Examining fuzzy logic's widespread success

Fuzzy logic attempts to capture within a mathematical framework the vagueness of such qualitative terms as "warm" or "fast." In essence, it allows engineers to specify the degree to which something occurs—how hot it may be outside or how rapidly a motor should turn – without setting a particular threshold value for crisply separating "warm" from "hot" or "moderate" from "fast." The boundaries between categories become, well, fuzzy.

Over the last few years, fuzzy logic has permitted engineers to design control systems for air conditioners, washing machines, rice cookers, and dozens of other devices, enabling the machines to respond smoothly and appropriately to changes in operating conditions. These successes have led to increased interest in applying fuzzy logic to more complicated "expert systems" that involve, for example, making a diagnosis or suggesting a course of action.

Whether it will really work in such situations has proved controversial. "Fuzzy logic is not adequate for reasoning about uncertain evidence in expert systems," argues computer scientist Charles P. Elkan of the University of California, San Diego. Proponents contend that no fundamental barrier stands in the way of deploying large systems based on fuzzy logic.

Elkan described his viewpoint in a paper presented at the Eleventh National Conference on Artificial Intelligence, held last week in Washington, D.C.

Fuzzy logic originated in the work of computer scientist Lotfi A. Zadeh, now retired from the University of California, Berkeley. In the 1960s, Zadeh introduced the notion of fuzzy sets as a way of circumventing the rigidity of traditional set theory, which asserts that a given object or item either does or does not belong to a particular set. For example, the number 3 belongs to the set of odd numbers but not the set of even numbers.

In contrast, in Zadeh's scheme, a given item may belong only partially to a particular fuzzy set. It may also belong simultaneously to several other fuzzy sets. Thus, in a given situation, a certain temperature may be 20 percent "cool" (and 80 percent "not cool") and 70 percent "just right."

Similarly, whereas conventional logic is based on the idea that a statement or proposition is either true or false, fuzzy logic deals with the degree of truth, expressed as an assigned value between zero and one. The choice is no longer just zero or one.

To build a fuzzy system, an engineer typically begins with a set of rules, often obtained from an expert. For example, a rule for operating an air conditioner may state: If the outside temperature is high and the humidity is high, then the air conditioner setting should be very high.

"A fuzzy system is simply a set of rules that matches inputs to outputs," explains fuzzy-logic theorist Bart Kosko of the University of Southern California in Los Angeles.

By expressing the relevant variables – temperature, humidity, and motor speed – as ranges of possible values corresponding to "warm," "high," and so on, and by focusing on how these distributions of values overlap, the designer can then apply fuzzy logic in the machine's control circuitry to obtain the optimum speed at which the motor should run for a given temperature and humidity.

Fuzzy controllers have been successful because they respond more flexibly than conventional controllers to changes in operating conditions, Elkan notes. But these systems have also been quite simple.

Typically, they use fewer than 100 rules – often just 20 or so, Elkan says. Only rarely does the conclusion resulting from the application of one rule serve as the premise of another rule. Moreover, there's usually a direct relationship between the input (temperature, humidity) and the output (motor speed) of the system.

"None of these features contributing to the success of systems based on fuzzy logic is unique to fuzzy logic," Elkan contends. Other numerical, rule-based techniques – if combined with methods for obtaining the required numerical values by trial and error – would work just as well, he argues.

Elkan, who has had practical experience in developing two large-scale expert systems using more conventional techniques developed in the artificial-intelligence community, also predicts that "technical limitations" of fuzzy logic will hinder its ability to solve larger, more complex problems.

Moving to larger systems will present difficulties, Kosko concedes. "But to claim that the inferential mechanism itself is faulty and [that these faults] will start showing up in larger cases is nonsense. There are problems with fuzzy systems, but this is not one of them."

"We know in theory that you can approximate any input-output relationship to any degree of accuracy with a finite set of fuzzy rules," he adds. "Whether in practice you can find the rules, that's an open question. That can be hard [in a large system], but it's not because of something wrong in the logic."

"The sorts of things that have been successful so far have been small," notes James C. Bezdek of the University of Western Florida in Pensacola. "But there's a lot of work going on now to figure out how to scale this up." *– I. Peterson*

'Good' lipoprotein shows its bad side

In the old days, it was easy to tell the hero of a Western (he wore white) from the villain (in black). Likewise, scientists could distinguish "good" lipoproteins, or fat-transporting substances, from "bad" ones by the density of these lipid-laden particles. High-density lipoproteins (HDL) – the good guys – ferried cholesterol away from artery walls to the liver for disposal, while low-density lipoproteins (LDL) deposited lipids, which could clog arteries.

Cardiologists urge their patients to exercise and eat healthful foods in order to lower LDL and raise HDL cholesterol counts, thereby reducing the risk of heart disease.

But Westerns aren't simple anymore, and neither is the lipoprotein story. "It's more complicated than we thought," says Aldons J. Lusis, a geneticist at the University of California, Los Angeles.

Using genetically engineered mice, he and his colleagues have discovered that large amounts of HDL containing apolipoprotein A-II actually increase the chances of developing coronary artery disease. "It clearly demonstrates that the HDLs are heterogeneous and that different HDL particles have different functions," he adds.

This not-so-good HDL protein usually coexists with another protein, called apolipoprotein A-I, in particles. Other scientists had shown that excess apolipoprotein A-I reduces the risk of atherosclerosis in mice.

To assess apolipoprotein A-II, the UCLA group created transgenic mice that contain about 10 copies of the gene specifying this molecule and consequently produce excess apolipoprotein A-II. The researchers then fed these mice either a low-fat or high-fat diet.

Even on a low-fat diet, the transgenic mice wound up with two to three times more HDL and total cholesterol than nontransgenic mice. The high-fat diet roughly doubled these cholesterol concentrations in the transgenic mice, UCLA's Craig H. Warden, Lusis, and their colleagues report in the July 23 SCIENCE.

More important, high concentrations of HDL cholesterol did not protect against atherosclerosis, says Lusis. Normally, mice on low-fat diets do not accumulate the fat-filled foam cells along vessel walls that lead to clogged arteries. But examination of the aortas of the transgenic mice revealed that this buildup occurred despite high HDL, implying that apolipoprotein A-II may do real harm in addition to interfering with apolipoprotein A-I's good work, says Lusis. Also, the high-fat diet led to bigger buildups in transgenic mice. *– E. Pennisi*