Editor: James Hendler
University of Maryland
hendler@cs.umd.edu

Computers Play Chess; Humans Play Go

Intelligent Readers,

Over the past few weeks, I've been trying to get my thoughts together about AI. I've been overdue on a couple of write-ups related to AI's 50th anniversary and have been experiencing terrible writer's block. Many thoughts, some of which have appeared in previous Letters from the Editor, have been crowding in my head, but I couldn't get them all to cohere. With the due date for this editorial looming, I was beginning to panic. What theme would pull it all together?

Then last night, I had a really weird dream. Really. It was a compendium of those anxiety dreams people get: I was late on getting a grant in, and suddenly everyone gave me stacks and stacks of unrelated junk that I had to write up by Friday. The elevator wouldn't stop at my floor. I was giving a talk, showed up early to have time to prepare my slides, and discovered I'd been given the wrong time and was talking "right now!" When I went to get my paper so that I could wing it, the paper was only one page of unreadable handwritten notes, with a sticky note saying "where are the other 16 pages you owe us?" So there I was, unprepared, not knowing what meeting this was, in front of the audience (in my bathrobe, of course). All I could think to utter was "Computers play chess; humans play Go." On that note, I awoke.

I lay in bed in the middle of the night pondering this cryptic phrase. It came to me that, with the usual imprecision of nightmares, this was the key to my conundrum—the phrase that could collect my scattered thoughts. See if you agree.

Computers play chess

In my January/February letter ("Fly, But Not Like an Eagle"), I discussed the idea that humans had learned to outfly birds in some ways but not in others, and that this was similar to the state of AI. In many ways, modern AI has enabled the computer to outperform humans—thus, in essence, achieving superhuman intelligence. Chess is the canonical example. The best chess-playing programs outperform the best humans; nowadays, a human had better play at the master's level or better to have a chance against even a mediocre chess program (despite predictions, not that long ago, that a computer would never play master's-level chess). The combination of brute-force computational power, huge memory, and good heuristic board evaluators make chess-playing programs formidable foes.

Chess isn't the only example; let's look at other domains in which AI has proved supreme over us puny humans. We might still have an edge when it comes to understanding language. However, we don't stand a chance next to an inverted index and a heuristic ranking function when it comes to recommending which of the billions of Web pages are most useful for a task at hand. Humans are often considered the masters of logic, but we seem to be better at illogic. If you want to take a set of first-order-logic axioms and reach a proof that requires a long and complex trail, you should use Snark, Vampyr, or one of the other modern theorem provers. Doing this saves you as much time in the logic domain as your calculator will in multiplying long series of multidigit numbers. Or consider data mining. At this point, a computer recognizes the regularities in a set of training data a lot better than a human does. In fact, almost any AI subfield has limited sets of problems where the computer sets us slow-thinking humans back on our heels.

In short, computers play chess.

Humans play Go

Go is very different from chess. Its combinatorics dwarf those of chess, and despite having relatively simple rules, it's said to be nontrivial even to write a program that can determine a finished game's winner. In the past few years, computer Go tournaments have been

We'd Like to Hear from You

Letters to the Editor:
Send letters, including a reference to the article in question, to intelligent@computer.org. Letters will be edited for clarity and length.

Articles:
If you're interested in submitting an article for publication, see our author guidelines at www.computer.org/intelligent/author.htm.

1541-1672/06/$20.00 © 2006 IEEE
Published by the IEEE Computer Society

awarding large payoffs, and prizes of over a million dollars have been offered for a computer program that can beat a professional player in a seven-game series. These prizes go unclaimed. In fact, the best computer Go players right now play at a rank of about 6 or 7 kyu, which is at the advanced-beginner level. While there are matches for players ranked in this range, humans are generally considered to be unworthy of serious competition until they reach significantly higher ranks. In short, computers are nowhere close to winning at Go.

While I'm not a Go player, I was interested in computer Go when I worked at DARPA in 2000, and I discussed the game with a number of good players. My understanding is that much of what makes Go so interesting, and so challenging, is that playing a good game requires a lot of pattern recognition (with abstraction) and that there are a lot of "nonlocal" effects—that is, a stone that's relatively far from the action's center can be crucial to the outcome. In fact, even talking about the center of action is wrong; strong players appear to be playing in several centers of action at the same time.

This latter hallmark of Go reveals, I think, where humans triumph over computers at this point in history. While computers can do deep searches and use complex rules in narrow domains, we win when it comes to either "synthesis" of results over multiple domains or recognition of a key aspect of a pattern that's spread over a wide set of entities. Our ability to handle language remains awe inspiring. Our ability to jump to the right conclusion against partially formed hypotheses is amazing. And, while computers might beat us at recognizing the regularities in data, the oddity in the data is what catches the human epidemiologist's eye.

In short, humans play Go.

The future of AI

Perhaps you can see where I'm heading. It seems to me that our old talk of human-level intelligence as AI's manifest destiny isn't quite right. We already see hard problems where the computer outperforms us, and I expect to see many more in the future. But we also see "easy" problems where we outperform the computer using this strange mass of wetware between our ears. Unfortunately, it doesn't seem to me that understanding the wetware will more likely explain how we think than understanding the hardware explains how the computer plays chess. Such understanding might yield important insights, but we still need to understand the "mind," not just the brain.

So, the future of AI must involve exploring and understanding the parts of human intelligence we haven't been looking at that much—the stuff at the heart of human thought. To do this, we need to stop looking for new ways to solve well-defined problems and start looking for ways to combine the things we know how to do, and then see if this helps us explore problems with more diversity and scope.

I'm encouraged to see a few developments that are taking us in the right direction. Ron Brachman, in his AAAI 2005 presidential address and in his role as a DARPA office director (see "Systems That Know What They're Doing," Nov./Dec. 2002, pp. 67–71), advocated large AI projects that would force people from different parts of AI to work together to achieve, in consortium, what no approach could achieve alone. I recently became a member of one of these DARPA-funded groups. I can testify that putting together a team of experts from different AI areas (planning, learning, and knowledge representation, among others) requires breaking new ground, and we hope it will produce unique results. However, these large-scale projects are expensive, the funding agents' needs dictate the scope, and the effort is high. So, although this is an important step in the right direction, it leaves a lot to be desired.

Another encouraging sign, although I admit to being biased, is the Semantic Web research that's going on. I don't mean all, or even most, of this research, which is often old AI in new clothes. I mean the research that's focusing on the needs of linking different people's models and representations—what I sometimes call the "Web side of the Semantic Web." This research is helping to create an infrastructure where people can try their techniques on other peoples' knowledge, where inconsistent and opposing views must necessarily be reconciled, and where the illogic that people are so good at comes into contact with the logic where machines rule.

Other encouraging signs include

- efforts to create large repositories of shared, processed data,
- efforts to create infrastructures allowing different agent architectures to play together, and
- a move from challenges that encourage overfitting to a single goal (forgive me, but think "RoboCup") to challenges that encourage more open performance and scoring—sort of like the free-skate event in the Olympics.

Editorial Board Transition

Austin Tate is stepping down from his role as Associate Editor in Chief of *IEEE Intelligent Systems*. We thank him for his years of dedication to the magazine in this role. He will continue to serve on our advisory board.

Unfortunately, despite these encouraging signs, I must admit that I have days where my waking thoughts are almost as pessimistic as my nightmare. AI continues to splinter into small communities that focus on performing incrementally better on narrower and narrower problems. Math envy continues to dominate many of our key journals, making it hard to publish papers on exactly the sorts of encouraging work I've just discussed. And, we seem to be finding more and more ways to blindly explore the elephant of AI, without getting any closer to figuring out what the elephant is.

But I also have days when my own naïve optimism about AI shines through and my nightmares fade before the light of exciting new results. As I've said in one way or another in these letters, I believe that for AI the best is yet to come. After all, no matter how we approach it, the AI problem remains one of the most exciting and challenging that the human race can ever pursue—understanding this human thing we call "thought." Exploring the things that humans can do, but that we can't yet even imagine how to get computers to perform, is still the hallmark and challenge of AI. Pursuing it will keep the next 50 years of our field just as exciting as the past half century has been. In short, it's time to learn how to play Go. ■