Artificial Neural Networks as a Model of the Human Brain

Have you ever wondered why there are tasks that are extremely simple for any human but incredibly difficult for computers? Artificial neural networks, or ANNs, were inspired by the central nervous system of humans. Like their biological counterpart, ANN's are built upon simple signal processing elements that are connected together into a large mesh. ANN's have been successfully applied to a number of problem domains:

- Classify data by recognizing patterns. Is this a tree on that picture?
- Approximate a target function–useful for predictions and forecasting. Will this storm turn into a tornado?

If this sounds a bit abstract, so let's look at some real-world applications. Neural networks can:

- identify faces,
- recognize speech,
- read your handwriting (mine perhaps not),
- translate texts,
- play games (typically board games or card games)
- control autonomous vehicles, robots and surely more things!

There are many ways of knitting the nodes of a neural network together, and each way results in a more or less complex behavior. Possibly the simplest of all topologies is the feed-forward network. Signals flow in one direction only; there is never any loop in the signal paths (figure 1). These different topologies, or node maps, are also known as the **neural network zoo** (figure 2). The topology is chosen according to the use-case of the ANN that we want to train.

Typically, ANN's have a layered structure. The input layer picks up the input signals and passes them on to the next layer, the so-called 'hidden' layer. (Actually, there may be more than one hidden layer in a neural network). Last, comes the output layer that delivers the result.

Neural networks must learn. Unlike traditional algorithms, neural networks cannot be 'programmed' or 'configured' to work in the intended way. Just like human brains, they have to learn how to accomplish a task. Roughly speaking, machine learning can be supervised and unsupervised.

Supervised learning can be used if a (large enough) set of test data with known results exists. Then the learning goes like this: Process one dataset. Compare the output against the known result. Adjust the network and repeat; this step is fundamental to the whole process of learning. Now we know a bit about the nature of artificial neural networks, but what exactly are they made of? And most importantly, how does a machine learn?

AppliedGo.Net
Asimovinstitute.org
C. Christoph, 2016
[Edited]

Figure 1.The topology of an artificial neural network

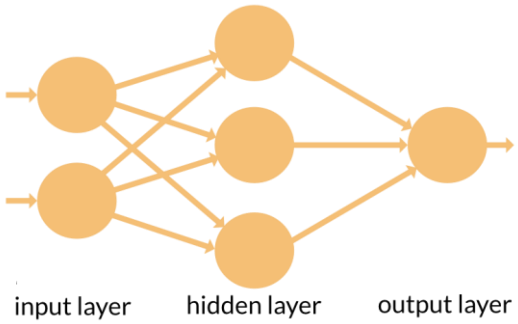

Figure 2.The ANN Zoo