


Rethinking artificial intelligence: Researchers hope to produce 'co-processors' for the human mind

December 7 2009, by David L. Chandler



The field of artificial-intelligence research (AI), founded more than 50 years ago, seems to many researchers to have spent much of that time wandering in the wilderness, swapping hugely ambitious goals for a relatively modest set of actual accomplishments. Now, some of the pioneers of the field, joined by later generations of thinkers, are gearing up for a massive 'do-over' of the whole idea.

This time, they are determined to get it right — and, with the advantages of hindsight, experience, the rapid growth of new technologies and

insights from the new field of computational neuroscience, they think they have a good shot at it.

The new project, launched with an initial \$5 million grant and a five-year timetable, is called the [Mind Machine Project](#), or MMP, a loosely bound collaboration of about two dozen professors, researchers, students and postdocs. According to Neil Gershenfeld, one of the leaders of MMP and director of MIT's Center for Bits and Atoms, one of the project's goals is to create intelligent machines — “whatever that means.”

The project is “revisiting fundamental assumptions” in all of the areas encompassed by the field of AI, including the nature of the mind and of memory, and how intelligence can be manifested in physical form, says Gershenfeld, professor of media arts and sciences. “Essentially, we want to rewind to 30 years ago and revisit some ideas that had gotten frozen,” he says, adding that the new group hopes to correct “fundamental mistakes” made in AI research over the years.

The birth of AI as a concept and a field of study is generally dated to a conference in the summer of 1956, where the idea took off with projections of swift success. One of that meeting's participants, Herbert Simon, predicted in the 1960s, “Machines will be capable, within 20 years, of doing any work a man can do.” Yet two decades beyond that horizon, that goal now seems to many to be as elusive as ever.

It is widely accepted that AI has failed to realize many of those lofty early promises. “Considering the outrageous optimism of much of the early hype for AI, it is no wonder that it couldn't deliver. This is an occupational hazard of many new fields,” says Daniel Dennett, a professor of philosophy at Tufts University and co-director of the Center for Cognitive Science there. Still, he says, it hasn't all been for nothing: “The reality is not dazzling, but still impressive, and many applications

of AI that were deemed next-to-impossible in the '80s are routine today,” including the automated systems that answer many phone inquiries using voice recognition.

Fixing what's broken

Gershenfeld says he and his fellow MMP members “want to go back and fix what's broken in the foundations of information technology.” He says that there are three specific areas — having to do with the mind, memory, and the body — where AI research has become stuck, and each of these will be addressed in specific ways by the new project

The first of these areas, he says, is the nature of the mind: “how do you model thought?” In AI research to date, he says, “what's been missing is an ecology of models, a system that can solve problems in many ways,” as the mind does.

Part of this difficulty comes from the very nature of the human mind, evolved over billions of years as a complex mix of different functions and systems. “The pieces are very disparate; they're not necessarily built in a compatible way,” Gershenfeld says. “There's a similar pattern in AI research. There are lots of pieces that work well to solve some particular problem, and people have tried to fit everything into one of these.” Instead, he says, what's needed are ways to “make systems made up of lots of pieces” that work together like the different elements of the mind. “Instead of searching for silver bullets, we're looking at a range of models, trying to integrate them and aggregate them,” he says.

The second area of focus is memory. Much work in AI has tried to impose an artificial consistency of systems and rules on the messy, complex nature of human thought and memory. “It's now possible to accumulate the whole life experience of a person, and then reason using these data sets which are full of ambiguities and inconsistencies. That's

how we function — we don't reason with precise truths," he says. Computers need to learn "ways to reason that work with, rather than avoid, ambiguity and inconsistency."

And the third focus of the new research has to do with what they describe as "body": "Computer science and physical science diverged decades ago," Gershenfeld says. Computers are programmed by writing a sequence of lines of code, but "the mind doesn't work that way. In the mind, everything happens everywhere all the time." A new approach to programming, called RALA (for reconfigurable asynchronous logic automata) attempts to "re-implement all of computer science on a base that looks like physics," he says, representing computations "in a way that has physical units of time and space, so the description of the system aligns with the system it represents." This could lead to making computers that "run with the fine-grained parallelism the brain uses," he says.

MMP group members span five generations of [artificial-intelligence](#) research, Gershenfeld says. Representing the first generation is Marvin Minsky, professor of media arts and sciences and computer science and engineering emeritus, who has been a leader in the field since its inception. Ford Professor of Engineering Patrick Winston of the Computer Science and Artificial Intelligence Laboratory is one of the second-generation researchers, and Gershenfeld himself represents the third generation. Ed Boyden, a Media Lab assistant professor and leader of the Synthetic Neurobiology Group, was a student of Gershenfeld and thus represents the fourth generation. And the fifth generation includes David Dalrymple, one of the youngest students ever at MIT, where he started graduate school at the age of 14, and Peter Schmidt-Nielson, a home-schooled prodigy who, though he never took a computer science class, at 15 is taking a leading role in developing design tools for the new software.

The MMP project is led by Newton Howard, who came to MIT to head this project from a background in government and industry computer research and cognitive science. The project is being funded by the Make a Mind Company, whose chairman is Richard Wirt, an Intel Senior Fellow.

“To our knowledge, this is the first collaboration of its kind,” Boyden says. Referring to the new group’s initial planning meetings over the summer, he says “what’s unique about everybody in that room is that they really think big; they’re not afraid to tackle the big problems, the big questions.”

The big picture

Harvard (and former MIT) cognitive psychologist Steven Pinker says that it’s that kind of big picture thinking that has been sorely lacking in AI research in recent years. Since the 1980s, he says “there was far more focus on getting software products to market, regardless of whether they instantiated interesting principles of intelligent systems that could also illuminate the human mind. This was a real shame, in my mind, because cognitive psychologists (my people) are largely atheoretical lab nerds, linguists are narrowly focused on their own theoretical paradigms, and philosophers of mind are largely uninterested in mechanism.

“The fading of theoretical AI has led to a paucity of theory in the sciences of mind,” Pinker says. “I hope that this new movement brings it back.”

Boyden agrees that the time is ripe for revisiting these big questions, because there have been so many advances in the various fields that contribute to artificial intelligence. “Certainly the ability to image the neurological system and to perturb the neurological system has made great advances in the last few years. And computers have advanced so

much — there are supercomputers for a few thousand dollars now that can do a trillion operations per second.”

Minsky, one of the pioneering researchers from AI’s early days, sees real hope for important contributions this time around. Decades ago, the computer visionary Alan Turing famously proposed a simple test — now known as the Turing Test — to determine whether a machine could be said to be truly intelligent: If a person communicating via computer terminal could carry on a conversation with a machine but couldn’t tell whether or not it was a person, then the machine could be deemed intelligent. But annual “Turing test” competitions have still not produced a machine that can convincingly pass for human.

Now, Minsky proposes a different test that would determine when machines have reached a level of sophistication that could begin to be truly useful: whether the machine can read a simple children’s book, understand what the story is about, and explain it in its own words or ask reasonable questions about it.

It’s not clear whether that’s an achievable goal on this kind of timescale, but Gershenfeld says, “We need good challenging projects that force us to bring our program together.”

One of the projects being developed by the group is a form of assistive technology they call a brain co-processor. This system, also referred to as a cognitive assistive system, would initially be aimed at people suffering from cognitive disorders such as Alzheimer’s disease. The concept is that it would monitor people’s activities and brain functions, determine when they needed help, and provide exactly the right bit of helpful information — for example, the name of a person who just entered the room, and information about when the patient last saw that person — at just the right time.

The same kind of system, members of the group suggest, could also find applications for people without any disability, as a form of brain augmentation — a way to enhance their own abilities, for example by making everything from personal databases of information to all the resources of the internet instantly available just when it's needed. The idea is to make the device as non-invasive and unobtrusive as possible — perhaps something people would simply slip on like a pair of headphones.

Boyden suggests that the project's initial five-year timeframe seems about right. "It's long enough that people can take risks and try really adventurous ideas," he says, "but not so long that we won't get anywhere." It's a short enough span to produce "a useful kind of pressure," he says. Among the concepts the group may explore are concepts for "intelligent," adaptive books and games — or, as Gershenfeld suggests, "books that think."

In the longer run, Minsky still sees hope for far grander goals. For example, he points to the fact that his iPhone can now download thousands of different applications, instantly allowing it to perform new functions. Why not do the same with the brain? "I would like to be able to download the ability to juggle," he says. "There's nothing more boring than learning to juggle."

Provided by Massachusetts Institute of Technology ([news](#) : [web](#))

Citation: Rethinking artificial intelligence: Researchers hope to produce 'co-processors' for the human mind (2009, December 7) retrieved 19 April 2024 from <https://phys.org/news/2009-12-rethinking-artificial-intelligence-co-processors-human.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private

study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.