

Fuzzy Relations and Cognitive Representations

Christian Freksa

Fuzzy set theory and fuzzy logics were able to demonstrate impressive achievements in control theory and in technical applications already in the 1970s; but Lotfi Zadeh's great concern was - and still is - to demonstrate the power of his radically different approach to representing human concepts as a representational foundation in Artificial Intelligence. With his approach of introducing graded compatibility values to describe relations between concepts and real-world entities, formal systems can characterize states of affairs in terms of a manageable number of concepts - much like humans who describe the world by concepts that are qualified through linguistic hedges, prosodic emphasis or attenuation, and many other subtle ways of describing situations in a differentiated way to capture their essential significance in a concise manner.

Such mechanisms enable humans to summarize complex events in a meaningful way. Without the ability to drastically eliminate details of events, people would be incapable of dealing with the complexity of the world. With this insight, Lotfi Zadeh described in the 1970s [7] his personal grand challenge for Artificial Intelligence: the ability to automatically summarize the content of a paper or a book as capable humans can do it. Zadeh realized that this would be extremely difficult to achieve by the dominating approaches in AI, at the time; the Fuzzy Set approach, in contrast, has a built-in approach to generalize from specific instances and to ignore insignificant details.

27.1 Conceptual Framework and Cognitive Requirements

The principles of a Fuzzy Set provide a rich framework within which we can discuss semantic relations in arbitrary knowledge representation systems. The relations between discrete linguistic labels, their finite or infinite domains of support, and the infinitely-valued membership values provide an excellent basis for describing and discussing representation-theoretical issues. These involve the symbols of language, the mental concepts associated with these symbols, and the entities that the symbols refer to.

At the time when Lotfi Zadeh introduced his notion of a linguistic variable with graded compatibility values, the cognitive psychologist Eleanor Rosch published seminal papers on concept categories and on more or less prototypical representatives of concept categories [4], [4]. The anthropologist Brent Berlin and the linguist

Paul Kay described systems of color naming in different cultures and their interrelationships [1]. These are just two examples of human conceptualization that provide strong support for approaches of graded applicability of concepts and against “black or white” categories in cognitive systems.

Nevertheless, the notion of a ‘fuzzy concept’ is prone to considerable confusion and has not been universally well received. I would like to offer a representation-theoretic explanation why the notion of a fuzzy concept appears counterintuitive even to people who fully agree with gradual applicability of concepts to specific states of affairs. I will do so by comparing demands of complex technical systems with those of cognitive systems and I will refer to Stephen Palmer’s notion of a knowledge representation system [3].

In describing and controlling technical systems, we are dealing with high-dimensional closed worlds, in which the ranges of the control variables are known. Although we cannot precisely describe all relations between potential control values, we can refer to them on an individual basis. In cognitive systems, however, we are confronted with multi-dimensional open worlds containing too many feature dimensions to enumerate and with open variable ranges for most of them. In these systems, we are less concerned with describing control values than with selecting, characterizing, and relating salient features; furthermore, due to the open characteristics of cognitive systems, we cannot specify contexts of applicability of concepts, in most cases; these must be derived implicitly. As a consequence, in cognitive systems we have to relate and compare feature dimensions and feature values and focus on those which are salient in a specific domain. It is neither desirable nor feasible to subdue the concepts we use to describe situations by specifying ranges of applicability in their specific contexts of use. For example, to understand what someone means by an ‘expensive shop’, we do not need to know the types of goods it sells or the price ranges of these goods; it is sufficient to know that ‘expensive’ is a (linguistic) value on the higher end of a prize scale, that this value is in contrast to ‘inexpensive’ and ‘medium-prized’, and that these values would be ordered *inexpensive* < *medium-prized* < *expensive*. Thus, the meaning is not derived from the hypothetical feature values in the specific situations in which they are used. Rather, concepts derive an important part of their meaning from their relation to other concepts.

In describing the world around us, we are concerned with the actual features and their values, rather than hypothetical feature values as in technical systems in which we modify the state of affairs. As a consequence, for cognitive systems, we require structures that support the concise representation of characteristic patterns describing given situations in comparison to contrasting situations.

Even if we are convinced that the relation between a symbolic cognitive notion and a set of real-world or hypothetical entities is best characterized by a fuzzy relation, there is a question with regards to the nature and the representation of the fuzziness. For cognitive science this is an important issue, as the representations and processes need not only produce certain effects (this may be sufficient for technical applications), but they also serve to understand the details of the effects and their underlying mechanisms.

More specifically, we need to decide whether (1) fuzzy relations are or should be explicitly encoded as parts of mental concepts or whether (2) fuzzy relations will be caused by implicit effects of the perceptual and representational mapping characteristics between real-world entities and discrete mental concepts. For both alternatives we can conceive of representational and algorithmic implementations, some of which may appear cognitively more plausible than others. The choice of alternatives has important implications for the notion of a 'fuzzy concept' in cognition.

(1) implies that the mental concept itself is fuzzy. E.g., when I invoke the notion *tall*, I will invoke an entire set of more or less applicable potential instances (concept grounding); these potential instances form a part of the meaning of the concept; as some are more and some are less applicable, we obtain a fuzzy set of applicability. This is a quantitative interpretation of concepts as in classical fuzzy set theory; i.e., context of applicability and the degree of applicability of a concept in that context are part of the meaning of the concept.

(2) permits the abstract mental concept *tall* to be crisp. I.e. when I invoke the notion *tall*, I relate it to connatural notions like *short*, *medium-sized*, and *not tall* without extending these notions to real or potential instances; the meaning of the concepts is not derived by grounding them in the physical world but by relating them to comparable and contrasting concepts in the conceptual domain [2]; [6]. This is a qualitative relational interpretation of concepts. The meaning does not depend on a specific quantitative context, i.e., the relation between *short*, *medium-sized*, and *not tall* (e.g. an ordering relation) is invariant wrt. the specific reference set (e.g. tallness of professional basket ball players vs. first graders).

Fuzziness is not an issue on the abstract concept level, as the significance of concepts is in the relation to other concepts rather than in the direct relation to instances; in this view, concepts like *tall* do not have a quantifiable meaning on the concept level. Fuzziness may become an issue once we apply the relations between concepts to relations between instances; but on the level of relations between concepts, fuzziness is no longer an omnipresent property that cognitive systems have to deal with all the time. For example, if our real-world domain consists of three entities ordered by tallness, say 5, 6, and 7 feet tall, respectively, we will be able to apply the ordering relation of the corresponding concepts directly and characterize them as *short*, *medium-sized*, and *tall*, respectively, with no fuzziness entailed. A special charm of this relational view is the implicit context-adaptivity of the concepts involved.

At the end, the debate on whether or not we want to talk about 'fuzzy concepts' or about fuzzy relations between concepts and entities may boil down to the philosophical question which parts of a representational system we decide to call a 'concept': are concepts platonic crisp entities in the mind that allow us to think and dream about real and imagined objects and structures abstractly, or do concepts include relevant aspects of these objects and structures as well as the corresponding compatibility functions more concretely. I personally prefer to consider concepts in a qualitative and platonic way; the fuzziness gets introduced when I attempt to apply my network of concepts to real-world situations. I will explain my reasons in the following.

27.2 How Did I Arrive at the Field? – My Personal Background

As a high school student I developed a strong interest in cybernetics; this led me to enroll in physics, mathematics, and informatics at the Technical University of Munich. Concurrently with my undergraduate studies I had the great opportunity to join an interdisciplinary team of sleep researchers at the Max Planck Institute for Psychiatry in Munich as a computer programmer. In 1975, I was admitted as a Ph.D. student to the EECS Department at UC Berkeley where my first contact was Dr. Lawrence W. Stark, professor of physiological optics and engineering.

When I introduced myself to Dr. Stark's research group by presenting my work on real-time EEG sleep stage classification I discussed the issue of classifying boundary cases and how we dealt with them in our Munich team. Dr. Stark immediately advised me to talk to Professor Zadeh, the inventor of the fascinating Fuzzy Set Theory; he added that Zadeh had been confronted with much antagonism from colleagues who did not understand the value and importance of his contribution.

I talked to Professor Zadeh and became a regular participant of his weekly research seminar and his informal gatherings at the 3 C's café. He, in turn, became my Ph.D. advisor. In the year that I was admitted to Berkeley, the Sloan Foundation-funded Berkeley cognitive science program started with a regular highly interdisciplinary seminar series. Eleanor Rosch and Stephen Palmer from psychology, George Lakoff and Chuck Fillmore from linguistics, John Searle and Hubert Dreyfus from philosophy, Paul Kay from anthropology, and Lotfi Zadeh from computer science, among other brilliant Berkeley scientists were regular participants. This was the most outstanding and exciting academic environment I could imagine to foster my interests in artificial intelligence and cognitive science. Brilliant and eloquent minds from different fields publicly interacted as self-confident human beings.

My own research as a graduate student in Artificial Intelligence was heavily influenced by the cognitive science seminar, by courses in cognitive psychology that I took from Stephen Palmer and Eleanor Rosch, by Lotfi Zadeh's Seminar on Expert Systems, by an exciting computer science colloquium series, and by special panel discussions of the Bay Area Circle on Artificial Intelligence organized by Lotfi Zadeh. I became particularly fascinated by the insight that high-level cognition seems to work rather reliably on the foundation of severely underspecified knowledge. For example, people are pretty bad at geometrically reconstructing spatial environments, even the ones they feel most familiar with; nevertheless people rarely get lost in these environments.

Whereas the basic issue of imprecise, incomplete, and fuzzy knowledge appear closely related in the engineering and cognition domains, I developed a strong opinion that natural cognitive systems use knowledge in a rather different way than what our engineering approaches aim at. The fuzzy set approach aims at precisiating knowledge to resolve uncertainty, and to make meaning more precise in order to control a continuous space of options; in cognition, a large body of problems consists of identifying existing situations which form discrete islands in the large space of theoretically possible configurations. Continuous-valued fuzzy membership values are of theoretical importance to describe on the meta-level how concepts and

potential instances are related; but for actually matching concepts and objects on the object level, weaker approaches should suffice.

The basic approach I pursued was to characterize labels of 'fuzzy concepts' not by grounding them in terms of feature values, but in terms of relations to other labels that characterize the same or similar feature dimensions. For example, we use the labels *short*, *medium-sized*, and *tall* in arbitrary contexts in the same ordering relation $short < medium-sized < tall$; from a cognitive point of view, it is important that we can use feature dimensions (here: *size*) to establish categories in a universally comprehensible way; in most non-technical situations (in which we do not use labels in a strictly defined way), it is not necessary to consider boundaries between concepts - and we can still get the main idea across; conceptually neighboring labels will be correctly resolved through the reference context. In other words: although the labels engage in a fuzzy relation to entities in a space of theoretical feature values, we can use them abstractly to form categories without having to decide precisely which entities belong into which category: it is important to conceptually distinguish between a mountain and a valley and it is helpful to assume that the mountain starts where the valley ends; but we do not have to decide exactly where the valley ends and where the mountain starts: the boundary region simply is not of interest for these two concepts. The important point is that when we categorize entities along a given dimension and categorize concept labels along the same dimension, we will obtain coherent patterns which can be easily matched when we take into account neighborhood relations. This will be the case, if the network of abstract concepts structurally matches the network of corresponding features.

Fuzzy set theory and knowledge representation theory serve as excellent frameworks to characterize the cognitive agenda of knowledge representation and reasoning. The agenda involves abstract mental concepts and specific real or imagined entities in a space of theoretically possible feature values.

27.3 Issues and Lessons to Be Learnt

Lotfi Zadeh advised his students against doing their doctoral theses in the area of fuzzy set theory when this area was faced with hostility from other areas; he wanted to protect his students against reduced professional opportunities.

In my view, Professor Zadeh's most important contributions to science are not the now classical fuzzy set theory and fuzzy logic; it is the epistemological framework that permits relating human concepts and knowledge to a large variety of theories and formalisms and to a large variety of application domains. Zadeh's approach of generalizing from existing theories opened avenues for asking new questions regarding the epistemological status of notions related to uncertainty such as compatibility, similarity, fuzziness, vagueness, probability, possibility, and for discussing and comparing these notions.

The personal lesson I learnt from the confrontation with these notions was that we should look much more closely at specific problem domains and the precise questions we want to answer. We need to investigate in which ways we can represent the

epistemological features of interest and to what extent we can create conditions in which we do not need to represent them on the problem solving level. By creating a knowledge processing environment which will ensure a proper treatment of specific aspects of knowledge we may be able to avoid to make certain features explicit.

We may safely abstain from explicitly representing certain epistemological features if we employ *intrinsic* representations of crucial aspects in the sense of Palmer ([3]). By doing so, we can guarantee that certain properties 'automatically' hold due to structural properties of the representation employed. A vivid example for an approach in which we may safely neglect the explicit representation of important knowledge is the spatial domain: an architect who represents the spatial layout of a building by a geometric floor plan on a 2-dimensional sheet of paper does not have to make geometric knowledge explicit (e.g. 'The sum of the angles in a planar triangle equals 180°'). The structural properties of the representation medium implicitly will guarantee that the rule holds; or, expressed negatively: it will be impossible to represent a triangle in this medium which will violate the rule of the sum of angles in a triangle. Similarly, we will be able to find representations of concepts and suitable processes which will automatically yield fuzzy relations without the need of explicitly characterizing them. The perceptual domain appears to be a suitable domain to explore intrinsic representations of fuzziness. Lotfi Zadeh provided a suitable theory to make explicit what is going on in such representations on the epistemological level.



Fig. 27.1. Spontaneous gathering of friends and admirers with Lotfi Zadeh on the occasion of the congress “Wissensbasierte Systeme” (Knowledge-based Systems) in Munich, Germany, 29 October 1985. Photographer: anonymous.

27.4 Epilogue

After receiving my Ph.D. on representing fuzzy knowledge by means of discrete relations from Berkeley, I returned to Europe full of excitement about continuing my research in cognitive knowledge representation, a field that European researchers had not yet seriously approached. It took 15 more years before we were able to establish computer science-based cognitive science in Germany. Lotfi Zadeh recognized my early frustration regarding lack of support for my research initiatives after leaving Berkeley. In 1981 he wrote to me: *My advice to you is to accept as a fact of life that you're in a conservative environment. I have no doubt, though, that eventually your ideas will prevail and receive the recognition they deserve. In short, don't give up your efforts no matter what.*

Thirty years later I can say, this was an excellent advice. As always.

Acknowledgement. I acknowledge generous support for the Spatial Cognition Research Center SFB/TR 8 at the Universities of Bremen and Freiburg by the German Research Foundation. I thank Thora Tenbrink for critical comments on this paper.

References

1. Berlin, B., Kay, P.: Basic Color Terms: Their Universality and Evolution. University of California Press, Berkeley (1969)
2. Freksa, C.: Communication about visual patterns by means of fuzzy characterizations. In: XXIIInd International Congress of Psychology, Leipzig (1980)
3. Palmer, S.E.: Fundamental aspects of cognitive representation. In: Rosch, E., Lloyd, B.B. (eds.) Cognition and Categorization, pp. 259–303. Lawrence Erlbaum, Hillsdale (1978)
4. Rosch, E.: On the internal structure of perceptual and semantic categories. In: Moore, T.E. (ed.) Cognitive Development and the Acquisition of Language, Academic Press, New York (1973)
5. Rosch, E.: Principles of categorization. In: Rosch, E., Lloyd, B.B. (eds.) Cognition and Categorization, pp. 27–48. Lawrence Erlbaum, Hillsdale (1978)
6. Tenbrink, T., Freksa, C.: Contrast sets in spatial and temporal language. *Cognitive Processing* 10, 322–324 (2009)
7. Zadeh, L.A.: Personal communication (1970s) (citation missing)