Ricky's Afterthought:

Artificial Intelligence "Par Excellence"

A. C. (Ricky) Metaxas

Life Fellow, St. John's College, Cambridge, UK
E-mail: acm33@cam.ac.uk

A few months ago I attended a lecture given by Demis Hassabis the co-founder and CEO of DeepMind, a neuroscience inspired Artificial Intelligence (AI) company bought by Google in 2014 and became subsequently Vice President of Engineering at Google. Young Hassabis studied computer science at Cambridge University, founded a number of companies, carried out a PhD in cognitive neuroscience at University College London, followed by post-doctoral positions at MIT and Harvard before he founded DeepMind.

The mission of DeepMind is to "solve intelligence" in a very generic way. Once intelligence is solved then, Hassabis argues, one can use it to solve literally anything. In DeepMind they try to create a general purpose learning machine, with the two key words being "learning" and "general". The algorithms they work with at DeepMind learn automatically from raw input data or raw experience. In other words, the system learns from itself. To distinguish it from normal AI, they called it Artificial General Intelligence (AGI). The key to building this system is "reinforcement learning framework", involving a loop made of an agent and its environment, where the two interact in just two ways: an agent gets fed back observations from an environment, and then based on these observations actions are directed back to the environment for more learning (and this loop continues thousands of times). The observations are coming in real time, they are incomplete and noisy and the agent attempts to build the best statistical model it can of the environment out there. Once the model is designed the second job of the system is to select the best action from a set of actions available to it at that moment which may involve, past experience, imagination or planning into the future.

Although early in the development of the DeepMind software many games, including chess, were challenged with great success, the developers of DeepMind were seeking a more complex game to tackle next. The one chosen was that of the ancient game of GO, a board game, involving two contestants placing black and white stones on a 90x90 grid. It is 3000 years old, 40 million people play the game and there are 10^{170} possible board positions. The aim of the game is to surround the opponent's stones, and the game ends when one player has amassed more squares than the other (in effect, the player with most territory wins). Apparently this game is very popular in Far East and if any youngster shows aptitude towards it, he or she are removed from the family and attend a special school to perfect the art of playing GO, and are closely supervised by a teacher who is a GO expert or past Master.

In addition, two neural networks were designed to train the system using data from past GO games. One, called the Policy neural network, mimics the moves of hundreds of thousands of average GO individuals playing the game, which meant that once it was fully trained the policy neural network would, from real experience, give the best 5 or 10 positions that should have been played at that particular stage of the game rather than the possible 200 moves. The DeepMind developers then used that neural network and played it thousands of times using the reinforcement learning framework to make sure that when a game was at last won the system stored all the "winning positions". The upshot of this

procedure was that the system learned at any particular instant in a future game which position to move next if it encountered a similar situation. Conversely, if the game was lost then the system will eradicate the moves that led to that particular game so that it got rid of the historical mistakes carried out in the past. A database was created using millions of games, and then they used the data from these games to train a second neural network, called the Valuer, which predicted with some accuracy which player was going to win. The combination of these two neural networks enabled one to make this seemingly intractable game more manageable.

The essence of the lecture was the development of DeepMind's software called AlphaGo which attempts to beat the top players in the world in the game of GO. First, the European champion was easily beaten 5-0 which gave the developers confidence that this system was ready to tackle the ultimate test by playing in March 2016 against South Korean world champion See Sedol in a $1 million challenge consisting of five games of GO. See Sedol was the acknowledged guru of the game, that is, the Roger Federer of GO. Before the challenge, the prediction was that See Sedol would win 5-0, however, to the consternation of all the pundits and local experts AlphaGo beat him 4/1. Over 500 million people watched that game live on their screens, and the whole of South Korea came to a standstill during that week. Presumably the majority of the people in the west were unaware that this was taking place.

Finally, Hassabis stressed that although this was a great success, the end aim of their company is not to just perfect the art of playing a game like GO, but they use GO as a training tool for tackling real problems, and currently the developers are addressing DeepMind towards Healthcare, the Energy Sector with emphasis on Renewables, and other scientific areas. One recent example Hassabis gave was energy usage at Google where he casually mentioned to their engineers that AlphaGO may help improve the overall energy usage in their data bases. They were not convinced stressing that their system was already highly optimised but in the end they acquiesced. The result was AlphaGO saved 40% of Google's energy usage amounting to millions of $!

By the way, as I write this piece I learned that in May 2017, an improved version of AlphaGO beat 19-year-old Chinese GO Master Ke Jie proving that computers are very close to mimicking how the brain functions. After the game, Ke Jie stated that last year AlphaGo played as a human brain would do, however, this year it became like God of Go. I am not advocating Microwave and RF heating researchers start playing board games, but with a bit of lateral thinking who knows where it can finally lead!

The lecture by Hassabis can be viewed in full at the following link:
https://www.youtube.com/watch?v=ZyUFy29z3Cw&feature=youtu.be