stimulation ultimately improved the athletes' jumping force by 70% and their coordination by 80%, compared with the sham group, Halo announced in February.

Troy Taylor, high-performance director for the USSA, is encouraged by the results — but concedes that they are preliminary.

PUSHING THE LIMITS OF ENDURANCE

Another study, presented on 7 March at the Biomedical Basis of Elite Performance meeting in Nottingham, UK, suggests that tDCS may reduce the perception of fatigue. Sports scientist Lex Mauger of the University of Kent in Canterbury, UK, and his colleagues found that stimulating the motor-cortex region that controls leg function allows cyclists to pedal longer without feeling tired.

The researchers stimulated the brains of 12 untrained volunteers before directing the athletes to pedal stationary bicycles until they were exhausted. Every minute, they asked the cyclists to rate their level of effort.

Volunteers who received tDCS were able to pedal two minutes longer, on average, than were those who were given a sham treatment. They also rated themselves as less tired. But there was no difference in heart rate or the lactate level in the muscles between the treatment and control groups. This suggests that changes in brain perception, rather than muscle pain or other body feedback, drove the improved performance.

Alexandre Okano, a biological engineer at Federal University of Rio Grande do Norte in Brazil, found similar increases in cyclists' performance when he stimulated the brain's temporal cortex, which is involved in body awareness and in automatic functions such as breathing[2]. This suggests that the temporal and motor cortices are connected in ways that are not understood, or that tDCS does not target locations in the brain precisely, Okano says.

These results support the notion that the brain manages exertion by collating feedback from the body and then slowing muscles to prevent fatigue, says Dylan Edwards, a neurophysiologist at Burke Medical Research Institute in White Plains, New York[3]. "Even when you think you're exercising as hard as you can, there is always some reserve of ability," he says.

TRICKY TESTS

But Horvath cautions that little is known about the long-term effects of stimulating the brain. And others are sceptical of the technique's potential to increase performance. Vincent Walsh, a neuroscientist at University College London, notes that the methods used in tDCS studies often differ between research groups — and might not always be optimized.

For instance, the fairly intense amount of electricity that Mauger's team used has been shown to sometimes have complex and unintended effects on the brain's activity[4].

Replicating such experiments is difficult because of variations in how people respond to brain stimulation. Some people do not respond at all; others might respond only when stimulated in a certain way. And even an individual's response can differ from day to day. Edwards says that it is important to map these differences if tDCS is to be used therapeutically or for other purposes. "We're moving toward customized prescription of brain stimulation," he says.

Nonetheless, the use of tDCS in sports is only likely to increase. Stimulating the motor cortex, for instance, seems to increase dexterity, so videogamers have been quick to take up the technique. And it is increasingly easy to acquire stimulation devices; Halo has begun to market its equipment for the express purpose of increasing athletic performance.

Taylor compares the use of brain stimulation by athletes to eating carbohydrates ahead of an athletic event, in the hopes of boosting endurance. "It piggybacks on the ability to learn," he says. "It's not introducing something artificial into the body."

But Edwards worries that the availability of tDCS devices will tempt athletes to try "brain doping", in part because there is no way to detect its use. "If this is real," he says, "then absolutely the Olympics should be concerned about it." ■

1. Cogiamanian, F. *et al. Eur. J. Neurosci.* **26,** 242–249 (2007).
2. Okano, A. H. *et al. Br. J. Sports Med.* **49,** 1213–1218 (2015).
3. Noakes, T. D. *Sports Med.* **37,** 374–377 (2007).
4. Batsikadze, G., Moliadze, V., Paulus, W., Kuo, M.-F. & Nitsche, M. A. *J. Physiol.* **591,** 1987–2000 (2013).

ARTIFICIAL INTELLIGENCE

What Google's winning Go algorithm will do next

AlphaGo's techniques could have broad uses, but moving beyond games is a challenge.

BY ELIZABETH GIBNEY

Following the defeat of one of its finest human players, the ancient game of Go has joined the growing list of tasks at which computers perform better than humans. In a 6-day tournament in Seoul, watched by a reported 100 million people around the world, the computer algorithm AlphaGo, created by the Google-owned company DeepMind, beat Go professional Lee Sedol by 4 games to 1. The complexity and intuitive nature of the ancient board game had established Go as one the greatest challenges in artificial intelligence (AI). Now the big question is what the DeepMind team will turn to next.

AlphaGo's general-purpose approach — which was mainly learned, with a few elements crafted specifically for the game — could be applied to problems that involve pattern recognition, decision-making and planning. But the approach is also limited. "It's really impressive, but at the same time, there are still a lot of challenges," says Yoshua Bengio, a computer scientist at the University of Montreal in Canada.

Lee, who had predicted that he would win the Google tournament in a landslide, was shocked by his loss. In October, AlphaGo beat European champion Fan Hui. But the version of the program that won in Seoul is significantly stronger, says Jonathan Schaeffer, a computer scientist at the University of Alberta in Edmonton, Canada, whose Chinook software mastered draughts in 2007: "I expected them to use more computational resources and do a lot more learning, but I still didn't expect to see this amazing level of performance."

The improvement was largely down to the fact that the more AlphaGo plays, the better it gets, says Miles Brundage, a social scientist at Arizona State University in Tempe, who studies trends in AI. AlphaGo uses a brain-inspired architecture known as a neural network, in which connections between layers of simulated neurons strengthen on the basis of experience. It learned by first studying 30 million Go positions from human games and then improving by playing itself over and over again, a technique known as reinforcement learning. Then, DeepMind combined AlphaGo's ability

KIM HONG-JI/TPX/REUTERS

Professional Go player Lee Sedol (centre) after his 4–1 defeat by Google's AlphaGo algorithm.

to recognize successful board configurations with a 'look-ahead search', in which it explores the consequences of playing promising moves and uses that to decide which one to pick.

Next, DeepMind could tackle more games. Most board games, in which players tend to have access to all information about play, are now solved. But machines still cannot beat humans at multiplayer poker, say, in which each player sees only their own cards. The DeepMind team has expressed an interest in tackling Starcraft, a science-fiction strategy game, and Schaeffer suggests that DeepMind devise a program that can learn to play different types of game from scratch. Such programs already compete annually at the International General Game Playing Competition, which is geared towards creating a more general type of AI. Schaeffer suspects that DeepMind would excel at the contest. "It's so obvious, that I'm positive they must be looking at it," he says.

DeepMind's founder and chief executive Demis Hassabis mentioned the possibility of training a version of AlphaGo using self-play alone, omitting the knowledge from human-expert games, at a conference last month. The firm created a program that learned to play less complex arcade games in this manner in 2015. Without a head start, AlphaGo would probably take much longer to learn, says Bengio — and might never beat the best human. But it's an important step, he says, because humans learn with such little guidance.

DeepMind, based in London, also plans to venture beyond games. In February the company founded DeepMind Health and launched a collaboration with the UK National Health Service: its algorithms could eventually be applied to clinical data to improve diagnoses or treatment plans. Such applications pose different challenges from games, says Oren Etzioni, chief executive of the non-profit Allen Institute for Artificial Intelligence in Seattle, Washington. "The universal thing about games is that you can collect an arbitrary amount of data," he says — and that the program is constantly getting feedback on what's a good or bad move by playing many games. But, in the messy real world, data — on rare diseases, say — might be scarcer, and even with common diseases, labelling the consequences of a decision as 'good' or 'bad' may not be straightforward.

Hassabis has said that DeepMind's algorithms could give smartphone personal assistants a deeper understanding of users' requests. And AI researchers see parallels between human dialogue and games: "Each person is making a play, and we have a sequence of turns, and each of us has an objective," says Bengio. But they also caution that language and human interaction involve a lot more uncertainty.

DeepMind is fuelled by a "very powerful cocktail" of the freedoms usually reserved for academic researchers, and by the vast staff and computing resources that come with being a Google-backed firm, says Joelle Pineau, a computer scientist at McGill University in Montreal. Its achievement with Go has prompted speculation about when an AI will have a versatile, general intelligence. "People's minds race forward and say, if it can beat a world champion it can do anything," says Etzioni. But deep reinforcement learning remains applicable only in certain domains, he says: "We are a long, long way from general artificial intelligence."

DeepMind's approach is not the only way to push the boundaries of AI. Gary Marcus, a neuroscientist at New York University in New York City, has co-founded a start-up company, Geometric Intelligence, to explore learning techniques that extrapolate from a small number of examples, inspired by how children learn. In its short life, AlphaGo probably played hundreds of millions of games — many more than Lee, who still won one of the five games against AlphaGo. "It's impressive that a human can use a much smaller quantity of data to pick up a pattern," says Marcus. "Probably, humans are much faster learners than computers." ■

MORE ONLINE

THE GO FILES

GOOGLE DEEPMIND

Updates from a battle of man versus machine over the Go board go.nature.com/j26o3w

MORE NEWS

- Horse-sized dinosaur sheds light on *T. rex*'s origins go.nature.com/ihlkx2
- Peculiar pattern found in 'random' prime numbers go.nature.com/hz11qo
- Genetic study links 'good' cholesterol mutation to heart disease go.nature.com/mjils9

NATURE PODCAST

Retrieving lost memories; nailing China's emissions; and is Alzheimer's transmissible? nature.com/nature/podcast