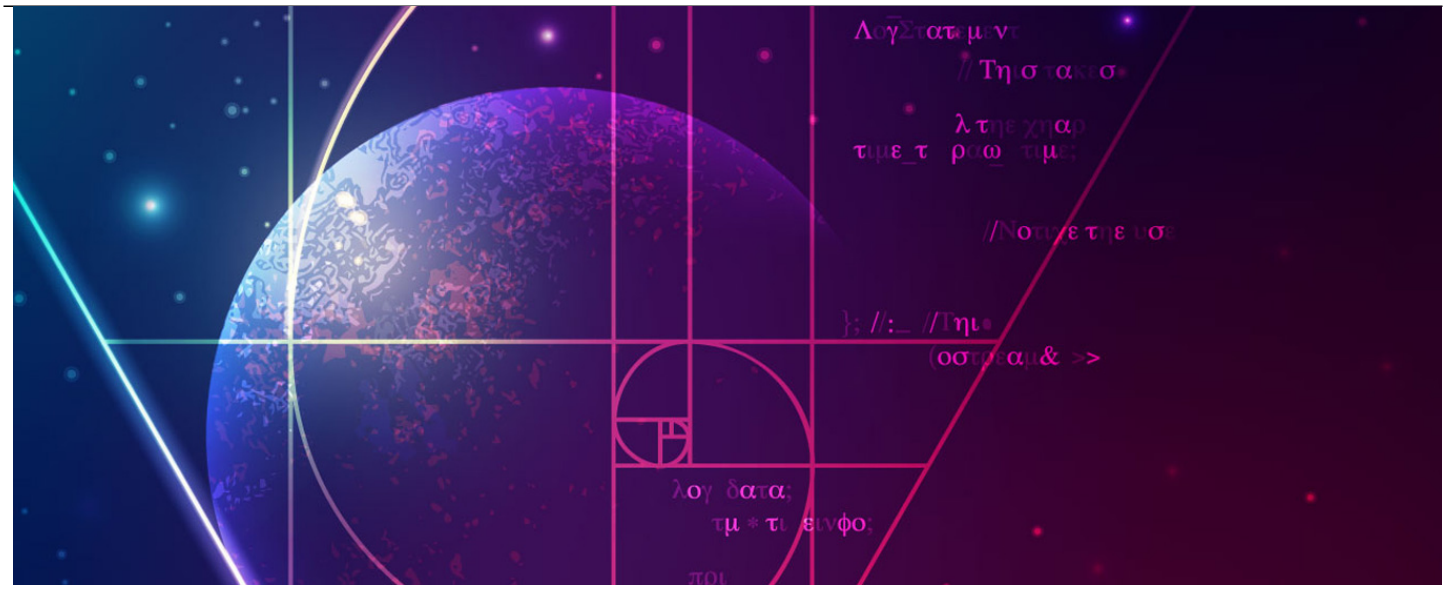
DeepMind ??????????????

DeepMind ??????????????

???Surya Prabha Vadlamani



Deep reinforcement Learning is an advanced type of machine learning in which an AI application solves very complex problems, usually to achieve a goal. On our blog, we recently discussed how deep reinforcement learning takes machine learning to another level of performance. With deep reinforcement learning, a machine can become more predictive. A machine learns from its mistakes, corrects them, and achieves complex goals such as winning a game or figuring out the fastest route for a self-driving car to travel amid constantly changing variables such as traffic patterns and weather conditions. Recently, DeepMind made a major breakthrough with deep reinforcement learning that has captured the attention of the business and technology world.

Let’s take a closer look.

DeepMind Unveils AlphaTensor

DeepMind, the artificial Intelligence subsidiary of Alphabet, unveiled AlphaTensor, the first artificial intelligence system for discovering novel, efficient and provably correct algorithms. AlphaTensor has discovered a faster way to do matrix multiplication, which a core problem in computing that affects thousands of everyday computer tasks.

Matrix multiplication involves multiplying numbers arranged in grids (or matrices) that might represent sets of pixels in images, air conditions in a weather model or the internal workings of an artificial neural network. To multiply two matrices together, a mathematician must multiply individual numbers and add them in specific ways to produce a new matrix. In 1969, mathematician Volker Strassen found a way to multiply a pair of 2 × 2 matrices using only seven multiplications, rather than eight.

Since Strassen published his algorithm in 1969, computer scientists have attempted to surpass its speed of multiplying two matrices. In fact, matrix multiplication is one of algebra's simplest operations (it's taught in high school math). But it is also one of the most fundamental computational tasks – and one of the core mathematical operations in today's neural networks.

DeepMind researchers tested AlphaTensor on input matrices up to 5 × 5. In many cases, AlphaTensor rediscovered shortcuts that had been devised by Strassen and other mathematicians, but in others it broke new ground. When multiplying a 4 × 5 matrix by a 5 × 5 matrix, for example, the previous best algorithm required 80 individual multiplications. AlphaTensor uncovered an algorithm that needed only 76.

AlphaTensor builds upon AlphaGo Zero, which, as we blogged, has demonstrated incredible superhuman powers beating even the most proficient experts on board games such as Go. AlphaTensor takes the AlphaZero journey further, moving from playing games to tackling unsolved mathematical problems.

Pushmeet Kohli, head of AI for science at DeepMind, said at a press conference, "If we're able to use AI to find new algorithms for fundamental computational tasks, this has enormous potential because we might be able to go beyond the algorithms that are currently used, which could lead to improved efficiency."

He added that to date, AI has not been very good at finding new algorithms for fundamental computational tasks. Automating algorithmic discovery using AI requires a long and difficult reasoning process — from forming intuition about the algorithmic problem to actually writing a novel algorithm and proving that the algorithm is correct on specific instances.

But AlphaTensor has achieved a breakthrough by discovering algorithms that are more efficient than the state of the art for many matrix sizes and outperform human-designed ones.

AlphaTensor begins without any knowledge about the problem, Kohli said, and then gradually learns what is happening and improves over time. "It first finds this classroom algorithm that we were taught, and then it finds historical algorithms such as Strassen's and then at some point, it surpasses them and discovers completely new algorithms that are faster than previously."

The implications are far-reaching. Matrix multiplication is used for processing smartphone images, understanding speech commands, generating computer graphics for computer games, data compression, and much more. Companies today use expensive GPU hardware to boost matrix multiplication efficiency; so, if a task can be completed slightly more efficiently, then it can be run on less powerful, less power-intensive hardware, or on the same hardware in less time, using less energy.

Even a small improvement in the efficiency of these algorithms could bring large performance gains, or significant energy savings. DeepMind found that the algorithms could boost computation speed by between 10 percent and 20 percent on certain hardware such as an Nvidia V100 graphics processing unit (GPU) and a Google tensor processing unit (TPU) v2.

For more detail, read the article that DeepMind published October 5 in *Nature*.

Contact Centific

To apply artificial intelligence, including deep reinforcement learning, to solve business problems, contact Centific. Learn about our work here.

-
-
-
-
-