

NEWS IN FOCUS



CANINE PSYCHIATRY Dogs provide genetic clues to human disorders **p.446**

PHYSICS Debate over meaning of Stephen Hawking's latest paper **p.448**

CLIMATE Developing nations struggle to keep carbon accounts **p.450**

CONSERVATION Songbird killing for restaurants becomes a hot issue in Cyprus **p.452**

NATURE VIDEO



Go, a complex game popular in Asia, has frustrated the efforts of artificial-intelligence researchers for decades.

ARTIFICIAL INTELLIGENCE

Google masters Go

Deep-learning software excels at complex ancient board game.

BY ELIZABETH GIBNEY

A computer has beaten a human professional for the first time at Go — an ancient board game that has long been viewed as one of the greatest challenges for artificial intelligence (AI).

The best human players of chess, draughts and backgammon have all been outplayed by computers. But a hefty handicap was needed for computers to win at Go. Now Google's London-based AI company, DeepMind, claims that its machine has mastered the game.

DeepMind's program AlphaGo beat Fan Hui, the European Go champion, five times out of five in tournament conditions, the firm

reveals in research published in *Nature* on 27 January¹. It also defeated its silicon-based rivals, winning 99.8% of games against the current best programs. The program has yet to play the Go equivalent of a world champion, but a match against South Korean professional Lee Sedol, considered by many to be the world's strongest player, is scheduled for March. "We're pretty confident," says DeepMind co-founder Demis Hassabis.

"This is a really big result, it's huge," says Rémi Coulom, a programmer in Lille, France, who designed a commercial Go program called Crazy Stone. He had thought computer mastery of the game was a decade away.

The IBM chess computer Deep Blue, which

famously beat grandmaster Garry Kasparov in 1997, was explicitly programmed to win at the game. But AlphaGo was not preprogrammed to play Go: rather, it learned using a general-purpose algorithm that allowed it to interpret the game's patterns, in a similar way to how a DeepMind program learned to play 49 different arcade games².

This means that similar techniques could be applied to other AI domains that require recognition of complex patterns, long-term planning and decision-making, says Hassabis. "A lot of the things we're trying to do in the world come under that rubric." Examples are using medical images to make diagnoses or treatment plans, and improving climate-change models. ▶

► In China, Japan and South Korea, Go is hugely popular and is even played by celebrity professionals. But the game has long interested AI researchers because of its complexity. The rules are relatively simple: the goal is to gain the most territory by placing and capturing black and white stones on a 19 × 19 grid. But the average 150-move game contains more possible board configurations — 10^{170} — than there are atoms in the Universe, so it can't be solved by algorithms that search exhaustively for the best move.

ABSTRACT STRATEGY

Chess is less complex than Go, but it still has too many possible configurations to solve by brute force alone. Instead, programs cut down their searches by looking a few turns ahead and judging which player would have the upper hand. In Go, recognizing winning and losing positions is much harder: stones have equal values and can have subtle impacts far across the board.

To interpret Go boards and to learn the best possible moves, the AlphaGo program applied deep learning in neural networks — brain-inspired programs in which connections between layers of simulated neurons are strengthened through examples and experience. It first studied 30 million positions from expert games, gleaning abstract information on the state of play from board data, much as

other programmes categorize images from pixels (see *Nature* 505, 146–148; 2014). Then it played against itself across 50 computers, improving with each iteration, a technique known as reinforcement learning.

The software was already competitive with the leading commercial Go programs, which select the best move by scanning a sample of simulated future games. DeepMind then combined this search approach with the ability to pick moves and interpret Go boards — giving

“Deep learning is killing every problem in AI.”

AlphaGo a better idea of which strategies are likely to be successful. The technique is “phenomenal”, says Jonathan Schaeffer, a computer scientist at the University of Alberta in Edmonton, Canada, whose software Chinook solved³ draughts in 2007. Rather than follow the trend of the past 30 years of trying to crack games using computing power, DeepMind has reverted to mimicking human-like knowledge, albeit by training, rather than by being programmed, he says. The feat also shows the power of deep learning, which is going from success to success, says Coulom. “Deep learning is killing every problem in AI.”

AlphaGo plays in a human way, says Fan. “If no one told me, maybe I would think the player was a little strange, but a very strong player, a real person.” The program seems to

have developed a conservative (rather than aggressive) style, adds Toby Manning, a lifelong Go player who refereed the match.

Google's rival firm Facebook has also been working on software that uses machine learning to play Go. Its program, called *darkforest*, is still behind commercial state-of-the-art Go AI systems, according to a November preprint⁴.

Hassabis says that many challenges remain in DeepMind's goal of developing a generalized AI system. In particular, its programs cannot yet usefully transfer their learning about one system — such as Go — to new tasks; a feat that humans perform seamlessly. “We've no idea how to do that. Not yet,” Hassabis says.

Go players will be keen to use the software to improve their game, says Manning, although Hassabis says that DeepMind has yet to decide whether it will make a commercial version.

AlphaGo hasn't killed the joy of the game, Manning adds. Strap lines boasting that Go is a game that computers can't win will have to be changed, he says. “But just because some software has got to a strength that I can only dream of, it's not going to stop me playing.” ■

SEE EDITORIAL P.437

1. Silver, D. *et al. Nature* 529, 484–489 (2016).
2. Mnih, V. *et al. Nature* 518, 529–533 (2015).
3. Schaeffer, J. *et al. Science* 317, 1518–1522 (2007).
4. Tian, Y. & Zhu, Y. Preprint at arXiv <http://arxiv.org/pdf/1511.06410.pdf> (2015).

GENOMICS

Dog DNA probed for clues to human psychiatric ills

Project will compare gene data to owners' assessments of how their companions behave.

BY HEIDI LEDFORD

Addie plays hard for an 11-year-old greater Swiss mountain dog — she will occasionally ignore her advanced years to hurl her 37-kilogram body at an unwitting house guest in greeting. But she carries a mysterious burden: when she was 18 months old, she started licking her front legs aggressively enough to wear off patches of fur and draw blood.

Addie has canine compulsive disorder — a condition that is thought to be similar to human obsessive-compulsive disorder (OCD). Canine compulsive disorder can cause dogs to chase their tails for hours on end, or to suck on a toy or body part so compulsively that it interferes with their eating or sleeping.

Addie may soon help researchers to

determine why some dogs are more prone to the disorder than others. Her owner, Marjie Alonso of Somerville, Massachusetts, has enrolled her in a project called Darwin's Dogs, which aims to compare information about the behaviour of thousands of dogs against the animals' DNA profiles. The hope is that genetic links will emerge to conditions such as canine compulsive disorder and canine cognitive dysfunction — a dog analogue of dementia and possibly Alzheimer's disease. The project organizers have enrolled 3,000 dogs so far, but hope to gather data from at least 5,000, and they expect to begin analysing DNA samples in March.

“It's very exciting, and in many ways it's way overdue,” says Clive Wynne, who studies canine behaviour at Arizona State University in Tempe.

Researchers have long struggled to find

genetic links to human psychiatric disorders by analysing DNA samples from thousands of people. Those efforts have in recent years met with some success in schizophrenia and depression. But for some conditions, including OCD, not a single robust genetic link has been sifted from the background noise of normal genetic variation.

Human studies are difficult in part because the species is so genetically diverse, says Wynne. Dogs, however, are more genetically homogeneous. Selected over thousands of years for particular characteristics, they display less genetic variation than do humans. Pure-bred dogs, in particular, have been rendered highly genetically consistent to achieve a homogenous appearance and behaviour.

Dogs also live side-by-side with humans, which some think can make them a better