Google DeepMind acquisition researchers working on a Neural Turing Machine

October 30 2014, by Bob Yirka

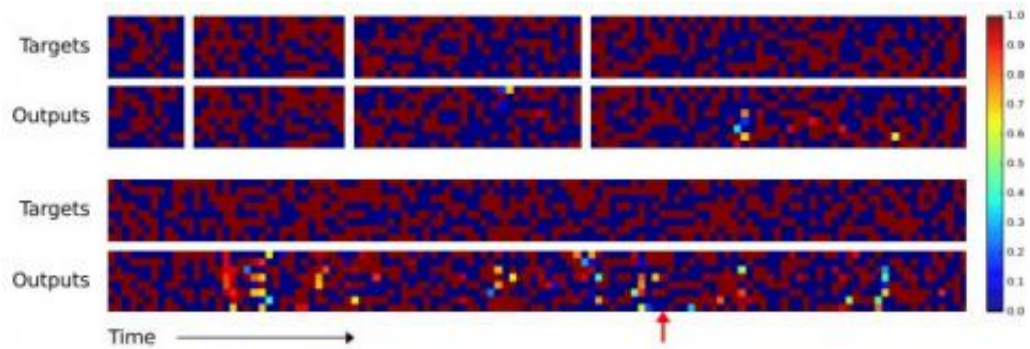

NTM Generalisation on the Copy Task. The four pairs of plots in the top row depict network outputs and corresponding copy targets for test sequences of length 10, 20, 30, and 50, respectively. The plots in the bottom row are for a length 120 sequence. The network was only trained on sequences of up to length 20. The first four sequences are reproduced with high confidence and very few mistakes. The longest one has a few more local errors and one global error: at the point indicated by the red arrow at the bottom, a single vector is duplicated, pushing all subsequent vectors one step back. Despite being subjectively close to a correct copy, this leads to a high loss. Credit: Neural Turing Machines, arXiv:1410.5401 [cs.NE]

Officials with Google have revealed that researchers working on a start-up recently purchased by the tech giant are working on building what they call a Neural Turing Machine—an artificial intelligence based computer system that seeks to fulfill the idea of a Turing Machine. Teams with the project (called DeepMind) have thus far uploaded two

papers to the *arXiv* preprint server—one describing the idea of their new machine, the other explaining related findings in Recurrent Neural Networks and Long-Short Term Memory Units.

A Turing machine (named for famed computer pioneer and deep thinker Alan Turning who came up with the idea back in 1936) as defined by Google is "a mathematical model of a hypothetical computing machine that can use a predefined set of rules to determine a result from a set of input variables." In other words a model of a computer that can learn the way we humans do. Over the past couple of decades, computer scientists have come closer to building such a machine using the idea of a neural network—interconnected nodes (neurons) which together represent data, and which can be reassembled to support changes (learning) to the network. But such machines to date have been missing one vital piece—external memory. Not in the traditional sense, of course, but in the sense that external memory can be used to store ideas or concepts that result from reconfiguration of neurons (learning).

One example would be where a collection of some nodes in a network together represent the idea of the game of basketball—the rules, the history, records made by noted players, etc., everything that it entails. External memory would mean storing the concept of a single word—basketball, the way it happens for us humans—when we hear the word we imagine players we rooted for, big games, or perhaps baskets we made as kids, and on and on. In this new effort, the researchers at DeepMind are trying to add that piece to a Neural Network to create a true real-world representation of a Turing Machine.

The team reports that they are making progress—they have all the pieces—a neural network, input/output and of course that external memory piece. They also report that the machine works when applied in very simple ways, and impressively, is able to outperform regular neural networks in several instances. That's the good news. The bad news, as the

team acknowledges, is that the team still has a very long way to go.

More information: Neural Turing Machines, arXiv:1410.5401 [cs.NE] arxiv.org/abs/1410.5401

Learning to Execute, arXiv:1410.4615 [cs.NE] arxiv.org/abs/1410.4615