

Catching fires in a neural net

Virginia Tressider

What do you do in your spare time? Bit of gardening? Swimming? Cooking? How about training a computer neural network to predict where bushfires might break out?



Credit: [Firstsignal/istock](#)

That's what Ritaban Dutta, from CSIRO's Computational Informatics Division, chose for an after-hours and weekend project. With the [increased risk of bushfires in Australia in mind](#), he got together with three colleagues from the University of Tasmania – Jagannath Aryal, Aruneema Das and Jamie Kirkpatrick – and started developing a tool that could predict where the risks were highest.

It didn't stay a part-time hobby for long. The group's [results](#) were published in *Nature Scientific Reports* in November 2013. That in itself is an indicator that this no ordinary spare-time project: fewer than 1 in 20 papers submitted to *Nature* are accepted.

The predictive tool the three researchers developed is an artificial neural network that transforms climate data into a map, and uses the results to learn and predict, in this case, bushfire hotspots.

So what exactly are neural networks? And how does this one work?


Biomimicry

Surprisingly, artificial neural networks – otherwise known as artificial intelligence – have been around for some 30 years (with many failures and setbacks, but without harming the world at large, despite the bad press in films and fiction).

Neural networks began as computational models that attempted to mimic animals' central nervous systems – in particular the brain. They worked like a series of interconnected neurons that, in the same way as the brain, could

convert inputs into values from the algorithms fed into the network. They were capable of [machine learning](#) and pattern recognition.

In modern neural networks, the biomimicry approach has largely been replaced by an emphasis on statistics and signal processing. However, both early and contemporary neural networks are based on the principles of non-linear, distributed, parallel and local processing and adaptation.



Credit: NASA

Neural networks fall into two categories: supervised and unsupervised networks.

Supervised networks are a type of ‘black box’ into which you feed information, and the system learns against known inputs. Basically, they’re a super-fast pattern finding and matching tool that ‘learns’ from examples, in the same way children learn to recognise dogs from examples of dogs, and exhibit the capability to generalise.

Unsupervised networks also learn, but are also capable of ‘deep learning’ through a two-stage process. First, the input is fed in, then the network – unsupervised – develops a ‘features matrix’. In other words, it develops a theory. [Google](#) is a leading proponent of deep learning, presumably as a means to leverage the vast streams of data it collects.

The CSIRO’s fire-hotspot predictive tool belongs to a class known as ‘recurrent’ neural networks, which can use their internal memory to process arbitrary sequences of inputs based on an architecture known as an [Elman network](#).

This three-layer network includes a set of ‘context units’ – feedback loops. These context units enable the network to train itself a stepwise fashion, ‘learning’ with each new ‘experience’. Because the network is trained one layer at a time, it can perform the high-level abstractions needed to plot a complex set of interactions – such as hand-writing recognition or, in this case, fire prediction.

Beyond bushfires

The next step for the researchers was to feed historic climate data collected from around Australia from 2001 to 2010, including soil moisture, wind speed, and humidity. They also used NASA’s satellite images of wildfires. The result was a tool that could pick bushfire hotspots with a 90 per cent level of accuracy and few ‘false negatives’.



Credit: NASA

Predicting where bushfires are likely to occur isn't the only environmental problem to which such networks can be applied.

Dr Dutta is also working on using neural networks in irrigation management – again using NASA satellite images, in this case, vegetation maps. The idea is to develop high-resolution, accurate pictures of the water balance of a whole farm, which can vary by location and time.

Elsewhere, researchers at Griffith University are using an artificial neural network to develop a method to rapidly estimate projected flood water levels. This would give small coastal communities more warning time of the need to evacuate.

The deep learning capacities of neural networks are necessary for getting real value from massive datasets such as those from NASA satellites. Conventional techniques of data processing are simply not up to the task. Thanks to artificial intelligence – and the human curiosity of three scientists – Australia has become more 'bushfire ready' at a time when that capacity is more critical than ever.

From ECOS online <http://www.ecosmagazine.com/?paper=EC14098>