COPING WITH THE IMPRECISION OF THE REAL WORLD

An Interview with Lotfi A. Zadeh

The tools researchers use to probe certain AI problems, says this Berkeley professor, are sometimes too precise to deal with the "fuzziness" of the real world.

Q. **Professor Zadeh, in this interview today, you agreed to talk mainly about the limits of traditional logic in dealing with many of the problems in the field of artificial intelligence (AI) and your approach toward helping to overcome those difficulties. Before getting into those issues, though, could you first give our readers a brief overview of what you see as the major areas for computer applications in the years ahead?**

ZADEH. In the years ahead, there will be three major areas of computer applications. One, in the traditional vein, is the use of computers for purposes of numerical analysis. Numerical analysis will be very important in a number of fields—particularly in scientific computations and simulation of large-scale systems.

For such purposes, there will be a need for larger and larger computers. This is especially true for applications in meteorology, in nuclear physics, in modeling of large-scale economic systems, in the solution of partial differential equations, and in the simulation of complex phenomena like turbulence, fluid flow, etc.

Area number two will be concerned with masses of data—large databases. This is the sort of thing that is playing and will be playing an important role in banking, insurance, records processing, information retrieval, etc. What will be important in these areas is not so much number-crunching capabilities as the capability to store massive amounts of data and to access whatever data are needed rapidly and at a reasonably low cost.

Furthermore, in these areas, computer networking, of course, will be playing an essential role. For you will have to have access not just to a single database but to a collection of interconnected databases. In response to this need, we will see many advances in computer networking during the next several years.

The third major area for computer applications is what has come to be known as *knowledge engineering.* This area has received considerable publicity during the past few years, particularly since the Japanese have highlighted it as an area of prime concern. This is a rapidly growing field in terms of importance and breadth of applications.

Knowledge engineering is one of the major areas of AI. And within knowledge engineering, a field of primary importance is that of expert systems. True, there may be exaggerated expectations of what expert systems can accomplish at this juncture, but as Jules Verne observed at the turn of the century, scientific progress is driven by exaggerated expectations.

So we have these three major areas for computer applications in the years ahead. All will be *growing* in importance. But knowledge engineering, I think, will be growing in importance more rapidly than the other two, because it is the youngest and, in a sense, the most pervasive of the three.

In saying that knowledge engineering is going to become very important, I don't want to imply that the other two will become less important. They will become more important also. But in relative terms, knowledge engineering will certainly be much more important than it is today.

Now, what I'm going to say will relate to this third area, rather than the first two.

Q. **Do you see supercomputers playing an important role only in the area of numerical analysis?**

ZADEH. Supercomputers pertain to all three areas: *numerical analysis, large databases,* and *knowledge engineering.* But they apply primarily to the first area: numerical analysis. There is at this point some controversy as to how the available research funds should be distributed between the efforts to build supercomputers and to build machines that will be AI oriented.

These are somewhat distinct efforts. The Japanese are pushing both of them. And in the United States, the emphasis on AI-oriented types of computers is just beginning to become strong, largely as a reaction to the Japanese effort. As you may know, Edward Feigenbaum of Stanford University is a leading advocate of the establishment of a U.S. National Center for Computer Technology as a rallying point for the U.S. effort.

Q. I'm not perfectly clear on the distinction between a supercomputer and an AI-oriented device, which I believe is sometimes referred to as a fifth generation computer. Could you clarify?
ZADEH. Fifth generation has become somewhat of a misnomer. Basically, the Fifth Generation Project is perhaps the most publicized of the several ambitious research programs undertaken by the Japanese. Another one of those programs is called the Supercomputer Project.

The Fifth Generation Project is not hardware oriented. Rather it is intended to exploit the advances in hardware that might be achieved under the other projects in the overall program.

For supercomputers, the emphasis is on large-scale computations relating to scientific and technological applications. For AI-oriented computers, on the other hand, the emphasis is shifted away from "data" to "knowledge."

What matters in the case of these AI-oriented computers is their ability *to infer* from the information resident in a large knowledge base—especially when this information is imprecise, incomplete, or not totally reliable.

We can expect that advances in the design of supercomputers—on both the hardware and the software levels—will have a significant impact on the architecture of AI-oriented machines.

Of course, the supercomputer is not a unique concept. It can take a variety of forms. Parallel processing may play a major role in AI-oriented applications—especially in pattern recognition, natural language processing, and inference from large knowledge bases. Whether this will actually happen within the next decade is a matter of conjecture.

During the past 20 years, parallel processing has been a highly promising area for research. Yet what could actually be accomplished with parallel processing has always lagged behind expectations. At this juncture, however, it's possible that important breakthroughs may be around the corner.

But there are many problems in AI that will not be helped to an appreciable extent by the availability of supercomputers, whether they will be von Neumann-type supercomputers or some other kind. The reason why this is so is because the limitation is not so much computing power, but our lack of understanding of some of the processes required to perform even simple cognitive tasks.

Lotfi A. Zadeh, a professor of electrical engineering and computer science, has been a member of the faculty at the University of California at Berkeley since 1959. Born in the Soviet Union of Iranian parents, Zadeh came to the U.S. in 1944 and studied at MIT and Columbia, from which he got his Ph.D. in 1949. After serving on the Columbia faculty during the 1950s, Zadeh went to Berkeley. Dissatisfied with the use of very precise mathematics to describe the sometimes highly imprecise real world, Zadeh, during the 1960s, developed the theory of fuzzy sets. In recent years, his controversial ideas on "fuzzy thinking" have begun to win a following in some quarters of the world.

Q. You say that there are many problems that won't be solved by the availability of supercomputers. Could you give an example of such a problem?
ZADEH. Let me start with a problem basic to most other problems: the problem of summarization. Now, when I talk about summarization, I am not talking about summarizing a short stereotypical story. That capability we have, thanks to Roger Schank, his associates at Yale University, and others within the AI community.

But what we have no understanding of, whatsoever, is how to summarize a nonstereotypical story that is not a short story.

Q. Could you give an example of a stereotypical story?
ZADEH. An example would be accounts of automobile accidents. They tend to be stereotypical. In other words, in the story there is an indication of what kind of accident it was, when it occurred, where it occurred, whether there were injuries, etc.

Q. I see. It has a predetermined structure?
ZADEH. Yes. When I say stereotypical, I mean it has a predetermined structure. So if you have a predetermined structure, then you can understand the story and you can summarize it.

Now, the reason why summarization is so difficult—it is far more difficult than machine translation from one language to another language—is because summarization requires understanding. And the ability of a computer to summarize a short, stereotypical story would be a little bit like the ability of a person who sees a story of that kind in the newspaper and doesn't understand completely the language in which the story is written. Nonetheless, he can discern a few words here and there and, on that basis, summarize the story. Many people could do that if they have some minimal competence in the language in which the story is written. But that minimal competence is completely inadequate when it comes to summarizing something that is not stereotypical and not short.

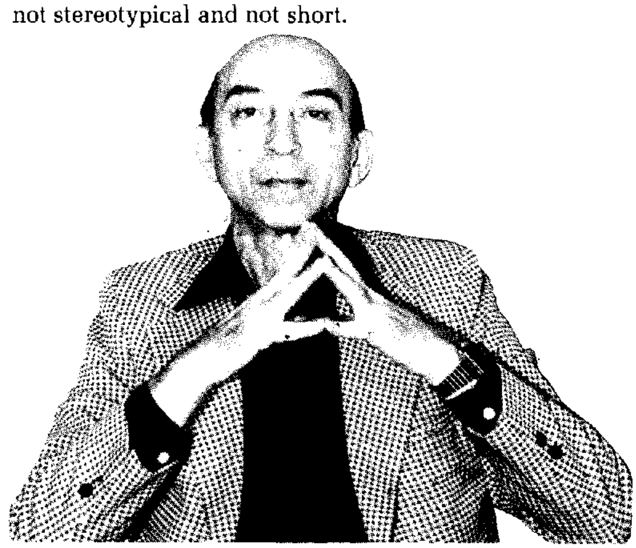

Q. You said that it is very difficult and often impossible at present to write a program to summarize a story. So why is that an important point?
ZADEH. The ability to summarize is an acid test of the ability to understand, which in turn is a test of intelligence and competence. Suppose I asked a person not familiar with mathematics to summarize a paper in a mathematical journal. It would be impossible for him to summarize it, because he doesn't understand what that paper is about, what the results are, what the significance is, and so forth.

So in a situation like that, it wouldn't help us to have a supercomputer. It wouldn't help us to have all the supercomputers in the world put together. That's not where the problem lies.

Q. Professor Zadeh, could you give me a few other examples of problems that will still defy solution even with major advances in supercomputer design?
ZADEH. Take the problem of identification of ethnic origin. Humans can do that. You can look at a person and say, "Gee, he looks Irish," or whatever. Now, it would be impossible at this point to write a program that would look at somebody's picture and identify the ethnic origin of that person. I don't want to say that it's impossible period. I'm merely saying that at this moment it's impossible.

Another problem is estimation of age. Assume that you look at somebody, and you say, "Well, this person must be around 35." Again, we cannot write a program that would enable a computer to do it at this point. And we can't put our finger on subtle differences, like the difference between a person who is 20 years old and somebody who is 10.

Q. Why is that? Why is it impossible to write a program to estimate a person's age from an analysis of physical features?
ZADEH. Because we don't understand too well how we arrive at assessments of that kind. In other words, in order to write a program, we have to have an understanding of how we do it. The limitation in problems of this kind is that we cannot articulate the rules that we employ subconsciously to make that kind of an assessment.

Of course, you know, there are certain things that might not be so difficult. It's not so difficult to differentiate between a person who is 70 years old and a person who is, for instance, 5 years old. But I'm talking about kinds of problems in the estimation of age that are not as trivial as that. And we know that you cannot base it entirely on wrinkles or color of hair. It's the totality of these things put together that enable us to make an assessment.

Still another problem of that same type is the problem of the identification of a musical tune. People can identify a tune if they hear just a few bars. They generally can guess who the composer is, even though they may never have heard this piece before. In other words, there's something about the way the music composed by a particular composer sounds that makes it possible for us to say, "Well, this is Mozart," even though we may never have heard that piece before.

Again, if somebody asked how you guessed that it was Mozart, you would not be able to put down on paper the criteria that you have employed.

Q. You're saying that this process of recognizing that a short burst of music is from Mozart, Beethoven, or someone else is something going on unconsciously or intuitively?
ZADEH. It's something that we can do without being able to articulate the rules. In other words, the decisional algorithms that we employ for this purpose are opaque rather than transparent.

The problems that I mentioned—ethnic origin identification, age estimation, composer identification, or tune recognition—all of these are problems in pattern recognition. Many of these problems are far from solution at this point. And they present right now a stumbling block to such applications as speech recognition for connected speech.

Q. What is the point of all these examples? What is the lesson here?
ZADEH. What I have said so far is intended merely to give the reasons for my feeling that the availability of supercomputers will not help us much in solving problems of that kind. This is the issue that I was really addressing myself to.

But at this point, all I'm trying to say is that the supercomputer effort and the AI-oriented type of computer effort to develop machines that can perform nontrivial cognitive tasks are not quite the same. There is some interaction between them, but they are qualitatively different.

Let us return now to area number three, *knowledge engineering.* What I have to say here will be at considerable variance with the widely held positions within the AI community.

AI, as we know it today, is based on two-valued logic—that is, the classical Aristotelian logic. And it is generally assumed that all you need as a foundation of AI is first-order logic.

Q. You say that AI today is based on two-valued logic. Is that a yes-or-no kind of logic?
ZADEH. Yes, that's right. It's a yes-or-no kind of logic. Actually, two-valued logic encompasses a variety of logical systems, all of which share the basic assumption that truth is two-valued. One of these logical systems is what is called first-order logic. And so the assumption that many people make is that first-order logic, perhaps with some modifications, is sufficient.

Q. Does a simple example come to mind of first-order logic?
ZADEH. Well, suppose you say, "All men are mortal. Socrates is a man. Therefore Socrates is mortal." This would be a very simple example of reasoning in first-order logic.

Briefly, within the AI community at this point, there are two camps. One camp, the conservative camp, takes the position that AI, and more generally knowledge engineering, should be based on logic, and in particular on first-order logic. One of the main proponents of this view is John McCarthy of Stanford University. Other prominent proponents include Nils Nilsson of SRI, Wolfgang Bibel of the University of Munich, Robert Kowalski of London, and Alain Colmerauer of Marseilles, France.

Now the other camp takes the position that logic is of limited or no relevance to AI. They believe that first-order logic is too limited to be able to deal effectively with the complexity of human cognitive processes. Instead of systematic, logical methods, this second camp relies on the use of ad hoc techniques and heuristic procedures. The prime exponents of this position are Roger Schank of Yale University and, more recently and less emphatically, Marvin Minsky of MIT.

Q. How does your position, Professor Zadeh, differ from those in the two AI camps you mentioned—the conservatives who believe in first-order logic and the other camp that believes logic has only limited relevance?

ZADEH. The position that I take—and this is really what differentiates me from most of the people in AI—is that we need logic in AI. But the kind of logic we need is not *first-order logic*, but *fuzzy logic*—that is, the logic that underlies inexact or approximate reasoning.

I feel this way because most of human reasoning—almost all of human reasoning—is imprecise. Much of it is what might be called common sense reasoning. And first-order logic is much too precise and much too confining to serve as a good model for common sense reasoning.

The reason why humans can do many things that present-day computers cannot do well or perhaps even at all is because existing computers employ two-valued logic.

To put it another way, the inability of today's computers to solve some of those problems I mentioned earlier is not that we don't have enough computing capacity. Rather, the computers we have today—in terms of both hardware and software—are not oriented toward the processing of fuzzy knowledge and common sense reasoning. This is where the problem lies in my view.

Q. Doesn't that view—the idea that computers don't mimic human thought processes very well—imply that there is something sacred about the way human beings think? Isn't it possible that the way humans think about things is not very good and that it might be possible to conceive an artificial way of thinking that is superior?

ZADEH. Of course, one could take the position, as some workers in AI do, that it is not essential to mimic the human mind in the design of AI systems. One argument is that when we design an aircraft, so goes one of the arguments, we don't design it like a bird. But somehow this plane-design analogy doesn't seem to pertain to the design of AI systems. For when we actually attempt to build AI systems that can perform humanlike cognitive tasks we invariably seem to come back to the human model. The human model is a pretty good one—better than many people thought.

Q. Could you elaborate a bit more on just what fuzzy logic is?

ZADEH. Let me do that. I will explain the main differences between fuzzy logic and classical two-valued logic. In classical two-valued systems, all classes are assumed to have sharply defined boundaries. So either an object is a member of a class or it is not a member of a class.

Now, this is okay if you are talking about something like mortal or not mortal, dead or alive, male or female, and so forth. These are examples of classes that have sharp boundaries.

But most classes in the real world do not have sharp boundaries. For example, if you consider characteristics or properties like tall, intelligent, tired, sick, and so forth, all of these characteristics lack sharp boundaries. Classical two-valued logic is not designed to deal with properties that are a matter of degree. This is the first point.

Now, there is, of course, a generalization of two-valued logic. And these generalized logical systems are called multivalued logics. So in multivalued logical systems, a property can be possessed to a degree.

Q. I'm not perfectly clear here. Consider the word "tall." Are you saying "tall" can take on multiple values?

ZADEH. Yes, tallness becomes a matter of degree, as does intelligence, tiredness, and so forth. Usually you have degrees between zero and one. So you can say, for example, that a person is tall to the degree 0.9. These degrees are grades of membership that may be interpreted as truth values.

In classical logic, there are just two truth values: true/false (or one and zero). In multivalued logical systems, there are more than two truth values. There may be a finite or even an infinite number of truth values, that is, an infinite number of degrees to which a property may be possessed.

In a three-valued system, for instance, something can be true, false, or on the boundary. Or you can have systems in which one has a continuum of truth values from zero to one.

Q. Who first developed multivalued logic?

ZADEH. The person best known in that connection is a Polish mathematician by the name of J. Lukasiewiecz. He first developed the concept of multivalued logic during the 1920s.

Q. Could you give an example or two of a situation that requires multivalued logic?

ZADEH. Well, you would need a multivalued system to be able to say something like "John is tall." For tall is a property that requires an infinity of truth values to describe it. So something as simple as "John is tall" would require multivalued logic—unless you arbitrarily establish a threshold by saying "somebody over 6 feet tall is tall, and those who are less than 6 feet tall are not tall." In other words, unless you artificially introduce some sort of a threshold like that, you will need multivalued logic.

But even though these multivalued logical systems have been available for some time, they have not been used to any significant extent in linguistics, in psychology, and in other fields where human cognition plays an important role.

Q. Why hasn't multivalued logic been used?
ZADEH. The reason such systems haven't been used is that multivalued logic doesn't go far enough. And this is where *fuzzy logic* enters the picture.

What differentiates fuzzy logic from multivalued logic is that in fuzzy logic you can deal with fuzzy quantifiers, like "most," "few," "many," and "several."

Fuzzy quantifiers have something to do with enumeration, that is, with counting. But they are fuzzy because they don't give you the count exactly, but fuzzily. For instance, you say "many" or "most."

In multi-valued logic you have only two quantifiers, "all" and "some," whereas in fuzzy logic you have all the fuzzy quantifiers. This is one of the important differences.

Q. Now are those the only fuzzy quantifiers?
ZADEH. Well, the ones that I mentioned are merely examples. In reality, there is an infinite number of fuzzy quantifiers. For example, you can say "not very many," "quite a few," "many more than 10," "a large number," "many," "few," or "very many." There is an infinite number of ways in which you can describe in an approximate fashion a count of objects.

Q. Besides the fuzzy quantifiers, what else distinguishes fuzzy logic from multivalued logic?
ZADEH. Another key difference is that in fuzzy logic *truth* itself is allowed to be fuzzy. So it is okay to say that something is "quite true." You can say "it's more or less true." You can also use fuzzy probability like "not very likely," "almost impossible," or "rarely." In this way, fuzzy logic provides a system that is sufficiently flexible and expressive to serve as a natural framework for the semantics of natural languages.

Furthermore, it can serve as a basis for reasoning with common sense knowledge, for pattern recognition, decision analysis, and other application areas in which the underlying information is imprecise. Within the restricted framework of two-valued and even multivalued systems, these problem areas have proved to be difficult to deal with systematically.

The crux of the problem, really, is the excessively wide gap between the precision of classical logic and the imprecision of the real world.

Q. Does fuzzy logic, then, provide a good match with the imprecise real world?
ZADEH. I don't wish to imply that fuzzy logic is in any sense an ultimate system. I do believe, however, that it is far better suited for dealing with real-world problems than the traditional logical systems.

Q. At this juncture, Professor Zadeh, could you encapsulate what you feel is your most important point so far?
ZADEH. Yes. It is that the limitation in knowledge engineering is not the unavailability of supercomputers. But it is the fact that computers—both their hardware and software—are based on a kind of logic that is not a good model for human reasoning.

Ultimately, the problem lies at the hardware level. For computers are basically digital devices: they deal with discrete bits of information. Fuzzy information, on the other hand, is not discrete. With fuzzy information, one thing merges into another.

Now an important point is this: even though present-day computers are based on two-valued logic, they can be programmed to process fuzzy information using fuzzy logic. But doing this does not represent an efficient use of the computational capabilities of present-day computers.

Q. You're saying there are problems with both the hardware and the software of existing computer systems, but that you can overcome them so that traditional computers can still handle fuzzy information?
ZADEH. Yes, you can overcome the hardware limitations of current computers with software. But doing that involves an inefficient use of computers.

The ability of the human mind to reason in fuzzy terms is actually a great advantage. Even though a tremendous amount of information is presented to the human senses in a given situation—an amount that would choke a typical computer—somehow the human mind has the ability to discard most of this information and to concentrate only on the information that is task relevant. This ability of the human mind to deal only with the information that is task relevant is connected with its ability to process fuzzy information. By concentrating only on the task-relevant information, the amount of information the brain has to deal with is reduced to a manageable level.

Q. So you can work with traditional computers, but they are inefficient because they are fundamentally incompatible with fuzzy information. What kind of computer could you use then?
ZADEH. At some point, we may be able to conceive computers that are radically different from existing computers in that the operations they perform are rooted in fuzzy rather than two-valued logic. In other words, they may need a different kind of hardware.

There has been talk about "chemical," "biological," or "molecular" computers.

Some imaginative thinkers are talking about computers of that kind, but we don't have them yet.

Q. What is a molecular computer?
ZADEH. It's difficult to say. But if you compare the way the human brain works with the way a modern computer works, I think that you will find there are some fundamental differences. The human brain, in a way that we don't understand too well at present, uses fuzzy logic.

So the hardware—if you may call it that—of the human brain is the kind of hardware that is effective for manipulating imprecise information. When one uses the term molecular computer, or biological computer or chemical computer, what one has in mind is something that approximates the way in which the human brain processes information. And therein lies a fundamental challenge: how to develop a better understanding of how the human brain processes this fuzzy information so effectively.

Q. How much is presently understood about how the brain processes information?
ZADEH. Scientists know a lot about the functioning of the brain at the neuron level. But how does the activity at that level aggregate into thinking processes? Trying to understand the brain at the neuron level is like trying to understand the functioning of a telephone system in a large city by examining the wiring of a telephone set. We can understand something on the microlevel but are unable to integrate that into an understanding of functioning on higher levels.

Q. I don't completely understand your reasons for writing off traditional computers. Maybe people haven't tried hard enough or long enough, as yet, to make them work for certain AI applications?
ZADEH. In fact, I'm not writing off traditional computers in regard to their ability to process fuzzy information. Rather, my position is that we do not have a good understanding at this point as to how to use them efficiently for handling fuzzy information.

In fact, I believe that there will be a growing number of applications of fuzzy logic in a wide variety of fields using present-day computers. But, ultimately, to achieve a higher level of efficiency, it may be necessary to employ computers that are specially designed for dealing with fuzzy information.

To give one example, speech recognition is a problem far from a satisfactory solution. We do have speech recognition systems with limited capability to understand speech. But all of these systems do not scale up. That is, they can not be merely modified and improved in an evolutionary way until they come close to the human ability to understand speech. So it's obvious that what is needed is an altogether different approach.

Q. I'm not perfectly clear when you say these systems won't scale up. Could you expand on that point?
ZADEH. Yes. In some situations we have a system that has a limited capability. But we see clearly how, by improving that system, we can raise its level of performance to a point where it can compete with humans in terms of certain abilities.

Now, within AI, this is generally not the case. That is, many AI systems do not scale up: They reach very quickly the limit of their ability. In other words, you cannot push them beyond that point.

Q. Could you give an example of a system that does not scale up?
ZADEH. A good example, to go back to what I said previously, are the programs that can summarize. These programs reach the limit of their ability very quickly—in terms of the length of the story they can summarize or in terms of the degree to which the story is nonstereotypical. And you cannot go beyond those limits without radically altering the approaches used.

Another example. Before we had integrated circuits, we depended on vacuum tubes in the design of our computers, and you could push the capabilities of those early computers only up to a point. We had to come up with something new, the concept of an integrated circuit, and eventually very large-scale integration (VLSI). Those breakthroughs greatly increased our capability to compute, to store, and more generally, to process information.

That was a situation that called for something radically different. And it wasn't a matter of evolution—but of revolution.

And so I think that we are faced with a somewhat similar situation in the case of computers that can perform high-level cognitive tasks. That is, we cannot hope to be able to solve the problems of the kind I mentioned earlier, in particular the problem of summarization, through evolutionary improvements in present computer hardware or software.

Q. In other words, some of the more difficult problems in AI won't be solved by innovations in supercomputer architecture—innovations like parallel processing?
ZADEH. That's right, such innovations aren't going to help much.

But there is more to this fuzzy logic than simply the enhancement of the ability of computers to solve various problems—to perform nontrivial cognitive tasks. Accepting fuzzy logic will also call for a certain fundamental shift in attitudes, particularly in theoretical computer science. At this point theoretical computer science is mathematical in spirit, in the sense that it is oriented toward the discovery and proof of results that can be stated as theorems.

Unfortunately, there is an incompatibility between precision and complexity. As the complexity of a system increases, our ability to make precise and yet nontrivial assertions about its behavior diminishes. For example, it is very difficult to prove a theorem about the behavior of an economic system that is of relevance to real-world economics.

What I anticipate in the future is a growing recognition of the necessity to find an accommodation with the pervasive imprecision of the real world. This change is needed to be able to make assertions that are not just nontrivial theorems, but something of relevance to practice. In computer science today, people use two-valued logic to establish certain results. But such results are often limited in their relevance to the real world—because they are excessively precise. In other words, we have to accord acceptance to assertions that do not adhere to high standards of precision.

This accommodation with imprecision will require the use of fuzzy logic. Gradually and perhaps rather slowly, there will be a growing acceptance of fuzzy logic as a conceptual framework for computer science.

Now, it is a little bit more difficult to articulate this particular position than some of the earlier things that I said. For this gets into issues that relate not just to computer science but, more generally, to science itself. Science at this point is based on two-valued logic. So what I'm talking about is a significant shift in attitude, not just in computer science, but more generally in scientific thinking.

At this point there is a long-standing and deep-seated tradition of according respectability to what is mathematical and precise. We may have to retreat from this tradition in order to be able to say something useful about complex systems and in particular about systems in which human reasoning plays an important role.

Q. Okay, then, there have been instances in the past where scientists have been too preoccupied with mathematics and precision and, as a result, have failed to come up with useful results. Does an example come to mind?

ZADEH. Yes. Take economics. Time and again, it has been demonstrated that what actually happens in the realm of economics is very different from what the experts predicted. These experts might be using large-scale econometric models, sophisticated mathematics, large-scale computers, and the like. Despite all that, the forecasts turn out to be wrong—very wrong.

Why? Two reasons. One is that economic systems are very complex. Second, and more important, human psychology plays an essential role in the behavior of such systems. And this complexity, together with human reasoning, makes the classical mathematical approaches, based on two-valued logic, ineffective.

So, again, to approximate the way humans can sort through large masses of data and arrive at some sort of a qualitative conclusion, it might be necessary to use fuzzy logic.

Q. Has fuzzy logic been able to solve some of the difficult problems in AI you mentioned earlier? Or is it still just a promise?

ZADEH. These problems are intrinsically complex, and fuzzy logic by itself does not provide a solution to them. Rather, it merely enhances our ability to do so without guaranteeing success. It's a little like finding a cure for cancer. You may develop a technique that may help in finding a cure but it doesn't guarantee a cure will be found.

Fuzzy logic, then, is a necessary but not sufficient condition to finding solutions to these problems. It is a tool that enhances our ability to deal with problems that are too complex and too ill-defined to be susceptible to solution by conventional means. It will be an ingredient of the tools that will eventually be used to solve these problems.

Q. Have you made any headway in persuading people that they needn't always be superprecise, that in fact such an approach may be an inappropriate approach for attacking certain types of problems?

ZADEH. It will be a slow process. It's not very easy to change some of the basic attitudes people have been educated with, like the attitude that we must be very precise and that we have to try to come up with results that can be stated as theorems. It's difficult to change these attitudes.

Let me draw an analogy with the way people dress. Classical logic is like a person who comes to a party dressed in a black suit, a white, starched shirt, a black tie, shiny shoes, and so forth. And fuzzy logic is a little bit like a person dressed informally, in jeans, tee shirt, and sneakers. In the past, this informal dress wouldn't have been acceptable. Today, it's the other way around. Somebody who comes dressed to a party in the way I described earlier would be considered funny.

Changes in attitude may take place not only in dress but also in science, music, art, and many other fields. And, in science, there may be an increasing willingness to realize that excessively high degree of formalism, rigor, and precision is counterproductive.

Freedom of expression in science could exhibit itself as a movement away from two-valued logic and toward fuzzy logic. Fuzzy logic is much more general and it gives you much more flexibility.

Q. How long will it take for traditional scientific attitudes about precision to change and fuzzy logic to take hold?

ZADEH. Well, I think it will take something on the order of perhaps a couple of decades. Fuzzy logic is making inroads, but it is not something that has coalesced into a broad movement. In other words, there are pockets. These pockets exist in various fields, and of course, there are some people who view these pockets with suspicion and hostility—just as some people who are conservative look with suspicion on those who dress informally.

The difficulty of persuading people has to do also with the question of where does respectability lie. Traditionally, respectability went along with being more mathematical, more precise and more quantitative. And these attitudes go back to Lord Kelvin who said that it's not really a science if it's not quantitative.

But fuzzy logic now challenges that. There are many things that cannot be expressed in numbers, for example, probabilities that have to be expressed as "very likely," or "unlikely." Such linguistic probabilities may be viewed as fuzzy characterizations of conventional numeric probabilities.

And so in that sense fuzzy logic represents a retreat. It represents a retreat from standards of precision that are unrealistic.

There are many parallels to that sort of thing in the history of human thought, where people didn't realize that the objectives they set were unrealizable.

Q. Does an example or two come to mind of a situation where scientists had to retreat from standards of precision that were not attainable?
ZADEH. Well, a good example of that sort of thing is statistical mechanics. People in the beginning of the nineteenth century were firm believers in the possibility of using the mechanics that were developed at that time by people like Lagrange and applying those mechanics to the solution of all sorts of problems involving the motion of bodies. But then they encountered the "two-body," "three-body," and "n-body" problems, and it became clear that they could not push this too far. That's where the groundwork was laid for statistical mechanics.

So statistical mechanics represented a retreat, a retreat in the sense that you say, "Well, I cannot say something precisely, but I'll say it statistically."

Now, the same thing happened in the case of the solution of differential equations. Today we freely accept numerical solutions. It is hard to realize that the idea of a numerical solution was not acceptable even as recently as perhaps 30–40 years ago.

Q. The rise of numerical analysis, then, constituted a retreat. Was it more of a brute force approach rather than an elegant, logical approach to the solution of differential equations?
ZADEH. Effectively, yes. People were simply not willing to say that, if you use the computer to come up with a numerical solution, you have really done something worthwhile. Somehow we tend to forget that things that are acceptable today were not acceptable 20 to 30 years ago.

Q. I can remember reading books on science of a few decades ago that always spelled science with a capital S.
ZADEH. Yes. It's that kind of veneration or worship I'm talking about. I sometimes use a word that offends people who take the more traditional view, and that word is *fetishism*—fetishism of precision and rigor in the context of classical logic.

There is also what might be referred to as "the curse of respectability in science." In trying to be respectable, scientists deny themselves the use of more flexible logical systems in which truth is a matter of degree.

Q. Is there anything that could be done to get certain people to stop worshiping precision?
ZADEH. I think it has to be a natural process. But because of the current emphasis on AI, and in particular on expert systems, there is a rapidly growing interest in inexact reasoning and processing of knowledge that is imprecise, incomplete, or not totally reliable. And it is in this connection that it will become more and more widely recognized that classical logical systems are inadequate for dealing with uncertainty and that something like fuzzy logic is needed for that purpose.

Q. Since you first developed the concept of fuzzy logic in the 1960s, Professor Zadeh, has there been much of a growth in interest? Have others picked up the banner?
ZADEH. Between then and now, somewhere between 3,000 and 4,000 papers have been written worldwide on fuzzy sets and their applications. And there are two regular journals: *Fuzzy Sets and Systems*, in English, and *Fuzzy Mathematics*, in Chinese. In addition, a quarterly entitled *Bulletin on Fuzzy Sets and their Applications* is published in France. The countries where most activity is taking place at this point are the Soviet Union, China, Japan, France, Great Britain, West Germany, East Germany, Poland, Italy, Spain and India. There has been less activity in the United States.

There is growing acceptance, but there is also considerable skepticism and in some instances hostility. At this point the largest number of researchers working on fuzzy sets is in China.

There appears to be more sympathy for other than two-valued systems in oriental countries, perhaps because their logic is not like Western, Cartesian logic. There is a greater acceptance of truth that is neither perfect truth nor perfect falsehood. This is particularly characteristic of Hindu, Chinese, and Japanese cultures.

Q. Professor Zadeh, that's the end of our questions. We on the editorial staff of *Communications* thank you warmly for giving our readers some of your views.
ZADEH. It was my pleasure.

RECOMMENDED READING

- Dubois and Prade. *Fuzzy Sets and Systems: Theory and Applications.* Academic Press, New York, 1980. This is a good, broad introduction to the subject.
- *Fuzzy Information and Decision Processes.* Elsevier North-Holland, New York, 1982. This is a recent volume of edited papers dealing with both theory and applications.
- *Fuzzy Sets and Possibility Theory.* Pergamon Press, New York, 1982. This is also a recent volume of edited papers dealing with both theory and applications.
- *Fuzzy Reasoning and Its Applications.* Academic Press, New York, 1981. A volume of edited papers.
- *Fuzzy Sets.* Plenum Press, New York, 1980.
- There are three periodicals devoted to fuzzy sets: *Fuzzy Sets and Systems*, Elsevier North-Holland, New York; *Bulletin on Fuzzy Sets and Their Application*, published in Toulouse, France; and *Fuzzy Mathematics* (in Chinese), published in Wuhan, China.