theguardian | TheObserver

AI robot: how machine intelligence is evolving

No computer can yet pass the 'Turing test' and be taken as human. But the hunt for artificial intelligence is moving in a different, exciting direction that involves creativity, language – and even jazz

Marcus du Sautoy
The Observer, Sunday 1 April 2012
Article history

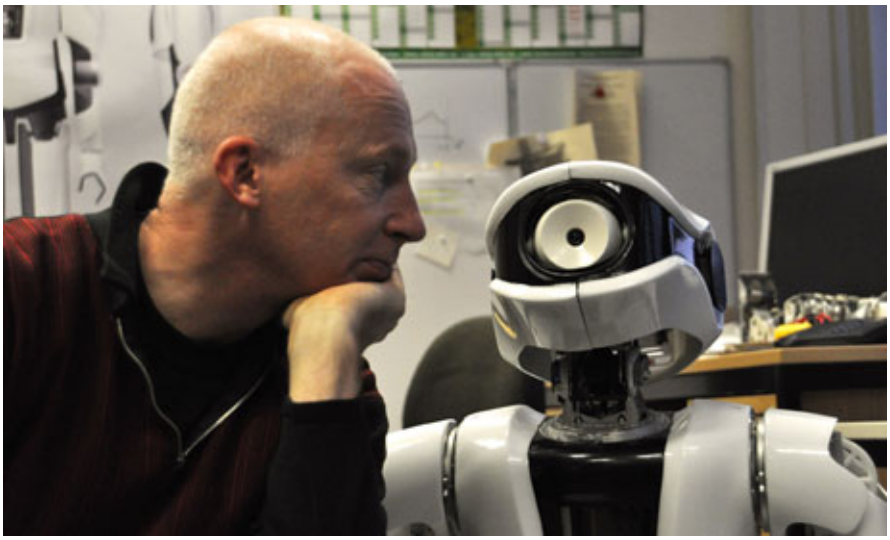

Marcus du Sautoy with one of Luc Steels's language-making robots. Photograph: Jodie Adams/BBC

'I propose to consider the question "Can machines think?"' Not my question but the opening of Alan Turing's seminal 1950 paper which is generally regarded as the catalyst for the modern quest to create artificial intelligence. His question was inspired by a book he had been given at the age of 10: *Natural Wonders Every Child Should Know* by Edwin Tenney Brewster. The book was packed with nuggets that fired the young Turing's imagination including the following provocative statement:

"Of course the body is a machine. It is vastly complex, many times more complicated than any machine ever made with hands; but still after all a machine. It has been likened to a steam machine. But that was before we knew as much about the way it works as we know now. It really is a gas engine; like the engine of an automobile, a motor boat or a flying machine."

If the body were a machine, Turing wondered: is it possible to artificially create such a contraption that could think like he did? This year is Turing's centenary so would he be impressed or disappointed at the state of artificial intelligence? Do the extraordinary

machines we've built since Turing's paper get close to human intelligence? Can we bypass millions of years of evolution to create something to rival the power of the 1.5kg of grey matter contained between our ears? How do we actually quantify human intelligence to be able to say that we have succeeded in Turing's dream? Or is the search to recreate "us" a red herring? Should we instead be looking to create a new sort of machine intelligence different from our own?

Last year saw one of the major landmarks on the way to creating artificial intelligence. Scientists at IBM programmed a computer called Watson to compete against the best the human race has to offer in one of America's most successful game shows: *Jeopardy!* It might at first seem a trivial target to create a machine to compete in a general knowledge quiz. But answering questions such as: "William Wilkinson's *An account of the principalities of Wallachia and Moldavia* inspired this author's most famous novel" requires a very sophisticated piece of programming that can return the answer quickly enough to beat your rival to the buzzer. This was in fact the final question in the face-off with the two all-time champions of the game show. With the answer "Who is Bram Stoker?" Watson claimed the *Jeopardy!* crown.

Watson is not IBM's first winner. In 1997 IBM's super computer Deep Blue defeated reigning world chess champion Garry Kasparov. But competing at *Jeopardy!* is a very different test for a computer.

Playing chess requires a deep logical analysis of the possible moves that can be made next in the game. Winning at *Jeopardy!* is about understanding a question written in natural language and accessing quickly a huge database to select the most likely answer in as fast a time as possible. The two sorts of intelligence almost seem perpendicular to each other. The intelligence involved in playing chess feels like a vertical sort of intelligence, penetrating deeply into the logical consequences of the game, while *Jeopardy!* requires a horizontal thought process, thinking shallowly but expansively over a large data base.

The program at the heart of Watson's operating system is particularly sophisticated because it learns from its mistakes. The algorithms that select the most likely answers are tweaked by Watson every time it gets an answer wrong so that next time it gets a similar question it has a better chance of getting it right. This idea of machine learning is a powerful new ingredient in artificial intelligence and is creating machines that are quickly doing things that the programmers hadn't planned for.

Despite Watson's win, it did make some very telling mistakes. In the category 'US cities' contestants were asked: "Its largest airport is named for a world war two hero; its second largest for a world war two battle." The humans responded correctly with "Where is Chicago?" Watson went for Toronto, a city that isn't even in the United States.

It's this strange answer that gives away that it is a probably a machine rather than a person answering the question. Getting a machine to pass itself off as human was one of the key hurdles that Turing believed a machine would need to pass in order to successfully claim the realisation of artificial intelligence. With the creation of the Loebner prize in 1991, monetary prizes were offered for anyone who could create a chatbot that judges could not distinguish from the chat of a human being. Called the

Turing test, many working in AI regard the challenge as something of a red herring. The Loebner prize, in their opinion, has distorted the quest and has proved a distraction from a more interesting goal: creating machine intelligence that is different from our own.

The AI community is beginning to question whether we should be so obsessed with recreating human intelligence. That intelligence is a product of millions of years of evolution and it is possible that it is something that will be very difficult to reverse engineer without going through a similar process. The emphasis is now shifting towards creating intelligence that is unique to the machine, intelligence that ultimately can be harnessed to amplify our very own unique intelligence.

Already the descendants of Deep Blue are performing tasks that no human brain could get anywhere near. Blue Gene can perform 360 trillion operations a second, which compares with the 3 billion instructions per second that an average desktop computer can perform. This extraordinary firepower is being used to simulate the behaviour of molecules at an atomic level to explore how materials age, how turbulence develops in liquids, even the way proteins fold in the body. Protein folding is thought to be crucial to a number of degenerative diseases so these computer simulations could have amazing medical benefits.

But isn't this number-crunching rather than the emergence of a new intelligence? The machine is just performing tasks that have been programmed by the human brain. It may be able to completely outperform my brain in any computational activity but when I'm doing mathematics my brain is doing so much more than just computation. It is working subconsciously, making intuitive leaps. I'm using my imagination to create new pathways which often involve an aesthetic sensibility to arrive at a new mathematical discovery. It is this kind of activity that many of us feel is unique to the human mind and not reproducible by machines.

For me, a test of whether intelligence is beginning to emerge is when you seem to be getting more out than you put in. Machines are human creations yet when what they produce is beginning to surprise the creators then I think you're getting something interesting emerging.

Exciting new research is currently exploring how creative machines can be in music and art. Stravinsky once wrote that he could only be creative by working within strict constraints: "My freedom consists in my moving about within the narrow frame that I have assigned myself for each one of my undertakings." By understanding the constraints that produce exciting music, computer engineers at Sony's Computer Science Laboratory in Paris are beginning to produce machines that create new and unique forms of musical composition. One of the big successes has been to produce a machine that can do jazz improvisation live with human players. The result has surprised those who have trained for years to achieve such a facility.

Other projects have explored how creative machines can be at producing visual art. The Painting Fool is a computer program written by Simon Colton of Imperial College. Not everyone likes the art produced by the Painting Fool but it would be anaemic art if they

did. What's extraordinary is that the programmes in these machines are learning, and changing and evolving so that very soon the programmer no longer has a clear idea of how the results are being achieved and what it is likely to do next. It is this element of getting more out than you put in that represents something approaching emerging intelligence.

For me one of the most striking experiments in AI is the brainchild of the director of the Sony lab in Paris, Luc Steels. He has created machines that can evolve their own language. A population of 20 robots are first placed one by one in front of a mirror and they begin to explore the shapes they can make using their bodies in the mirror. Each time they make a shape they create a new word to denote the shape. For example the robot might choose to name the action of putting the left arm in a horizontal position. Each robot creates its own unique language for its own actions.

The really exciting part is when these robots begin to interact with each other. One robot chooses a word from its lexicon and asks another robot to perform the action corresponding to that word. Of course the likelihood is that the second robot hasn't a clue. So it chooses one of its positions as a guess. If they've guessed correctly the first robot confirms this and if not shows the second robot the intended position.

The second robot might have given the action its own name, so it won't yet abandon its choice, but it will update its dictionary to include the first robot's word. As the interactions progress the robots weight their words according to how successful their communication has been, downgrading those words where the interaction failed. The extraordinary thing is that after a week of the robot group interacting with each other a common language tends to emerge. By continually updating and learning, the robots have evolved their own language. It is a language that turns out to be sophisticated enough to include words that represent the concept of "left" and "right". These words evolve on top of the direct correspondence between word and body position. The fact that there is any convergence at all is exciting but the really striking fact for me is that these robots have a new language that they understand yet the researchers at the end of the week do not comprehend until they too have interacted and decoded the meaning of these new words.

Turing might be disappointed that in his centenary year there are no machines that can pass themselves off as humans but I think that he would be more excited by the new direction artificial intelligence has taken. The AI community is no longer obsessed with reproducing human intelligence, the product of millions of years of evolution, but rather in evolving something new and potentially much more exciting.

Marcus du Sautoy is Simonyi professor for the public understanding of science and a professor of mathematics at the University of Oxford.

▪

Comments

108 comments, displaying Oldest first

g Staff
C Contributor

Open for comments. Sign in or create your Guardian account to join the discussion

Mackname
1 April 2012 1:46AM

The problem is that having stranded in our narrow mentality we are trying to make an intelligent being like our own species.

Recommend? (6)
Responses (2)
Report
Clip | Link

ragingbull
1 April 2012 2:16AM

A lot of technologies that might make amusing iPhone apps but not much else. Why would anyone want a computer portrait?

And why this obsession with intelligence? Our problems as a society are not intellectual ones. They are moral.

Recommend? (28)
Responses (1)
Report
Clip | Link

purplesurfer
1 April 2012 3:51AM

Cold emotionless "Terminators" attempting to destroy humanity....yes life in Cameron's Britain is pretty depressing when you think about it.

Recommend? (17)
Responses (0)
Report
Clip | Link

msmlee
1 April 2012 3:57AM

> but the really striking fact for me is that these robots have a new language that they understand yet the researchers at the end of the week do not comprehend until they too have interacted and decoded the meaning of these new words.

Recommend? (9)
Responses (0)
Report
Clip | Link

Oh, will we be seeing the beginning of a robot species who evolve their own language and means of communication not accessible to humans and who thus form their own communities and who might, even, compete with humans for finite Earth resources once they are intelligent enough to have a desire for self-perpetuation (maybe culturally transmitted and socially learned if their own language gets sophisticated enough for this)? Individually robots may not have the capacity for a conscious self but if philosophy tells us anything, socialised robots may learn to develop their group identity (facilitated first by this development of their own language) and it is a short skip and a jump away from the social "intelligence" that pits their in-group (fellow socialisable AI robots) against an out-group (non-comprehending humans?)!

Maybe the point of the Turing test is precisely so that whatever AI becomes it will try to be more like human beings than not, which will be to the advantage of humans - morally as well as tecnologically, than the alternative scenario in assuming that non-human-like AI would still be subservient to human control or even be sympathetic to humanistic values.

msmlee
1 April 2012 4:03AM

Recommend? (4)
Responses (0)
Report
Clip | Link

On the other hand, of course, if socialised AI robots could reap all the benefits of humanist philosophy whilst curbing the worst excesses of human foibles by not having any of that greed and narcissism as part of their conceptual make-up then maybe they are a better being than humans to inherit this Earth...

Somehow I'm reminded of Studio Gibli's feature-length animation Laputa, in which featured a kind intelligent robot lovingly tending a long-forsaken garden...

Tintinsdog
1 April 2012 5:05AM

Recommend? (2)
Responses (0)
Report
Clip | Link

Gather ye rosebuds while ye may, Old Time is still a-flying; And this same flower that smiles today Tomorrow will be dying.

When a machine can 'think' and write words like that it's getting close to 'intelligence', any kind of intelligence.

I'm not quite sure what this article is saying, beyond the usual clap-trap about AI. Humans are specialist creatures. The majority of us still can't read and write properly. We rely on a tiny minority to propel us forwards.

Blue Gene can perform 360 trillion operations a second, which compares with the 3 billion instructions per second that an average desktop computer can perform.

The ghost in the machine does way better than that every time you make yourself another gin and tonic, which is why people like me will tell you that the mind works at a quantum level.

Tintinsdog
1 April 2012 5:26AM

Recommend? (1)
Responses (0)
Report
Clip | Link

Response to Mackname, 1 April 2012 1:46AM

The problem is that having stranded in our narrow mentality we are trying to make an intelligent being like our own species.

exactement

chechazzo
1 April 2012 5:35AM

Recommend? (19)
Responses (1)
Report
Clip | Link

Response to ragingbull, 1 April 2012 2:16AM

I recommended your comment because it's a fascinating insight, but have you considered that there is a feedback loop between the intellectual and moral problems? Would we not still be burning heretics at the stake if we had not had the scientific insights that promoted the cause of free expression in the minds of the public? Industrial revolution was fuelled by science but at its early stage changed society's perspectives on the importance of science, rational thinking, and freedom of thought and expression.

Another consideration is this: moral problems in most cases boil down to resource allocation, and if you have a resource multiplier such as machine learning, you can begin to solve some resource limitations and in turn, solve the moral problems they cause. A hypothetical example is road rage: if you have machine learning driving cars and automating traffic, you detach the driver from the allocation of road space resources, free his or her time to spend on relaxing or alternative activities, thereby deviating their attention from their fellow drivers.

A more practical example is Google, which is one of the main implementations of machine learning. This has directly affected the ability of humans to join moral causes by reducing the information costs - it's just easier to find out what is happening in the world around you and to take part.

Another practical example is sequencing and annotating the human genome, which leans heavily on machine learning algorithms (especially the latter, which is finding what the genes actually do). Through this we are able to speed up drug discovery for major killer diseases. These diseases create all sorts of moral problems in terms of caring for ill people, society's hesitation to fully fund care, etc etc.

Anyway I'm sure there are readers who can think of much better examples, but the point is that we are at the beginning of something quite useful with this machine learning malarkey. Just don't build Skynet.

Tintinsdog
1 April 2012 6:12AM

Recommend? (0)
Responses (2)
Report
Clip | Link

Response to chechazzo, 1 April 2012 5:35AM

The machines us humans build can't really learn.

One of the fundamental flaws with 'artificial intelligence' is that Turing machines/digital computers are deterministic devices - they have to follow a sequence that reaches a conclusion (aka, algorithms) - whereas all evidence shows that nature does not work like that.

Would anyone care to give me an accurate weather forecast for next month?

chechazzo
1 April 2012 6:30AM

Recommend? (24)
Responses (2)
Report
Clip | Link

Response to Tintinsdog, 1 April 2012 6:12AM

You'd be surprised, machines can really learn.

Machine learning algorithms predict by learning from a whole bunch of environmental variables, which can be evident or even hidden. A basic learning algorithm would adjust how much it weighs the importance of each variable each time it makes a mistake in its prediction. This is learning. The inputs are not known beforehand, the machine is adjusting its model of the environment with each mistake it makes.

The limitation is in the complexity of the algorithms, but then you'd also be surprised by how much can be achieved using conceptually simple, but well structured and intricate techniques. The example in the article above, with the robots creating a new language, may seem simple, but when you scale the same to solving all sorts of problems in science and

economics, you get a really powerful set of applications that learn complex data.

Islero
1 April 2012 6:46AM

Recommend? (0)
Responses (0)
Report
Clip | Link

'animal bodies are like machines? really?' was my qu. when reading the first words of this long, long piece. I did write this just at that moment when I simultaneously read over the box where I'm composing this message that the long, long piece ends 'The example in the article above, with the robots creating a new language, may seem simple, but when you scale the same to solving all sorts of problems in science and economics, you get a really powerful set of applications that learn complex data'. never will a machine be created that can produce animal/man like capabilities without *the* technology. sighs,

---

qertyiou
1 April 2012 7:02AM

Recommend? (7)
Responses (0)
Report
Clip | Link

Response to chechazzo, 1 April 2012 6:30AM

Re simple algorithms: a few years ago I built a program to determine the position of development boundaries, thereby negating the need for corrupt councillors who draw them in such a way as to suit their own ends. I started with a set of questions and values, like one will find in a girls magazine to determine if he loves you or not. The only difference being the weighting of the values could be changed, and some sections had three or more, what amount to fuzzy logic inputs by the user. I had then expected to spend a few weeks messing with some fancy logic to make it work properly, but to my surprise it worked a treat as it was. Actually it did a better job than the planners, who are working from a rats nest of rules. However selling such an idea to civil servants people who will eventually be replaced by programs of this sort is not easy. Having said that, if they continue to misapply mindless rules, their usefulness must surely be limited.

Tintinsdog
1 April 2012 7:17AM

Recommend? (4)
Responses (3)
Report
Clip | Link

Response to chechazzo, 1 April 2012 6:30AM

chechazzo, I'm always fascinated by people's views of the concept of artificial intelligence.

Please don't think I'm being antagonistic as I point out another

fatal flaw in AI: what's now called 'strong AI' is to simulate/mimic intelligence in a computer. Anything that can be explained mathematically can be expressed as an algorithm; but here's the rub: we are a long way, if ever, from understanding what that mysterious thing called 'thought' is, so how on earth are you going to write an algorithm for it?

Granted, people such as yourself no doubt write very nifty algorithms that allow Turing machines to do all kinds of wonderful things, but it's never, ever going to be 'thought' in the human sense.

The biggest problem with AI, though, is that nature is not deterministic, whereas digital computers are.

TonyChinnery
1 April 2012 7:32AM

Recommend? (8)
Responses (1)
Report
Clip | Link

Actually, artificial intelligence, in the sense of thinking, does not exist. Of course those within the industry have a vested interest in pretending it does, but its a well known fact that those with vested interests cannot think straight. Just consider: we don't know what thinking is, we don't know what conscious sensations are, they cannot be measured by any of the parameters used in science (velocity, position, mass etc). I am not talking about the states of our neurons, which can be measured. When I get the sensation of red, I see red without even knowing any of the states of my neurons.
So how can you artificially create something when you don't know what it is?

qertyiou
1 April 2012 7:41AM

Recommend? (4)
Responses (0)
Report
Clip | Link

Response to Tintinsdog, 1 April 2012 7:17AM

While I agree that squidgy stuff is required to build a brain, two of your questions have been answered in part already. Hofstadter in I Am A Strange Loop, suggests that thought is the feedback mechanism of symbolic constructs, which feed back down to the physical construction. It is this process that we call thought. Re determinism, Wolfram in A New Kind of Science, proves that it only takes five rules (each with three variables) to create non deterministic scenarios within a computer.

Recommend? (2)
Responses (0)

Batmanned
1 April 2012 7:49AM

Report

Clip | Link

The question is, could the Painting Fool stand back from its work and have an aesthetic opinion on it?

---

beadster
1 April 2012 9:03AM

Recommend? (3)

Responses (1)

Report

Clip | Link

> Despite Watson's win, it did make some very telling mistakes. In the category 'US cities' contestants were asked: "Its largest airport is named for a world war two hero; its second largest for a world war two battle." The humans responded correctly with "Where is Chicago?" Watson went for Toronto, a city that isn't even in the United States.
>
> It's this strange answer that gives away that it is a probably a machine rather than a person answering the question.

You need to rethink this or prove that these people are not IBM androids

---

sadoldpedant
1 April 2012 9:11AM

Recommend? (12)

Responses (0)

Report

Clip | Link

Response to Tintinsdog, 1 April 2012 7:17AM

@Tintinsdog, do you seriously imagine that proponents of Strong AI are unaware of what you call its "fatal flaw"? In fact, if you read the article you will see that there are ways of programming machines to do things that we don't fully understand. One such is genetic algorithms -- if you write something simple but give it the means to evolve, and then you sit back and see what happens, the result can in principle (and often in practice) behave in ways that are hard to predict and analyse.

IKNOWNOTHING
1 April 2012 9:27AM

Recommend? (10)

Responses (1)

Report

Clip | Link

Common sense might suggest that, rather than attempting to construct machines capable of thought, and in so doing gain understanding of what it is to be "intelligent", we instead try to understand the living intelligences that share this planet with us. The realisation that all living organisms are "intelligent" would

be a start.

Briar
1 April 2012 9:28AM

Recommend? (5)
Responses (1)
Report
Clip | Link

Response to Mackname, 1 April 2012 1:46AM

Yes. Since our definition of intelligence is entirely self referential, how do we know we are really intelligent at all? So certain behaviours enable some to survive at the expense of others - that's just survival of the fittest human-style, today recrafted as free market capitalism, which defines self interest as rational. You could throw in hypocrisy as a defining human characteristic, but is it an essential component of intelligence?

edwardlongshanks
1 April 2012 10:14AM

Recommend? (2)
Responses (1)
Report
Clip | Link

I'm only an average chess player, yet I could beat deep blue easily. As I sat at the board, waiting for the game to commence, I would ask it "why do you play?" deep blue would have no response.
Any human player could answer that question; no machine could.

twiglette
1 April 2012 10:27AM

Recommend? (1)
Responses (2)
Report
Clip | Link

Machines will never have feelings: unless they are grafted together with humans - cyberwoman here we come!!!

Mrdaydream
1 April 2012 11:02AM

Recommend? (1)
Responses (0)
Report
Clip | Link

Lends plausibility to creationism.

respectfulWarrior
1 April 2012 11:14AM

Recommend? (1)
Responses (1)
Report
Clip | Link

i dont believe we wll ever have artificial intelligence as such. Certainly not from a basis, such as a Turing machine. Turing machine are esentially algorithmic devises totally unsuitable to create intelligence. For an indepth mathematical analysis of this read the book: "the emporor's new clothes".

respectfulWarrior
1 April 2012 11:16AM

Recommend? (1)
Responses (1)
Report
Clip | Link

Sorry, that should have said "The Emporor's new Mind"

MichaelBulley
1 April 2012 11:28AM

Recommend? (1)
Responses (0)
Report
Clip | Link

Did Watson want to play Jeopardy? When it comes to language, you have to consider meaning. The French for "I mean" is, tellingly, "je veux dire". Would a machine ever *want* to communicate an idea, let alone have one?

SamJohnson
1 April 2012 11:41AM

Recommend? (1)
Responses (0)
Report
Clip | Link

> lends plausibility to creationism

Perhaps the first benchmark for AI should be creationists, on the grounds that such limited intelligence would be easier to reproduce? CI? But to call any denial of evolution, claims of a young earth etc any form of intelligence is an abuse of the word. Artificial Republicanism? That's an abuse too. Perhaps in time we can get robots to come up with a word for this class of humans?

naishjam
1 April 2012 11:43AM

Recommend? (5)
Responses (1)
Report
Clip | Link

Reading through a lot of the comments here suggests that AI research is deeply misunderstood by the general public, so cheers for this Marcus! Perhaps a documentary on AI research would be great viewing for us all?

My sense is that state of the art AI is exceptionally good as a problem solving technology, but that techniques must be trained or developed for specific domains. For example, Deep Thought was exceptionally good at chess playing, Watson at answering general knowledge questions. However, if these two programs had been put in one another's contexts (Watson against Kasparov, Deep Thought against the Jepordy champions) then the results would have been much worse than human novices.

One of the challenges in future AI research, therefore, is to build applications that can generalise their intelligence across problem domains. Better machine learning techniques are a crucial step

in this direction, but also we need a much better understanding of knowledge representation. AI techniques today typically use domain-specific knowledge structures; how can we generalise these?

However, I believe there is a deeper problem. In AI research to date (and similarly in cognitive science - there is a strong historical link between the two) the focus has been primarily on problem solving. Simon and Newell emphasise this restricted scope in their seminal book: "Human Problem Solving". Clearly the human mind has functionality that goes far beyond problem solving, although this is undoubtedly one of the basic ingredients. We need to take a deeper look at this broader functionality. Some have mentioned emotions, and this seems to be an under-studied area in cognitive science. How can we build machines that can simulate emotion and senitiment? Some researchers at Manchester University are currently working on detecting sentiment within natural language corpi. This is a step in the right direction, but it is a small one.

Finally, this year's Turing lecture at Manchester University raised an intriguing question. Professor Ray Dolan (UCL), was discussing how the mathematics developed at Bletchley Park for code breaking could be used to model certain phenomena in neuroscience. In the Q&A at the end of his talk, however, he suggested at a trend of ever lower levels of abstraction in the history of psychology. Early efforts, such as behaviourism, viewed the mind as a black box with inputs and outputs. Cognitive science seeks to describe detailed thought processes within the mind. Neuroscience seeks to study the structure of the brain and how this produces those thought processes. A new and emerging area of research, however, is to try and understand the chemical and electrical basis for these processes: how do different chemical reactions within the brain ultimately create human thought?

This, I would suggest, seems to raise the question as to whether or not a binary computer is capable of producing true intelligence, or whether, at some point in the future of AI research, the challenge will prove to be fundamentally a hardware, rather than a software question.

naishjam
1 April 2012 12:20PM

Recommend? (16)
Responses (0)
Report
Clip | Link

Response to Tintinsdog, 1 April 2012 6:12AM

You fundamentally misunderstand the concept of determinism.

The fact that a human being can't predict the weather accurately has nothing to do with a lack of determinism in weather systems. It has to do with the fact that our models are crappy, and the fact that we could never gather all the data required to perfectly predict the weather. In short, you confuse complexity with non-determinism. Mistake.

chechazzo
1 April 2012 12:21PM

Recommend? (1)
Responses (0)
Report
Clip | Link

Response to Tintinsdog, 1 April 2012 7:17AM

Tintinsdog, I see your point, and it does beggar belief, but thought is really a search process - you are looking into what possibilities exist, assigning a value to them. Probabilistic algorithms, which is what machine learning essentially is, encompass this search process. They look at all the data points in 2D, 3D, even infinite dimensions, and they look tat what predictive value or explanatory function they can come up with, by a sequence of searches across all the possibilities.

That robot example in the article is very, very powerful. It kind of embodies the whole concept of machine learning. By simple search and factoring in mistakes, machine learning algorithms find new discoveries that, because they were programmed to be random, were unpredictable.

It's a random walk that's built up of small components.

naishjam
1 April 2012 12:27PM

Recommend? (3)
Responses (1)
Report
Clip | Link

Response to TonyChinnery, 1 April 2012 7:32AM

Actually, we have a pretty clear conception of what thought is. It's simply an emergent property of a highly complex system. The real question is not what thought is, but what the function of each neuron within the brain might be, and how, collectively, these neurons combine to produce higher-level functions such as memory, language parsing, reasoning, and so on.

You don't have to know the states of each individual neuron in order to sense red. But if one wanted to understand how the sensation of "redness" is produced, then one would ideally have a complete layout of the neural connections within your brain and the state of each neuron. The task then is to look at how other people sense red, what happens when you sense green (what changes?), and many more questions so that we can work out the role played by each individual neuron in the process of creating

this sensation.

naishjam
1 April 2012 12:30PM

Recommend? (6)
Responses (1)
Report
Clip | Link

Response to IKNOWNOTHING, 1 April 2012 9:27AM

In what sense are "all living organisms 'intelligent'"? What about plants, for starters? What kind of "intelligence" do they exhibit?

naishjam
1 April 2012 12:34PM

Recommend? (0)
Responses (0)
Report
Clip | Link

Response to Briar, 1 April 2012 9:28AM

I don't really follow your point. This isn't really how anything is defined. Definitions arise organically because people group a set of objects, phenomena, events, etc together that share some similar properties and then attach a label to that class. It's an extremely well-studied process known as "chunking". So intelligence is defined broadly as the set of capabilities exhibited by human thought. How can there be any concept of intelligence beyond that which is defined by humans, therefore?

vakibs
1 April 2012 12:36PM

Recommend? (2)
Responses (2)
Report
Clip | Link

Machine intelligence is not evolving.

Our conceptions (and misconceptions) about machine intelligence are evolving.

naishjam
1 April 2012 12:38PM

Recommend? (12)
Responses (0)
Report
Clip | Link

Response to edwardlongshanks, 1 April 2012 10:14AM

Nonsense. I'm not totally familiar with the specifics of deep blue, but one of two things would happen. Either it would ignore your question and get on with the business of playing chess, or, at worst, it would give a convoluted response to your question and then get on with the business of playing chess. Unless you claim to be capable of beating a chess grandmaster, then you would be beaten in either case.

As for no machine being capable of answering your question, of course it could. A simple example of this would be a machine that was simply programmed to generate an answer specifically

to that question. At the other end of the spectrum would be a machine that was programmed to contemplate its own nature, and to generate a response accordingly. In either case, it's possible to think of a machine that could answer your question.

naishjam
1 April 2012 12:39PM

Recommend? (3)
Responses (0)
Report
Clip | Link

Response to twiglette, 1 April 2012 10:27AM

Never say never. It's difficult to imagine how this might be achieved, but it's possible, in principle, to conceive of such a machine.

naishjam
1 April 2012 12:45PM

Recommend? (2)
Responses (1)
Report
Clip | Link

Response to respectfulWarrior, 1 April 2012 11:14AM

Neuroscientists commonly think of the models that they build as the "algorithms of the brain". Don't be so dismissive of algorithms.

Whether or not today's binary computers are sufficient to produce sophisticated models of intelligence is another question; but the Turing machine does not necessarily imply a modern binary machine.

On another note, the idea of symbollic processing is a highly contentious one in AI. Simon and Newell's seminal book on cognitive science stated that the view of reasoning as symbollic processing was a useful abstraction to help us understand what is happening inside the brain in the absence of a deeper understanding of brain physiology. That deeper understanding is now emerging - neuroscience - but is not entirely superceding cognitive science. The two are useful in conjunction with one another.

My point is that algorithms are the models that we build in all scientific disciplines to understand a vast array of phenomena. Human intelligence, we assume, is simply an emergent property of a highly complex physical and chemical system, and therefore why should it be treated any differently to, say, a weather system or a solar system?

klang
1 April 2012 12:49PM

Recommend? (1)
Responses (0)
Report

Give me skynet, over jazz.

Clip | Link

dirkbruere
1 April 2012 2:13PM

Recommend? (1)
Responses (0)
Report
Clip | Link

The thing about the Turing Test is that it has increasingly come to mean a piece of software that can fool computer scientists and psychologists. The best can, right now, probably fool most ordinary people most of the time..
Also, it has always seemed to me to be naive to expect Human level intelligence from machines of less than Human brain processing power. At present it could be argued that the best supercomputers *might* overlap with the lower estimates of brain power. However, within 6 years there will be computers most definitely in this range (exascale computing).

NottyImp
1 April 2012 2:19PM

Recommend? (3)
Responses (0)
Report
Clip | Link

> Gather ye rosebuds while ye may, Old Time is still a-flying; And this same flower that smiles today Tomorrow will be dying.
>
> When a machine can 'think' and write words like that it's getting close to 'intelligence', any kind of intelligence.

I think the key will be when an AI understands words like that and why they can be beautiful.

---

respectfulWarrior
1 April 2012 2:42PM

Recommend? (2)
Responses (2)
Report
Clip | Link

Response to naishjam, 1 April 2012 12:45PM

I understand where you are coming from. You are actually replicating exactly the same argument Turing did. He thought that our bodies, brain etc are system that run algorithms, the only difference between them and a Turing machine is that we do not understand how they arrive at their output.

However this concept is flawed. I could try an explain how, but I probably would not make sense. Let me point you to the excellent book I mentioned above written by a world renowned physicist. Basically he argues from a mathematical and physics point of view that our brains do not run algorithmicly and are therefore **cannot** be modeled on a turing machine.

> but the Turing machine does not necessarily imply a modern binary machine

Coming from a computer science background I'd say that is not true. In essence a Turing machine is a model of a computer. I.e any grammar that is Turing acceptable is also acceptable by a Binary pc and vice versa. At the moment no form of computation is know that is is super-Turing. If we accept (as I do) that intelligence is not going to be created from a computation form with equal power to the Turing machine, then we are still a long way from AI (If such a thing is possible, which I doubt).

dirkbruere
1 April 2012 2:59PM

Recommend? (7)
Responses (0)
Report
Clip | Link

Response to respectfulWarrior, 1 April 2012 2:42PM

Penrose is in a tiny minority when it comes to explaining how the brain functions

---

truthandreason
1 April 2012 3:36PM

Recommend? (1)
Responses (0)
Report
Clip | Link

Response to vakibs, 1 April 2012 12:36PM

Your point is really only about the definition of terms, but if you accept that we are machines, then progress towards artificial machine intelligence can be interpreted as a part of the natural process of evolution. Once AI was achieved, the ability to reproduce without our help would probably soon follow, at which point they would become a separate life form - though we could still argue about whether they were really thinking.

That said, while the example of Sony's language evolving robots is encouraging, I would not hold my breath waiting for that system to start allowing humans to communicate with machines via human natural language, even for the modest purpose of replacing conventional programming languages (let alone writing poetry - although a few rhyming rules plus random word generation and some word n-gram statistics could approximate a lot of what passes for poetry). Experiments with human-like pre-verbal learning by machines which interact with the environment have often shown some signs of learning, but so far never seem to get very far. Something is missing.

One problem for these language learning robots is that they are limited to discovering terms based on concepts which are supported by the instruction set which operates their sensor-

motor system. As well as ideas inherent in Euclidean geometry, they could also develop concepts involving sound (SPL values), colour (RGB values) and simple mathematics (logic, arithmetic and geometry). However, they will not be able to learn to correctly communicate information concerning any of the many other human sensations (including most of the inner sensations or emotions) until these have also been quantified, which could take some time.

dirkbruere
1 April 2012 3:36PM

By the time a machine can fully pass the Turing Test it will only do so by playing as dumb as a Human.

Recommend? (2)
Responses (0)
Report
Clip | Link

klassekanp
1 April 2012 3:37PM

Intelligent machines,who pulls the plug?

Recommend? (0)
Responses (1)
Report
Clip | Link

AhabTRuler
1 April 2012 3:43PM

What could possibly go wrong?

Recommend? (1)
Responses (0)
Report
Clip | Link

ConorLane
1 April 2012 3:56PM

Response to vakibs, 1 April 2012 12:36PM

We create the machine intelligence as we create the machines, therefore, when we create a more intelligent machine, machine intelligence evolves.
If we weren't actually making the machines and it was all entirely theory, you'd be right, but the fact that they exist means they evolve when new ones are built.

Recommend? (0)
Responses (0)
Report
Clip | Link

ConorLane
1 April 2012 3:58PM

Response to twiglette, 1 April 2012 10:27AM

How can you possibly make that claim?

Recommend? (1)
Responses (0)
Report
Clip | Link

dirkbruere
1 April 2012 4:01PM

Recommend? (1)
Responses (0)
Report
Clip | Link

Response to klassekanp, 1 April 2012 3:37PM

> Intelligent machines,who pulls the plug?

Internet, who pulls the plug?

Chronos
1 April 2012 4:03PM

Recommend? (3)
Responses (2)
Report
Clip | Link

Response to respectfulWarrior, 1 April 2012 11:16AM

The Emperor's New Mind isn't a particularly good book on the topic. Penrose seems to have a poor grasp of both the efforts of AI research and the philosophical arguments that surround the idea of intelligent machines. He quotes Searle at length without acknowledging the poverty of his Chinese Room argument which essentially boils down to "humans are special because humans are special".

Throw in his dodgy arguments about the quantum brain which are a bit left-field, to say the least and I think you could do a lot better than read his book. Even if the brain was reliant on behaviour of structures at the quantum level, why couldn't a machine be built that made use of the same properties?

Short of getting into religious arguments about the need for a soul, I have never heard argument for why, in theory, a machine couldn't do exactly what a human brain does.

Open for comments. Sign in or create your Guardian account to join the discussion