The Paradoxical Success of Fuzzy Logic*

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

November 1993

Abstract

Applications of fuzzy logic in heuristic control have been highly successful, but which aspects of fuzzy logic are essential to its practical usefulness? This paper shows that an apparently reasonable version of fuzzy logic collapses mathematically to two-valued logic. Moreover, there are few if any published reports of expert systems in real-world use that reason about uncertainty using fuzzy logic. It appears that the limitations of fuzzy logic have not been detrimental in control applications 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 used successfully in many real-world applications, but the coherence of the foundations of fuzzy logic remains under attack. Taken together, these two facts constitute a paradox. A second paradox is that almost all the hundreds or thousands of successful fuzzy logic applications are embedded controllers, while most of the thousands of theoretical papers on fuzzy methods deal with knowledge representation and reasoning. This paper attempts to resolve these paradoxes. More concretely, the aim of the paper is to identify which aspects

*Published in *IEEE Expert*, August 1994. An earlier version with the same title appeared in *Proceedings of the Eleventh National Conference on Artificial Intelligence* (AAAI-93), Washington, D.C., July 1993, MIT Press (pp. 698–703).

of fuzzy logic render it so useful in practice and which aspects are inessential. The conclusions here are based on a mathematical result, on a survey of the literature on the use of fuzzy logic in heuristic control and in expert systems, and on practical experience developing deployed expert systems.

This paper is organized as follows. First, Section 2 states and discusses the theorem mentioned above, which is that only two truth values are possible inside an apparently reasonable system of fuzzy logic. In an attempt to understand how fuzzy logic can be useful despite this potential collapse, Sections 3 and 4 examine the main practical uses of fuzzy logic, in expert systems and heuristic control. The 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 have remained 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 the logical connectives "and" $(\wedge)$, "or" $(\vee)$, and "not" $(\neg)$, degree of truth is defined as follows:

Definition 1: Let A and B be arbitrary assertions. Then

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

Depending how the phrase "logically equivalent" is understood, Definition 1 yields different formal systems. A system of fuzzy logic is intended to allow an indefinite variety of numerical truth values. However, for many notions of logical equivalence only two different truth values are possible given the postulates of Definition 1.

Theorem 1: Given the formal system of Definition 1, if $\overline{A \wedge \overline{B}}$ and $B \vee (\overline{A} \wedge \overline{B})$ are logically equivalent, then for any two assertions A and B, either $t(B) = t(A)$ or $t(B) = 1 - t(A)$. ■

A direct proof of Theorem 1 is given in the appendix. It can also be proved using similar results couched in more abstract form due to Dubois and Prade [1980; 1988]. The following result in particular is related.

Proposition: [Dubois and Prade, 1988] Let $\mathcal{P}$ be a finite Boolean algebra of propositions and let τ be a truth-assignment function $\mathcal{P} \rightarrow [0,1]$, supposedly truth-functional via continuous connectives. Then for all $p \in \mathcal{P}$, $\tau(p) \in \{0,1\}$. ■

The link between Theorem 1 and this proposition is that $\overline{A \wedge \overline{B}} \equiv B \vee (\overline{A} \wedge \overline{B})$ is a valid equivalence of Boolean algebra. Theorem 1 is stronger in that it only relies on one particular equivalence, while the proposition of Dubois and Prade is stronger in that it applies to any connectives that are truth-functional and continuous as defined in their paper.

The equivalence used in Theorem 1 is rather complicated, but it is plausible intuitively, and it is natural to apply in reasoning about a set of fuzzy rules, since $\overline{A \wedge \overline{B}}$ and $B \vee (\overline{A} \wedge \overline{B})$ are both re-expressions of the classical implication $A \rightarrow B$.[1] It was chosen for this reason, but the same result can also be proved using many other ostensibly reasonable logical equivalences.

It is important to be clear as to what exactly Theorem 1 says, and what it does not say. On the one hand, 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 given the equation $\overline{A \cap \overline{B}} = B \cup (\overline{A} \cap \overline{B})$, because Definition 1 can be understood as axiomatizing degrees of membership for fuzzy set intersections, unions, and complements.

On the other hand, the theorem does not necessarily apply to versions of fuzzy logic that modify or reject any of the 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 modeled by any operator in the Sugeno class [Sugeno, 1977], and when disjunction or conjunction are modeled by operators in the Yager classes [Yager, 1980].

Of course, the last postulate of Definition 1 is the most controversial one. In

[1] Note however that Theorem 1 does not depend on any particular definition of implication in fuzzy logic. New definitions of fuzzy implication are still being proposed as new applications of fuzzy logic are investigated, recently for example by Dubois and Prade [1992].

order to preserve a continuum of degrees of truth, one naturally wants to restrict the notion of logical equivalence. In intuitive descriptions, fuzzy logic is often characterized as arising from the rejection of the law of excluded middle: the assertion $A \vee \neg A$. Unfortunately, rejecting this law is not sufficient to avoid collapse to just two truth values. Intuitionistic logic [van Dalen, 1983] rejects the law of excluded middle, but the formal system of Definition 1 still collapses when logical equivalence means intuitionistic equivalence.[2] Of course, collapse to two truth values is avoided when one only admits the equivalences generated by the operators minimum, maximum, and complement to one. However, these equivalences are just the axioms of de Morgan, which only allow restricted reasoning about collections of fuzzy assertions.

3 Fuzzy logic in expert systems

The basic motivation for fuzzy logic is clear: many ideas resemble traditional assertions, but they are not 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. 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 uncertainty, and if so, which varieties of uncertainty. The question of interest here is more empirical: whether or not fuzzy logic is in practice an adequate formalism for uncertain reasoning in knowledge-based systems.

A thorough search of the technical literature, using the INSPEC and Computer Articles databases of over 1.3 million papers published since 1988, reveals no published report of an expert system that uses fuzzy logic as its primary formalism for reasoning under uncertainty and that is clearly stated (in the abstract available online) to be deployed and in routine use. Many theoretical papers on using fuzzy logic in expert systems have been published (see the volume edited by

[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 in the appendix can be carried over directly. Dubois and Prade [1988] note that if all the properties of a Boolean algebra are preserved except for the law of excluded middle, their proposition no longer holds. This observation is compatible with a collapse assuming only the equivalences of intuitionistic logic, because although intuitionistic logic rejects the law of excluded middle, it admits a doubly negated version of the law, namely $\neg\neg(\neg\neg A \vee \neg A)$.

Kandel [1992] for example) and several prototype systems have been described (by Graham [1991] for example) but it is hard to find reports of fielded systems doing knowledge-intensive tasks such as diagnosis, scheduling, or design.

Recent conferences give a representative view of the extent to which fuzzy logic is actually applied in current commercial and industrial knowledge-based systems. All the systems in actual use described at the 1992 IEEE International Conference on Fuzzy Systems are controllers as opposed to reasoning systems. No applications of fuzzy logic in knowledge-based systems were reported at the 1993 IEEE Conference on Artificial Intelligence for Applications. Fuzzy logic is used in some way in three of sixteen deployed systems described at the 1993 AAAI Conference on Innovative Applications of Artificial Intelligence: the CAPE, DODGER, and DYCE systems [Cunningham and Smart, 1993; Levy *et al.*, 1993; Pierson and Gallant, 1993]. However none of these systems uses the operators of fuzzy logic for reasoning about uncertainty. Input observations are assigned degrees of membership in fuzzy sets but inference with these degrees of membership uses other formalisms.

DYCE is one of several knowledge-based systems developed and fielded over the last five years by a team at IBM. Other systems deployed by this team are used for software and hardware diagnosis, for data analysis, and for operator training [Gallant and Thygesen, 1993; Hekmatpour and Elkan, 1993]. These systems have varying architectures and cope with different varieties of uncertainty. Experience with them suggests that fuzzy logic is rarely suitable in practice for reasoning about uncertainty. The basic problem is that the ways in which items of uncertain knowledge are combined must be carefully controlled to avoid incorrect inferences. Fixed, domain-independent operators like those of fuzzy logic do not work.

The correct propagation of degrees of certainty must take into account the content of the uncertain propositions being combined. This is necessary both when the uncertain propositions constitute shallow knowledge and when they constitute deep knowledge. In the case of shallow knowledge, which may be defined as knowledge that is valid only in a limited context (for example a correlation between a symptom and a fault), how degrees of uncertainty are combined must be adjusted to take into account unstated background knowledge.

A simple example shows what the difficulty is. Consider a system that reasons in a shallow way using a notion of "strength of evidence," and assume that as in many expert systems, this notion is left primitive and not analyzed more deeply. (Certainly "strength of evidence" is an intuitively meaningful concept which may or may not be probabilistic, but which is definitely different from "degree of truth.") For concreteness, suppose the context of discourse is a collection of melons, and in this context by definition $watermelon(x) \leftrightarrow redinside(x) \wedge$

$greenoutside(x)$. For some melon m, suppose that $t(redinside(m)) = 0.5$ and $t(greenoutside(m)) = 0.8$, meaning that the evidence that m is red internally has strength 0.5 and m is green externally with strength of evidence 0.8. Are the rules of fuzzy logic adequate for reasoning about this particular type of uncertainty? They say that the strength of evidence that m is a watermelon is $t(watermelon(m)) = \min\{0.5, 0.8\} = 0.5$. However, implicit background knowledge in this context says that being red inside and green outside are mutually reinforcing pieces of evidence towards being a watermelon, and m is a watermelon with strength of evidence over 0.5.

Deep knowledge can be defined as knowledge that is detailed and explicit enough to be general-purpose and valid in multiple contexts. Knowledge that is deep should be usable in complex chains of reasoning. However Theorem 1 says that if more than two different truth values are assigned to the input propositions of long chains of inference using the rules of fuzzy logic and one plausible equivalence, then it is possible to arrive at inconsistent conclusions. Fuzzy logic cannot be used for general reasoning under uncertainty with deep knowledge.

The fundamental 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, whatever the type of uncertainty represented by $t(\cdot)$. For example, in the case of probabilistic uncertainty the rule $t(A \wedge B) = t(A) \cdot t(B)$ is valid if and only if A and B represent independent events. In general, for probabilistic uncertainty all one knows is that $\max\{0, t(A) + t(B) - 1\} \leq t(A \wedge B) \leq \min\{t(A), t(B)\}$.

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 by 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, 1974]. 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 continues to meet resistance. For example, at the 1991 International 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. A fuzzy logic controller is embedded in the automatic transmission of the 1994 Honda Accord, but advertising brochures use the phrase "grade 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 a paper by 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.

- 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.[3]

- 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. Human engineers or learning algorithms can do this tuning; neural network methods have been especially successful [Keller and Tahani, 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.

The question which naturally arises is which of the 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 outputs 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.

[3]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 normally exist which map between numerical values and qualitative values. In fuzzy logic terminology, these components are said to defuzzify outputs and implement membership functions.

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 or a human tuned numerical values for those formalisms, as is customary when using fuzzy logic. A quote from the originator of fuzzy heuristic control is relevant here [Mamdani and Sembi, 1980]:

> . . . it should be remarked that the work on process control using fuzzy logic was inspired as much by Waterman and his approach to rule-based decision making as by Zadeh [1973] and his novel theory of fuzzy subsets.

The reference is to Waterman [1970].

Several knowledge representation formalisms that are rule-based and numerical have been proposed besides fuzzy logic. For example, well-developed systems were presented by Sandewall [1989] and Collins and Michalski [1989]. To the extent that numerical factors can be tuned in these formalisms, they should 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 has been studied by Tanaka and Sugeno [1992], and methodological problems in constructing models of complex systems based on deep knowledge have been considered by Pedrycz [1991]. Controllers with intermediate variables, thus with chaining of inferences, were investigated by 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. 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. 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, the models constructed tend to be highly application-specific and not generalizable [Bourne *et al.*, 1991].

Many recent expert systems, therefore, are neither rule-based nor model-based in the standard way [Hekmatpour and Elkan, 1993]. 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 that are 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 by 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 [Watrous and Shastri, 1987; Elman, 1990], and have already been used in control applications [Ku *et al.*, 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 to two-valued logic under an apparently reasonable condition, and despite the inadequacy of fuzzy logic for general inference with uncertain knowledge. The difficulties with fuzzy logic identified above have not been harmful in practice because current fuzzy controllers are far simpler than other knowledge-based systems. First, Theorem 1 is not an issue for present fuzzy controllers because they do not perform chains of inference, and they are developed informally, with no formal reasoning about their rules that applies equivalences such as the one used in the statement of Theorem 1. 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 general-purpose, the inadequacy of fuzzy logic for general reasoning about uncertainty 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, as fuzzy controllers are scaled up, the technical difficulties identified in this paper can be expected to become important in practice.

Theorem 1 is a crisp demonstration of one of several deep difficulties of scale in artificial intelligence: the problem of maintaining consistency in long sequences of reasoning. Other difficulties of scale 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 many colleagues for useful comments on earlier versions of this paper.

References

[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.

[Cunningham and Smart, 1993] Adam Cunningham and Robert Smart. Computer aided parts estimation. In *Proceedings of the Fifth Innovative Applications of Artificial Intelligence Conference*, pages 14–25, Washington, D.C., July 1993.

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

[Dubois and Prade, 1980] Didier Dubois and Henri Prade. New results about properties and semantics of fuzzy set-theoretic operators. In Paul P. Wang and S. K. Chang, editors, *Fuzzy Sets: Theory and Applications to Policy Analysis and Information Systems*, pages 59–75. Plenum Press, 1980.

[Dubois and Prade, 1988] Didier Dubois and Henri Prade. An introduction to possibilistic and fuzzy logics. In Philippe Smets *et al.*, editor, *Non-Standard Logics for Automated Reasoning*. Academic Press, 1988.

[Dubois and Prade, 1992] Didier Dubois and Henri Prade. Gradual inference rules in approximate reasoning. *Information Sciences*, 61(1–2):103–122, 1992.

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

[Gallant and Thygesen, 1993] George Gallant and Janet Thygesen. Digitized Expert PICTures (DEPICT): An intelligent information repository. In *Proceedings of the Fifth Innovative Applications of Artificial Intelligence Conference*, pages 50–60, Washington, D.C., July 1993.

[Graham, 1991] I. Graham. Fuzzy logic in commercial expert systems–results and prospects. *Fuzzy Sets and Systems*, 40(3):451–72, April 1991.

[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.

[Kandel, 1992] Abraham Kandel, editor. *Fuzzy expert systems*. CRC Press, Boca Raton, Florida, 1992.

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

[Ku *et al.*, 1992] C.-C. Ku, K.Y. Lee, and R.M. Edwards. Improved nuclear reactor temperature control using diagonal recurrent neural networks. *IEEE Transactions on Nuclear Science*, 39(6):2298–2308, 1992.

[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.

[Levy *et al.*, 1993] Arthur J. Levy *et al.* *Dodger*, a diagnostic expert system for the evaluation of nondestructive test data. In *Proceedings of the Fifth Innovative Applications of Artificial Intelligence Conference*, pages 107–117, Washington, D.C., July 1993.

[Mamdani and Sembi, 1980] E. H. Mamdani and B. S. Sembi. Process control using fuzzy logic. In Paul P. Wang and S. K. Chang, editors, *Fuzzy Sets: Theory and Applications to Policy Analysis and Information Systems*, pages 249–265. Plenum Press, 1980.

[Mamdani, 1974] E. H. Mamdani. Application of fuzzy algorithms for control of simple dynamic plant. *Proceedings of the Institution of Electrical Engineers*, 121(12):1585–8, 1974.

[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.

[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.

[Pierson and Gallant, 1993] Donald D. Pierson and George J. Gallant. Diagnostic yield characterization expert (DYCE): A diagnostic knowledge-based system shell for automated data analysis. In *Proceedings of the Fifth Innovative Applications of Artificial Intelligence Conference*, pages 152–160, Washington, D.C., July 1993.

[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.

[Waterman, 1970] D. A. Waterman. Generalization learning techniques for automating the learning of heuristics. *Artificial Intelligence*, 1:121–170, 1970.

[Watrous and Shastri, 1987] R. L. Watrous and L. Shastri. Learning phonetic features using connectionist networks. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, pages 851–854, 1987.

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

[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.

A Proof of Theorem 1

Theorem 1: Given the formal system of Definition 1, if $\overline{A \wedge \overline{B}}$ and $B \vee (\overline{A} \wedge \overline{B})$ are logically equivalent, then for any two assertions A and B, either $t(B) = t(A)$ or $t(B) = 1 - t(A)$.

Proof: Given the assumed equivalence, $t(\overline{A \wedge \overline{B}}) = t(B \vee (\overline{A} \wedge \overline{B})$. 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, which are

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

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 $t(B) = 1 - t(A)$. ∎