Advancing Neural Turing Machines: Learning a Solution to the Shortest Path Problem

Author: Conor Sexton
Supervisor: Dr. Rozenn Dahyot
MSc in Computer Science

May, 2017

Abstract

The growth of the internet and advances in processing power have seen the dawn of a machine learning age. Concepts and models, new and old, are being combined and reworked to produce interesting and useful results. As we struggle to make sense of the vast amount of data that we now store, there is a pressing need to develop models that can make sense of it for us.

One such model is the neural Turing machine. Combining a neural network with an addressable memory bank, it emulates a human's ability to maintain a working memory. Initial research has shown that by using this memory, it can learn its own algorithms to solve problems, based solely on input and output examples.

This work measures the neural Turing machine's ability to learn a solution to a problem more complex than those it has seen before - the shortest path problem from graph theory. A large number of problems can be expressed using graphs, so showing that a neural Turing machine can reason over their structure shows that they can reason over a much larger set of problems.

By analysing the performance of the neural Turing machine against a long short-term memory network, this work shows that the memory has a positive effect on the model's ability to both solve the problem, and understand graph structures.