How an evolving technology might be used to support or develop new business – mac...

What is machine learning?

Machine Learning is a way of teaching computers to teach themselves. Despite the concept being almost as old as computers themselves, it has only recently become a feasible method for business due to huge increases in computational power, portable computers such as mobiles and higher network speeds.

How does Machine Learning help us?

The main principle of Machine Learning is for a programmer to create a system where the computer can use different methods to find the best solution to a problem. An example used in the current era is for the rapidly developing world of driverless cars. It would be an almost impossible task to create a program that could drive perfectly, as there are so many rules, and more exceptions to the rules that would have to be programmed in. Every detail about the breaking of a car, the distances of obstructions, what to do when other cars drive badly, would all have to be tested and then made into a set of rules for an automated car to follow. However, cars can now have extremely powerful onboard computers capable of reading the world around them and then deciding “ on their own” how to deal with what they see.

In a video showing how an evolving neural network can teach a car to drive with very simple conditions, we can see how effective this tool could be when used for much larger problems. The video shows how the car, with 3 inputs simply saying if there is anything to the left, in front, or to the right of the car, is able to navigate a track of tight corners simply by ‘ evolving’ its rules. To program the vehicle in this scenario wouldn’t have been all that difficult due to the simplicity of the problem, but Neural Networks can take

huge problems and find the best solution, and then evolve themselves to fit scenarios that a programmer might not have anticipated, making them very versatile programs that can expand to fit with the needs they face.

How can a neural network evolve?

Using the example from the video, the car has 3 inputs, and can only output a speed and turning angle for the vehicle. To begin with, a random set of rules are created that take the 3 inputs, manipulate them and then create outputs. As expected, the initial results are erratic and extremely bad – but as multiple random rules are used, the ones that work the best are taken to a second round and the ones that travelled the least distance are thrown out. Using a mutation factor, the winning vehicle rules are once again randomly altered, usually by a small amount such as 2-5%. Over several hundred rounds, these small changes can slowly push the rules in the direction required for a flawless track completion.

The video shows that by the 33rd generation, the car is able to complete the track. However, some imperfections are noted, such as the vehicle only breaks when turning left and being close to the right wall. This is one of the downsides with machine learning as a rule set can work well but it can have strange features and may not be the most effective way of completing the problem.

To understand this issue with Machine Learning, we must investigate the natural world where the very concept of machine learning stems from. Evolutionary Biologist and Author, Richard Dawkins, refers to this as getting stuck on a peak on “ mount improbable”. His theory states that an

evolutionary path of a creature can find a working solution, but by the time it climbs up the ‘ peak’, it is too far down for a mutation to find a better path which would lead it higher up the mountain. The diagram attached here shows how the evolution of certain creature’s eyes, such as the Nautilus which has evolved such that it cannot improve as it lacks a lens in its eye. Other creatures which did not climb this path, instead found other routes which took them higher, but as creatures do not devolve, it is now stuck.

This problem can also be found with evolving networks as they use the same principal of mutation, testing and the survival of the fittest. The only real way to avoid a network getting stuck in a local optima is to simply run the network again, from the start, with completely new starting values. Once this has been done enough times to suit the complexity of the desired solution, and no better solutions have been found, it can be assumed that an optimum solution has been found.

Concerns for a business using Machine Learning

For a business planning on using this technology, they must understand that the first solution they get out may not be the best one. Heavily complex tasks where this technology might be employed , such as automated cars, stock predictions and market advertising based on consumer trends will have to be tested multiple times to find the safest, most accurate and least likely to fail solution. If this part of the process is ignored, slight anomalies could exist creating radically incorrect results.

After all, when a neural network is fed data that it hasn’t solved before, or isn’t expecting as it’s far from what it would usually react to, the results

could be fatal. If the company Tesla, the market leader of electric cars and the biggest name in the automated car industry, were to release a fully automated car and hadn't done thousands of hours of testing in conditions that push the possibilities to the extreme, they could lose millions in lawsuits and potentially drop their value in the billions. That is why currently some of their cars have the ability to drive themselves, but it is required to have a driver alert and ready to take the wheel for if something does go wrong – even if the cars can drive themselves without an issue for 99. 9% of journeys.

Why business needs to adopt before it's too late

Internet Giant, Google, has been known to adopt technology early on to get ahead of everybody else. So when they recently did a live demonstration of Google Duplex , potentially the most advanced virtual assistant ever created using an extremely advanced neural network, other companies should be jumping on the technology to improve their business. Google Duplex should scare competitors such as Siri on the Apple iPhone, as the number of things it can do vastly exceed it. One of the flagship parts of Duplex is the virtual assistant being able to book appointments with a voice so human that people didn't realise it was a computer. It can understand the nuances of human speech and respond to questions so well that it feels like you are talking to a person. Right now it may not be able to hold an in depth conversation, but a simple minute long conversation to book an appointment goes unnoticed by the other conversation participant. Without machine learning, I believe this software would have been impossible to make. Just as Google has overtaken other search engines due to the enormous amount of

data they can collect to show better results, they will soon be overtaking many other companies using neural network technology.

The fascinating part of this rapidly growing field is that it can be applied to a vast amount of problems to improve them and remove the human error part of that job. Just as machines took the jobs of loom workers, and robotics took the jobs of factory laborers, neural networks will quite possibly take the most jobs of any technology up to this point. Train systems such as the London Underground may well be entirely automated in the few years because of it, pilots may no longer be required, the easier jobs done by GPs could be removed from their work as highly advanced machines are able to scan patients and identify issues faster than any human can.

In short, this technology is only just beginning. The benefits we could have from it are so great that it is hard to theorise how much they could do if given enough time, resources and guided training to do things we don’t need to do anymore. Eventually machine learning will let companies make more money, have fewer human errors, be more helpful to consumers/clients and require less human input while doing it, all while improving people’s lives.