

# Google's AI reasons its way around the London Underground

DeepMind's latest technique uses external memory to solve tasks that require logic and reasoning — a step toward more human-like AI.

Elizabeth Gibney

13 October 2016



Mary Evans Picture Library

DeepMind's AI uses external memory to accomplish tasks that require reasoning, such as learning to navigate the London Underground.

Artificial-intelligence (AI) systems known as neural networks can [recognize images](#), [translate languages](#) and even [master the ancient game of Go](#). But their limited ability to represent complex relationships between data or variables has prevented them from conquering tasks that require logic and reasoning.

In a paper published in *Nature* on 12 October<sup>1</sup>, the Google-owned company DeepMind in London reveals that it has taken a step towards overcoming this hurdle by creating a neural network with an external memory. The combination allows the neural network not only to learn, but to use memory to store and recall facts to make inferences like a conventional algorithm. This in turn enables it to tackle problems such as navigating the London Underground without any prior knowledge and solving logic puzzles. Though solving these problems would not be impressive for an algorithm programmed to do so, the hybrid system manages to accomplish this without any predefined rules.

Although the approach is not entirely new — DeepMind itself reported attempting a similar feat in a preprint in 2014<sup>2</sup> — “the progress made in this paper is remarkable”, says Yoshua Bengio, a computer scientist at the University of Montreal in Canada.

## Memory magic

A neural network learns by strengthening connections between virtual neuron-like units. Without a memory, such a network might need to see a specific London Underground map thousands of times to learn the best way to navigate the tube.

DeepMind's new system — which they call a 'differentiable neural computer' — can make sense of a map it has never seen before. It first trains its neural network on randomly generated map-like structures (which could represent stations connected by lines, or other relationships), in the process learning how to store descriptions of these relationships in its external memory as well as answer questions about them. Confronted with a new map, the DeepMind system can write these new relationships — connections between Underground stations, in one example from the paper — to memory, and recall it to plan a route.

DeepMind's AI system used the same technique to tackle puzzles that require reasoning. After training on 20 different types of question-and-answer problems, it learnt to make accurate deductions. For example, the system deduced correctly that a ball is in a playground, having been informed that "John picked up the football" and "John is in the playground". It got such problems right more than 96% of the time. The system performed better than 'recurrent neural networks', which also have a memory, but one that is in the fabric of the network itself, and so is less flexible than an external memory.

Although the DeepMind technique has proven itself on only artificial problems, it could be applied to real-world tasks that involve making inferences from huge amounts of data. This could solve questions whose answers are not explicitly stated in the data set, says Alex Graves, a computer scientist at DeepMind and a co-author on the paper. For example, to determine whether two people lived in the same country at the same time, the system might collate facts from their respective Wikipedia pages.

Although the puzzles tackled by DeepMind's AI are simple, Bengio sees the paper as a signal that neural networks are advancing beyond mere pattern recognition to human-like tasks such as reasoning. "This extension is very important if we want to approach human-level AI."

*Nature* | doi:10.1038/nature.2016.20784

## References

---

1. Graves, A. *et al.* *Nature* <http://dx.doi.org/10.1038/nature20101> (2016).
2. Graves, A., Wayne, G. & Danihelka, I. Preprint at <http://arxiv.org/abs/1410.5401> (2014).