Interview with Ken Forbus

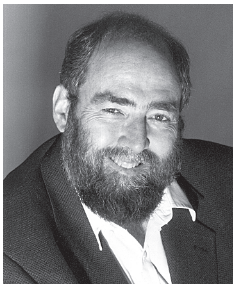

Kenneth D. Forbus is the Walter P. Murphy Professor of Computer Science and Professor of Education at Northwestern University. His research interests include qualitative reasoning, analogy and similarity, sketch understanding, spatial reasoning, cognitive simulation, reasoning system design, articulate educational software, and the use of AI in computer gaming. He received his degrees from MIT (Ph.D. in 1984). He is a Fellow of the American Association for Artificial Intelligence, the Cognitive Science Society, and the Association for Computing Machinery. He serves on the editorial boards of Cognitive Science, the AAAI Press, and on the Advisory Board of the Journal of Game Development.

KI: Ken, you are conducting research in analogy since a long time. Why is this topic of so much interest to you?

I think analogy is one of the core operations in human cognitive processing. There is evidence that the principles of Gentner's structure-mapping theory are applicable from medium-level vision to everyday reasoning to cognitive development to scientific discovery. That's quite a range of phenomena. Our computational modeling efforts do not yet span anything close to that range. But that's one of my long-range goals. So you can see why it can take a long time.

KI: What would be the appeal of an „analogical machine" for computer science?

One of the deepest scientific mysteries in AI is how common sense reasoning works. People are the most flexible and robust reasoning systems we know of. By contrast, today's software tends to be brittle and special-purpose. To be sure, today's automated reasoning systems can be effective in solving large, complex problems, but only if people formulate both the task and the knowledge required very carefully. Small changes in the formulations of problems or input knowledge can lead to significant errors or orders of magnitude differences in performance or both. If we could build software systems that worked more like people do, we might be able to overcome the brittleness and narrowness we see in today's software reasoning systems. Dedre and I suspect that people use analogy heavily in our everyday reasoning and learning. If we are right, then that could lead to significant improvements in reasoning systems.

KI: Could you give a specific example from your ongoing projects how incorporating analogy into a reasoning system can overcome the narrowness and brittleness of classical systems.

One example is our work on learning to solve physics problems. In the 1970s, people who built programs to solve physics problems hand-authored the knowledge in them. By contrast, what we are doing is giving our system worked solutions, similar to those students see in the worked examples in textbooks. By accumulating worked solutions, one can solve a large range of new problems by analogy.

People use analogy heavily in everyday reasoning and learning

We have some papers coming out about this project in AAAI and in the International Qualitative Reasoning Workshop this year. In the US, there are tests called „AP" tests, meaning Advanced Placement, which students can take to show that they already know the subject of college courses and thus get credit for them. The company which administers these tests, the Educational Testing Service (ETS), actually ran an external evaluation last summer of our system, to test its ability to transfer what it had learned. This was something of a milestone; ETS has worked with researchers many times in the past, but I think this is the first time they have directly trained and tested an AI system. Our system did quite well, averaging over 70% correct after training. It started out knowing no equations at all. It was able to use prior worked examples to figure out how to apply equations to new circumstances and derive solutions.

We are still finishing the full analysis of the errors the system made, but we can already say the following. The failures were not due to analogical processing problems: It appears to have always retrieved reasonable analogs, and mapped them correctly. The failures all appear to be caused by the hand-coded parts of the system, which could not be extended via learning. In other words, they were brittle. This is why we want to extend analogical learning throughout our architecture. We want to make as many parts of the system automatically extensible via analogical learning.

KI: This sounds very impressive. What are your long-term research goals?

My long-term research goal is to understand how minds work by building them. For a long time, we have had to focus on component processes (e.g., qualitative reasoning, spatial reasoning, and analogical matching, retrieval, generalization) and small-scale systems that capture particular expertise (e.g., engineering thermodynamics). While some effort is still going into refining these components, most of our effort is now focused on building larger-scale integrated systems. This in turn has two parts. One part is improving the communication bandwidth between people and machines, via our work on natural language and on sketch understanding. The other is exploring ideas in cognitive architecture.

Our Companions cognitive architecture is the first attempt to create a cognitive architecture organized