The collaboration game

A major shift in the way people interact with computers is coming. And it is something that we badly need. The problems we face in our societies are growing ever more complex, but our human cognitive capacities remain unchanged. Modern information technology helps, to be sure. But the current model of 'software as tool' is ultimately limited. Times change, and our software needs to change with them, ideally without the intervention of a priesthood of technical experts. I believe as artificial intelligence advances, a new model, 'software as collaborator', will become possible, with tremendous potential benefits. Collaborators adapt to each other, playing off each other's strengths, so that the whole is greater than the sum of the parts. Software collaborators could be designed to be enough like people that this mutual adaptation is possible, and that we can understand and trust their contributions. However, we should also be able to design them without certain human frailties. People tend to only look for evidence that confirms their hypotheses (called 'confirmation bias'), and have other things on their minds, such as their life outside of work. Software collaborators that did not share these frailties could become valuable complements to individuals and to teams.

We are still a long way from being able to build software collaborators, but there is important progress being made in many fronts in artificial intelligence. For example, IBM's Watson shows how a combination of AI techniques can combine synergistically to perform question-answering at a level that no one thought was possible even a few years ago. Machine reading techniques were used to assimilate vast collections of documents into internal representations that supported multiple forms of reasoning. Machine learning techniques were used to determine which strategies were likely to succeed for different types of questions. Massive hardware power was harnessed to provide real-time response, capable of performing at the level of the best humans at its task. Such a system takes a step towards the collaborator model, by adapting to the human world, instead of humans adapting to the IT world.

But this is only a first step. Collaborators engage in dialogue, with follow-up questions being interpreted with respect to the ongoing conversation. Such dialogues can include sketching and gestures, as well as text and speech (called 'multimodal dialogues'). Many researchers are working on sketch understanding, vision for understanding gestures and facial expressions (and Microsoft's Kinect will catalyse even more work in this area), and dialogue understanding. Collaborators work for long time spans, ranging from hours to years, tracking changing information, updating models to maintain situational awareness, and learning as they go.

Building robust systems that can reason and learn over a vast range of knowledge remains an exciting open challenge. Many in the artificial intelligence community are addressing this question, from a variety of perspectives. Cognitive architectures offer one intriguing approach, in trying to model cognition in the 'large', as opposed to narrow technical areas. Often this work is performed in collaboration with other cognitive scientists, since understanding how people reason, learn, and interact provides valuable clues for creating intelligent systems.

Watson's enormous computing requirements may seem to limit the potential for future systems, which will require even more computation than it used. However, yesterday's supercomputer is tomorrow's smartphone, and within a few years of Deep Blue's victory at chess in 1997, there were programs that performed at similar levels without special hardware. So assuming artificial intelligence – and computer science and engineering more broadly – remains on-course, we should be able to create software collaborators.

Kenneth D Forbus
Chair
Cognitive Science Society