

Rethinking artificial intelligence

Broad-based MIT project aims to reinvent AI for a new era. By going back and fixing mistakes, researchers hope to produce 'co-processors' for the human mind.

David L. Chandler, MIT News Office

December 7, 2009

[email](#)

[comment](#)

[print](#)

[share](#)

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



today's news

'Micro-ants': Tiny conveyor belts for the 21st century



Graphic: Christine Daniloff

A new method of moving tiny particles using magnetic polymer beads and magnetic fields could find uses in microchips and in medicine

New evidence links sirtuins and life extension

December 15, 2009

Their infinite wisdom

December 14, 2009

related

[Mind Machine Project](#)

[Neil Gershenfeld](#)

[Marvin Minsky](#)

[Patrick Winston](#)

[Ed Boyden](#)

[Gerald Jay Sussman](#)

[Henry Lieberman](#)

[Sebastian Seung](#)

[MIT Computer Science and Artificial Intelligence Laboratory \(CSAIL\)](#)

[Center for Bits and Atoms](#)

[The Open Mind Initiative](#)

[RALA](#)

[ARCHIVE: "Papert misses 'big ideas' from early years of artificial intelligence"](#)

[ARCHIVE: "Hope on the horizon: Digital fabrication"](#)

[ARCHIVE: "Burgeoning field of neurotech spawns brainy businesses"](#)

tags

[Artificial intelligence](#)

[Brain and cognitive sciences](#)

[Center for Bits and Atoms](#)

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-Nielsen, 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."

Comments

[Log in to write comments](#)

robin7013 - Rethinking AI

2009-12-07 05:04:49

I am happy to read that there is a new ray of hope among AI researchers to make AI practical. Hope they will win the race this time.

monimonawar - AI and Linguistics

2009-12-07 09:23:54

I'm a Master's student in Linguistics and have devoted some of my time studying possible interfaces between Linguistics and AI. This is something that could produce great results if looked closely! We definitely have to look at the big picture, keeping the brush strokes in mind as well :)

Keith

2009-12-07 09:24:16

So, the plan is to roll back the clock 30 years and try out abandoned ideas while avoiding past mistakes, so that they can wipe the slate clean and revolutionize the field of AI? And create intelligent machines that will "understand" children's books in 5 years?

I'm glad that they're avoiding the hype and grandiose promises that lead to people talking about the "death" of AI when it didn't deliver.

Seriously, though, it looks like there are a lot of incredibly smart people in the group, and I'm sure some amazing things will come out of it regardless of contradictory and vaguely defined goals. I'm waiting with bells on.

romantercero - brain pda...

2009-12-07 10:42:58

I cant wait for my cognitive augmentation headphones!

But seriously. There is something very troubling in this article. I can't trust a man that presents juggling as as boring activity to learn.

Good day sir!

smokfrog - Human

2009-12-07 10:43:07

Definitely the harbinger of a new era, pending success. But humanizing machines with a specified ultimate goal of dehumanizing humans? Not the lofty ideals I would hope for in a project of this caliber.

natch - "Nothing more boring"?

2009-12-07 13:24:43

"There's nothing more boring than learning to juggle."

Wow. When I was learning to juggle, the interesting part wasn't the action of repeatedly throwing, catching, and dropping balls. It was the internal dialogue I was having with myself about the nature of learning, what is possible, the difference between what is possible and what one thinks is possible to learn, how people challenge themselves... what kind of close-minded simpleton would think that learning to juggle is boring? Maybe the kind who has held back AI research for the last 50 years.

neuroscience2011 - Revisiting A.I.

2009-12-08 05:00:11

Smokfrog has a point, and juggling is boring? Usually that is an indication of someone that tried to learn to perform an activity and failed at it. No reward equals boring then.

Dehumanizing individuals could be bad if gone too far, but the project will most likely have a few shortcomings to prevent overall dehumanization of the human population. It seems to depend on the progress of the group and how far they decide to take things in each direction. I am not sure it is accurate to say they are replaying the past with A.I. To do so would mean they would have to repeat the same mistakes over again, inviting disaster. Therefore, it should be a worthwhile project in the end with some interesting results.

YODspica - Creative Computer Intelligence

2009-12-08 05:00:28

This is an interesting development, however it is difficult for humans to create machine intelligence derived from concepts based on humans.

This research could lead to perhaps achieve an outcome that concentrates on creating intelligence that would appear intelligent to humans but not necessarily be true machine intelligent.

syntience - AI took a wrong turn in 1650

2009-12-08 05:01:03

Rolling back 30 years is not enough. The main problem is the Reductionist stance. Intelligence is a holistic (in the sense of Epistemology as opposed to the "crystals and aromatherapy" sense) and emergent effect. For more info, see (2+ years old site) <http://artificial-intuition.com> and the (more recent) videos at <http://videos.syntience.com>.

And "Modeling thought" is wrong; the problem goes deeper than that. We must use Model Free Methods the way the life sciences do it. In fact, AI should have been a life science from day one, and not part of computer science.

- Monica Anderson

realdreamer - the art in artificial

2009-12-09 05:39:32

I caught the end of a program the other day on the radio about AI, and the gist of this had to do with people falling in love with robots, and a host of "ethical problems" that might be part of this. It seemed "way out", but not so way out. Maybe "string theory" is part of this, and I wonder some time, given the humor and scary side of this, whether somebody is "pulling the strings". Is there an intelligence guiding this entire show, as in, this symphony has a conductor?

Then again, if this were true, would we be the robots? We don't yet know really what consciousness is and people are talking about thought as energy. But what drives consciousness? Why do we attend to only a portion of what we see and hear? It is obvious that given an event, observed by many, we each perceive it differently, and we take in differently. Do we each have individual filters that have to do with our own stories? Probably. But I think pondering this goes deeper. Is consciousness itself, somehow, directed?

Another question is, how far do we really want to go, if we have choices, and we do act this way? Do we want artificial everything? We do have now, artificial limbs, and parts, and the program was saying we are already approaching that sci fi notion of cyborgs, that are part human, part machine.

I think we need to remember that being human, has to have a component of humanity to it. And I often wonder, why it is, we need so much and that we humans are always reaching for more, often without realizing how much we have, right here, right now.

gssdrone1

2009-12-09 05:41:04

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."

You need to have an object of study...what is intelligence? This question is the beginning.

tessera - The blind leading the blind

2009-12-10 04:45:10

Another 50 years going down wrong paths won't lead anywhere either. With their preconceived

biases these guys have already decided how and what to pursue, trying to force Nature down paths of their choosing, so this is a sure guarantee that they've already missed the boat. Only a non-deterministic non-hierarchical process can be the basis of thought, and you can't do that on a Turing machine. Good luck, you'll need it...

jonathanbruder - The blind feeding the trolls

2009-12-13 15:45:28

"With their preconceived biases these guys have already decided how and what to pursue..."

"Only a non-deterministic non-hierarchical process can be the basis of thought..."

LULZ. Sounds like a preconceived bias to me.

I would be astounded if you read the release above. It would be difficult to find a proclamation that more closely matches the sort of thinking this project was established to overcome.

You can't shoehorn neurocognitive processes into a single model. Some might be non-deterministic and non-hierarchical; others necessarily will not.

By using a quilt of independent models, these researchers aim to cover a wider swath of cognitive functionality than a single model could provide for. I don't know if they'll be successful, but I'd guess they'll make some ground.

Of course, you could always scroll up twenty inches and read what I'm repeating from.

santi - funny how inconsistent you are :)

2009-12-13 15:48:22

Speaking of people with "preconcieved biases", apparently you already think you know what can and cannot be the basis of thought :)

Additionally, once you place a turing machine in a non-deterministic environment (like the real world), you can easily get non-deterministic behaviors out of one. But anyway, since apparently you have already made up your mind. I won't try to convince you any further.

jdsetls - AI and linguistics again

2009-12-10 04:46:46

This in my opinion is quite the point. I have a diploma in linguistics and finishing a PHD in economics : Economy of thought in preventing and eradicating software bugs.

Our mind has a background of different preprocessing and blind processing like the pre-conscious working at full speed and often for nothing. Guies like Kepler or Euler had it working for something, which made all the difference.

However Kepler had not the user's manual of this gift and Euler had. So Kepler's teachings were boring and he had few pupils. He was always disgrressing.

Euler played with his children and grand children, then went to his cabinet. All was ready in his head. He published over eight thousands mathematics articles.

So I wonder if AI should not work on connotations (subjective, cultural, implicit and/or emotional coloration in addition to the explicit or denotative meaning of any specific word or phrase in a language, i.e. emotional association with a word.), multiple meanings and probability, exactly as the article says, in an uncertain world, accepting contradictions, using methods to deal with uncertainty and managing the vast field of the implicit.

Jacques De Schryver jdsetls@aim.com

jgord - new start

2009-12-10 04:48:00

I applaud these people. This is the breath of fresh air that all science can learn from. I have no quarrel with their efforts in the past, but often a blank sheet of paper opens up new possibilities. It is most encouraging that these scientists have the maturity and creativity to re-evaluate without being forced to do so by outside forces, i.e., being "proved" wrong, or being forced aside by the popular trend of the moment. In so doing we all benefit from their

years of accumulated experience while bring new ways of thinking into the mix.

Might I suggest that they also consider trying to fit our new social trends of openness and collaboration into their process. One model of which I was very skeptical 15 years ago has proven itself very powerful and productive, the software open source movement. The productivity of the past close to the vest practices of both academia and industry is being rapidly bypassed by these fresh approaches. The \$5 million funding is a lot of money to an individual, but not that much for a significant project, these new approaches may provide huge leverage.

It looks like they wish to bring in researchers from many disciplines to look at this from fresh perspectives, both process and discovery. One person they might add to the mix is a "crackpot". By this I mean some person that does not necessarily have a lot of extra letters after their name, but one of the 'garagistes' with a head full of ideas. Most of them may be dead ends, but it keeps minds fresh, and if you can apply an open source type of model you will have many helpers to examine all ideas at an early stage, rather using the old adversarial model which waits to publish and defend.

Bintu

2009-12-11 05:06:28

It might generate lot of interest in open source community to work on such complex project, if it also announces prizes for different module of this bigger goal. Something like netflix prize.

rslippert - Conceptual (Shared) Knowledge Representation

2009-12-11
10:38:35

>...brain augmentation..cut.. something people would simply slip on like a pair of headphones...cut..

OK...yes...Proceed by placing headphones on a blind man so that this person hears the "Conceptual Language" of vision. This at once accomplishes many things; 1) the blind can begin to see, 2) we will begin to understand how to represent knowledge at it's lowest level conceptual representation. To fully understand the science of "Conceptual Knowledge Representation" the blind person will tell us where we are getting things wrong.

Now having said that, the idea of the need for an "ecology of models" is that there is some form of (two way)communication between these models. This communication is likely to be where "how to interpret" involves semantic network like re-alignment of concepts via some common shared conceptual language. So 3) We also begin to understand how to build distributed knowledge "semantic networks".

Finally, the information on system needs from another system is selective. The blind person will want to, for example, focus or turn the camera. From a much larger context, the eyes see, but must know how to see, and what to filter out. There are two aspects of this; 4) external control, and 5) The Blocking of internal knowledge not desired by the controlling external function. In other words, the controlling function must tell the eyes to report only moving objects (that kind of thing).

wlburke - Network Analyst AT&T

2009-12-13 15:46:20

Departure

Be it sight, sound, the smell, the touch.
There's something,
Inside that we need so much,
The sight of a touch, or the scent of a sound,
Or the strength of an arquebus deep in the ground.
The wonder of flowers, to be covered, and then to burst up,
Thru tarmac, to the sun again,
Or to fly to the sun without burning a wing,
To lie in the meadow and hear the grass sing,
To have all these things in our memories hoard,
And to use them,
To help us,

To find...

<http://www.youtube.com/watch?v=wGEye0b5JXw>

rjm612 - Rethinking AI

2009-12-13 15:47:39

At the risk of seeming too radical I'm proposing that the problem is not the results or lack of results over the last 50 years, but the concept itself.

Measuring intelligence in human beings and animals is controversial. This must mean that our understanding of what intelligence is or indeed if it exists and can be identified is at best very poor.

How can we create artificial intelligence if we really can't identify real intelligence?

Therefore the Turing Test and other tests of artificial intelligence would be better deemed computer science problems, like the problems of mathematics (eg the four colour problem).

Personal view, but if this group could do what Turing and others did 50 years ago and propose a new set of problems that extend our concept of computer science and computation that would be much more useful.

quantum_flux - Might as well call it "artificial stupidity" inste

2009-12-13
15:47:48

Learning to live with vague things is not very intelligent though, but the precision of understanding physics and mathematics is.

ferrouswheel

2009-12-13 15:48:12

While I'm glad some people are getting the funding to work on these AI challenges, it'd be nice if they didn't make such claims as being the first to approach the challenge from this direction.

I work on the open-source artificial general intelligence platform OpenCog and we've been around for the last two years working on exactly the same ideas.

<http://opencog.org/>

Unfortunately we don't have the financial support that MIT has, so progress is slow going. On the other hand, we're open source and community based so our research will remain available and not disappear into the archives of academia as so many AI projects are prone to do.

prabhakardbagalkot - How the Brain works....

2009-12-13 15:48:46

Sirs

This topic has been one of the area which has been with in my mind for long time. It is good news that MIT is doing work in this area.

... I am an old student of IIT Kharagpur ,India. Our institute was formed with ideas derived from MIT and IIT Illinois, Cal Tech ,and many others during 1950's. You are our elder Gurus.

... I have studied Biology and Maths together in the Pre professional programmes. Later studied Mechanical engg and Machine tools, CNC, FEM etc.

...I am always interested to know how the molicular structures ,floating or partly flexible, in the brian, with neurons, electric pulses, and a kind of biological memory storage device,...able to store ,analyse ,restructure,conclude,find alternative thoughts, thought process, etct ,etc, and logics with in work.

I am sure this will help all in the world over both in Biology,Medecine and Computers, Help to repair biological circuits within the brain, medulla oblangata,and the nerve cells, etc etc

Wish you all great sucess and Happy Christmass.

I am also a self learner/parent of univ student ,in India. Study the MIT OCW web regularly.

Regards

Prabhakar.D.Bagalkot .DGM Projects, Tata Steel India...MTech Mechanical engineering ,IIT Kharagpur India....

btilmann

2009-12-13 15:49:10

I am puzzled the "Big Picture" game plan fails to mention fundamental advances in computer technology other than.... "there are supercomputers for a few thousand dollars now that can do a trillion operations per second."

The very nature of computers are about to change...and it has little to do with speed per se. The coming Memristor technology should be part of your game plan.

"Memristor theory was formulated and named by Leon Chua in a 1971 paper. Chua extrapolated the conceptual symmetry between the resistor, inductor, and capacitor, and inferred that the memristor is a similarly fundamental device."

"On April 30, 2008 a team at HP Labs announced the development of a switching memristor. Based on a thin film of titanium dioxide, it has a regime of operation with an approximately linear charge-resistance relationship. These devices are being developed for application in nanoelectronic memories, computer logic, and neuromorphic computer architectures."

The implications this technology holds for the future of AI is profound and should be a key component of the MIT "Big Picture" game plan.

btilmann

fairlane32

2009-12-13 15:49:25

Computers need to learn "ways to reason that work with, rather than avoid, ambiguity and inconsistency."

Agreed. Machines need to be able to learn to spot those ambiguities, and inconsistencies, and react accordingly, as in having them be minimized.

djayjp - **Delusive Grandiosity and arrogance once again**

2009-12-13
15:49:56

So they claim to be abandoning the unrealistic and grandiose aims of their forebears, yet the part about the brain co-processor is OBSCENE. Was that a joke? This is the field that has been struggling for half a century and they propose this?!

Anyway, the key to reaching true general/strong intelligence is by working from the bottom up. Forget about any notion of creating a "mind" and start with the basics-- applied intelligence. Can your model produce the same intelligent behaviour of ants, bees, mosquitoes, cats, dogs? Once this is done, THEN things like natural language (in the most basic sense) and imagination and wonder/existential questioning will come in time.

Another route that might be interesting in relation to being able to produce a simulation of something that can fight and flight a prey and predator is to use the chess-like, many solutions computed, best solution chosen by statistics, of e.g., Deep Blue regarding pathfinding. See: "March of the Froblins" tech demo as well for AI on the GPU. The \$400 radeon 5870 can do almost 3 teraflops single precision (0.5 double).

