to refine the parameters of a fuzzy controller iteratively can compensate for the arbitrariness of the fuzzy logic operators as applied inside a limited domain.

The common assumption that heuristic controllers based on fuzzy logic are successful because they use fuzzy logic appears to be an instance of the *post hoc, ergo propter hoc* fallacy. The fact that using fuzzy logic is correlated with success does not entail that using fuzzy logic causes success. In the future, the technical limitations of fuzzy logic identified in this paper can be expected to become important in practice. Other general dilemmas of artificial intelligence work can also be expected to become critical—in particular, the issue of designing learning mechanisms that can solve the credit assignment problem when the simplifying features of present controllers are absent.

Acknowledgements. The author is grateful to several colleagues for useful comments on earlier versions of this paper, and to John Lamping for asking if Theorem 1 holds when equivalence is understood intuitionistically.

References

[Abbott, 1988] K. Abbott. Robust operative diagnosis as problem solving in a hypothesis space. In *Proceedings of the National Conference on Artificial Intelligence*, pages 369–374, 1988.

[Bourne *et al.*, 1991] J. R. Bourne *et al.* Organizing and understanding beliefs in advice-giving diagnostic systems. *IEEE Transactions on Knowledge and Data Engineering*, 3(3):269–280, September 1991.

[Burkhardt and Bonissone, 1992] David G. Burkhardt and Piero P. Bonissone. Automated fuzzy knowledge base generation and tuning. In *Proceedings of the IEEE International Conference on Fuzzy Systems*, pages 179–188, San Diego, California, March 1992.

[Collins and Michalski, 1989] A. Collins and R. Michalski. The logic of plausible reasoning: A core theory. *Cognitive Science*, 13(1):1–49, 1989.

[Dontas and Zemankova, 1990] K. Dontas and M. Zemankova. APPLAUSE: an implementation of the Collins-Michalski theory of plausible reasoning. *Information Sciences*, 52(2):111–139, 1990.

[Driankov and Eklund, 1991] Dimiter Driankov and Peter Eklund. Workshop goals. In *IJCAI'91 Workshop on Fuzzy Control Preprints*, Sydney, Australia, August 1991.

[Elman, 1990] Jeffrey L. Elman. Finding structure in time. *Cognitive Science*, 14(2):179–211, 1990.

[Gaines, 1983] Brian R. Gaines. Precise past, fuzzy future. *International Journal of Man-Machine Studies*, 19:117–134, 1983.

[Hekmatpour and Elkan, 1992] Amir Hekmatpour and Charles Elkan. A multimedia expert system for wafer polisher maintenance. Technical Report CS92-257, Department of Computer Science and Engineering, University of California, San Diego, 1992.

[Hekmatpour and Elkan, 1993] Amir Hekmatpour and Charles Elkan. Categorization-based diagnostic problem solving in the VLSI design domain. In *Proceedings of the IEEE International Conference on Artificial Intelligence for Applications*, pages 121–127, March 1993.

[Karim and Rivera, 1992] M. N. Karim and S. L. Rivera. Comparison of feed-forward and recurrent neural networks for bioprocess state estimation. *Computers and Chemical Engineering*, 16:S369–S377, 1992.

[Keller and Tahani, 1992] J. M. Keller and H. Tahani. Backpropagation neural networks for fuzzy logic. *Information Sciences*, 62(3):205–221, 1992.

[Kosko, 1990] Bart Kosko. Fuzziness vs. probability. *International Journal of General Systems*, 17(2–3):211–240, 1990.

[Lee, 1990] C. C. Lee. Fuzzy logic in control systems–parts 1 and 2. *IEEE Transactions on Systems, Man, and Cybernetics*, 20(2):404–435, March 1990.

[Mamdani and Assilian, 1975] E. H. Mamdani and S. Assilian. An experiment in linguistic synthesis with a fuzzy logic controller. *International Journal of Man-Machine Studies*, 7:1–13, 1975.

[McCarthy and Hayes, 1969] John McCarthy and Patrick J. Hayes. Some philosophical problems from the standpoint of artificial intelligence. In *Machine Intelligence*, volume 4, pages 463–502. Edinburgh University Press, 1969.

[Newquist, 1988] H. P. Newquist. Struggling to maintain. *AI Expert*, 3(8):69–71, 1988.

[Pearl, 1988] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann Publishers, Inc., 1988.

[Pedrycz, 1991] Witold Pedrycz. Fuzzy modelling: fundamentals, construction and evaluation. *Fuzzy Sets and Systems*, 41(1):1–15, 1991.

[Sammut and Michie, 1991] Claude Sammut and Donald Michie. Controlling a "black box" simulation of a space craft. *AI Magazine*, 12(1):56–63, 1991.

[Sandewall, 1989] Erik Sandewall. Combining logic and differential equations for describing real-world systems. In *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning (KR'89)*, pages 412–420, 1989.

[Smets, 1991] Philippe Smets. Varieties of ignorance and the need for well-founded theories. *Information Sciences*, 57–58:135–144, 1991.

[Sugeno *et al.*, 1989] Michio Sugeno *et al.* Fuzzy algorithmic control of a model car by oral instructions. *Fuzzy Sets and Systems*, 32(2):135–156, 1989.

[Sugeno, 1977] Michio Sugeno. Fuzzy measures and fuzzy integrals—a survey. In Madan M. Gupta, George N. Saridis, and Brian R. Gaines, editors, *Fuzzy Automata and Decision Processes*, pages 89–102. North-Holland, 1977.

[Tanaka and Sugeno, 1992] K. Tanaka and M. Sugeno. Stability analysis and design of fuzzy control systems. *Fuzzy Sets and Systems*, 45(2):135–156, 1992.

[Uomori *et al.*, 1990] Kenya Uomori, Autshi Morimura, Hirohumi Ishii, Takashi Sakaguchi, and Yoshinori Kitamura. Automatic image stabilizing system by full-digital signal processing. *IEEE Transactions on Consumer Electronics*, 36(3):510–519, August 1990.

[van Dalen, 1983] Dirk van Dalen. *Logic and Structure*. Springer Verlag, second edition, 1983.

[von Altrock *et al.*, 1992] C. von Altrock, B. Krause, and Hans J. Zimmermann. Advanced fuzzy logic control of a model car in extreme situations. *Fuzzy Sets and Systems*, 48(1):41–52, 1992.

[Yager, 1980] Ronald R. Yager. On a general class of fuzzy connectives. *Fuzzy Sets and Systems*, 4:235–242, 1980.

[Yager, 1992] Ronald R. Yager. Implementing fuzzy logic controllers using a neural network framework. *Fuzzy Sets and Systems*, 48(1):53–64, 1992.

[Yamakawa and Hirota, 1989] Special issue on applications of fuzzy logic control to industry. *Fuzzy Sets and Systems*, 32(2), 1989. Takeshi Yamakawa and K. Hirota (editors).

[Yamakawa, 1989] Takeshi Yamakawa. Stabilization of an inverted pendulum by a high-speed fuzzy logic controller hardware system. *Fuzzy Sets and Systems*, 32(2):161–180, 1989.

[Zadeh, 1973] Lotfi A. Zadeh. Outline of a new approach to the analysis of complex systems and decision processes. *IEEE Transactions on Systems, Man, and Cybernetics*, 3:28–44, 1973.