

**The
Economist**

Subscribe



■ Print edition | Science and technologyAug 30th 2018

SOFTWARE that can learn is changing the world, but it needs supervision. Humans provide such oversight in two ways. The first is to show machine-learning

algorithms large sets of data that describe the task at hand. Labelled pictures of cats and dogs, for instance, allow an algorithm to learn to discriminate between the two. The other form of supervision is to set a specific goal within a highly structured environment, such as achieving a high score in a video game, and then let the algorithm try out lots of possibilities until it finds one that achieves the objective.

These two approaches to “supervised learning” have led to breakthroughs in artificial intelligence. In 2012 a group of researchers from the University of Toronto used the first method to build AlexNet, a piece of software that in a competition recognised one in ten more images than its closest competitors. In 2015 researchers at DeepMind, a British AI firm owned by Alphabet, used the second method to teach an algorithm to play Atari video games at superhuman levels, an advance that led to later triumphs at Go, a board game.

ADVERTISING

[inRead](#) invented by Teads

Get our daily newsletter

Upgrade your inbox and get our Daily Dispatch and Editor's Picks.

Latest stories

“The Other Side of the Wind” dazzles, puzzles and intrigues

PROSPERO

Are gender quotas good for business?

THE ECONOMIST EXPLAINS

A fillip for Ukrainian hopes of spiritual independence

ERASMUS

See more

Such breakthroughs underpin much of the excitement in AI today. But supervised learning has weaknesses. Human guidance is expensive, involving manual tasks such as labelling data or designing virtual environments. Once complete, that guidance cannot be used for other lessons. Nor is supervised learning very realistic. The real world does not often label things or provide explicit signals about the progress that a learner is making. Both AlexNet and DeepMind's game-playing agents require millions or billions of examples or simulations to work on—and powerful computers that use lots of electricity. “If you are going to do this with every new [training] task, you are going to need dozens of nuclear power plants doing nothing else,” says Pierre-Yves Oudeyer, an AI researcher at Inria, the French national institute for computer science in Paris.

I wonder...

If AI is going to really take off, then something more is needed. Dr Oudeyer says that requirement is driving interest in one of the fundamental mechanisms used by humans to learn about the world: curiosity. Instead of training algorithms with functions created by humans, Dr Oudeyer and others have spent the past 20 years developing artificial agents that use their own intrinsic reward systems to inspect the world around them and gather data. Such work is starting to come into its own.

The first generation of curious AI used “prediction error” to motivate the agent. The software would explore the environment it was required to study, whether physical or virtual, looking for things that deviated significantly from what it predicted it would find. In other words, it searched for novel data. Using prediction error worked, but it had a big flaw. An agent looking at passing cars, for instance, might become obsessed with the sequence of the colours of each car, because its prediction about what colour would come next is almost always wrong. That serves no useful purpose. Nor would a curious robot repeatedly throwing itself down the stairs for the sheer informatic thrill of it, rather than learning to walk its way down.

This problem is fixed by concentrating on the rate at which an agent's prediction error changes, rather than on the error itself. Using this process, a robot watching the sun rise and set will see its prediction errors start high but decrease over time, as it learns about the actual properties of a physical system. Using the rate of change in a prediction-error system as a signal for the agent to move on to something else is equivalent to giving it a boredom threshold. If the robot trying to

work out the pattern of colours of passing cars were to use such a system it would make errors at a steady rate, and get bored.

Dr Oudeyer has tried out his curiosity algorithms in practical pursuits. In June his group tested one on 600 primary schoolchildren at a number of public and private schools in the Aquitaine region of France. The idea was to model each child's learning in mathematics and present each pupil with exercises in a way that optimises their learning. The system, called KidLearn, treats each child as its own curious agent, and adapts the learning content to suit that child's level of understanding and progress. Unlike other software, KidLearn does not rely on data gathered from other children as its guide but is tuned primarily by a child's curiosity. Dr Oudeyer's researchers will shortly report on how well their system performs.

Researchers in Silicon Valley have been embracing curiosity, too. In a recent paper Deepak Pathak and his colleagues at the University of California, Berkeley and OpenAI, a non-profit research firm backed by Elon Musk, showed that curiosity-driven learning works well across a range of virtual environments, despite the fact that their agent was told nothing about the video games it was playing, nor given any signal when it died in the game or reached a higher level.

The curious agent displayed some interesting behaviour. It learned to achieve higher scores in Breakout, a block-breaking game, because the higher the score the more complicated the pattern of blocks becomes, and the more the agent's curiosity was satisfied. When two curious agents played Pong they learned to rally so long that they crashed the game because they found rallying was more interesting than winning. Dying is also boring. "The agent avoids dying in the games since that brings it back to the beginning of the game, an area it has already seen many times and where it can predict the dynamics well," the researchers said in a recent paper.

There are other ways to bestow machines with the urge to explore. Kenneth Stanley, a researcher at Uber's AI lab in San Francisco, mimics evolution. His system starts with a set of random algorithms, chooses the one that looks good for the task at hand, then generates a set of algorithms derived from it. Eventually it arrives at an algorithm that is most suited for the job. Evolution, Mr Stanley notes, can yield serendipitous results that goal-driven optimisation cannot. Biological evolution was not explicitly curious about flying, and yet it still managed to come up with birds.

All this suggests that a more complete set of learning algorithms is emerging. Artificial agents that are driven by curiosity or evolution could look after the earlier stages of learning. They are also more suited to sparse environments devoid of much data. Once something interesting has been found, supervised learning could take over to ensure particular features are learned exactly. Last week, in a video-game competition in Vancouver, AI agents created by OpenAI, using the most advanced supervised-learning techniques available, were crushed by humans in DOTA 2, a strategy game. More curious modes of learning might have helped AI play the long-term parts of the game, in which there are few reward signals and no changes in score.

"I'd hate to die twice. It's so boring," were the death-bed words of Richard Feynman, an American theoretical physicist. His last salute to curiosity followed a lifetime probing the inner workings of the universe, finding new things to model and to understand. That very human inclination can motivate machines as well as man.

This article appeared in the Science and technology section of the print edition under the headline "Head full of brains, shoes full of feet"

 **Print edition | Science and technology**

Aug 30th 2018

Reuse this content

About *The Economist*



Redefining helpful finance

Rules on bank lending in poor neighbourhoods are being rethought

America's Community Reinvestment Act is sensitive, but still up for revision

<https://www.economist.com/science-and-technology/2018/08/30/a-sense-of-curiosity-is-helpful-for-artificial-intelligence>

All the young prudes

Naked Europe covers up

A man out of time

Rahm Emanuel will not seek re-election as mayor of Chicago

DEMOCRACY IN AMERICA

Open Future: Debate

Should affirmative action be scrapped?

Tell us what you think of Economist.com

[Leave feedback](#)

Need assistance with your subscription?

[Contact us](#)

[Classified ads](#)

[Subscribe](#)

[Group subscriptions](#)

[Contact us](#)

[Help](#)

Keep updated



Sign up to get more from *The Economist*

Get 3 free articles per week, daily newsletters and more.

Email address

Sign up

About *The Economist*

Advertise

Reprints

Careers

Media Centre

Terms of Use

Privacy

Cookie Policy

Manage Cookies

Accessibility

Modern Slavery Statement

Copyright © The Economist Newspaper Limited 2018. All rights reserved.