УДК 811.111: 004.896

Karulin V., Basovets I.

Neural Networks

Belarusian National Technical University
Minsk, Belarus

We all know that from the beginning of the 20th century, science fiction writers dreamed of the intelligence of robots, which would be identical to the human. In the 21st century, we almost reached the point of creating a full-fledged artificial intelligence and the key to this is neural networks creation.

Artificial neural networks are computer programs assembled from hundreds, thousands, or millions of artificial brain cells (called units) that learn and behave in a remarkably similar way to human brains. Each unit has its own function.

A typical neural network has anything from a few dozen to hundreds, thousands, or even millions of artificial neurons called units arranged in a series of layers, each of which connects to the layers on either side. Some of them, known as input units, are designed to receive various forms of information from the outside world that the network will attempt to learn about, recognize, or otherwise process. Other units sit on the opposite side of the network and signal how it responds to the information it’s learned; those are known as output units. In between the input units and output units are one or more layers of hidden units, which, together, form the majority of the artificial brain. Most neural networks are fully connected, which means each hidden unit and each output unit is connected to every unit in the layers either side. The connections between one unit and another are represented by a number called a weight, which can be either positive (if one unit excites another) or negative (if one unit suppresses or inhibits another). The higher the weight, the more influence one unit has on another. (This corresponds to the way actual brain cells trigger one another across tiny gaps called synapses.)

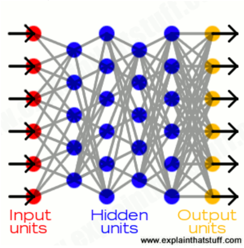

A fully connected neural network is made up of input units (red), hidden units (blue), and output units (yellow), with all the units connected to all the units in the layers either side. Inputs are fed in from the left, activate the hidden units in the middle, and make outputs feed out from the right. The strength (weight) of the connection between any two units is gradually adjusted as the network learns [1].

Information flows through a neural network in two ways. When it's learning (being trained) or operating normally (after being trained), patterns of information are fed into the network via the input units, which trigger the layers of hidden units, and these in turn arrive at the output units. This common design is called a feedforward network. Not all units "fire" all the time. Each unit receives inputs from the units to its left, and the inputs are multiplied by the weights of the connections they travel along. Every unit adds up all the inputs it receives in this way and (in the simplest type of network) if the sum is more than a certain threshold value, the unit "fires" and triggers the units it's connected to.

For a neural network to learn, there has to be an element of feedback involved – just as children learn by being told what they're doing right or wrong. In fact, we all use feedback, all the time. Think back to when you first learned to play a game like ten-pin bowling. As you picked up the heavy ball and rolled it down the alley, your brain watched how quickly the ball moved and the line it followed, and noted how close you came to knocking down the skittles. Next time it was your turn, you remembered what you'd done wrong before, modified your movements accordingly, and hopefully threw the ball a bit better. So you used feedback to compare the outcome you wanted with what actually happened, figured out the difference between the two, and used that to change what you did next time. The bigger the difference between the intended and actual outcome, the more radically you would have altered your moves.

You learn how to do skillful things like this with the help of the neural network inside your brain. Every time you throw the ball wrong, you learn what corrections you need to make next time. Neural networks learn things in exactly the same way, typically by a feedback process called backpropagation. This involves comparing the output a network produces with the output it was meant to produce, and using the difference between them to modify the weights of the connections between the units in the network, working from the output units through the hidden units to the input units – going backward, in other words. In time, backpropagation causes the network to learn, reducing the difference between actual and intended output to the point where the two exactly coincide, so the network figures things out exactly as it should [1].

Modern examples of neural networks are robots in Boston dynamics able to walk. Or a nice picture for your profile in social media can be the product of neural-style which combines content and style from two different images. Another famous and very popular among students of our university example of neural networks is Google Translate which is getting brainier. The online translation tool recently started using a neural network to translate between some of its most popular languages – and the system is now so clever that it can do this for language pairs on which it has not been explicitly trained. To do this, it seems to have created its own artificial language. Traditional machine-translation systems break sentences into words and phrases, and translate each individually. In September 2017, Google Translate unveiled a new system that uses a neural network to work on entire sentences at once, giving it more context to figure out the best translation. This system is now in action for eight of the most common language pairs.

Although neural machine-translation systems are fast becoming popular, most only work on a single pair of languages, so different systems are needed to translate between others. With a little tinkering, however, Google has extended its system so that it can handle multiple pairs – and it can translate between two languages when it hasn't been directly trained to do so. For example, if the neural network has been taught to translate between English and Japanese, and English and Korean, it can also translate between Japanese and Korean without first going through English. This capability may enable Google to quickly scale the system to translate between a large number of languages [2]. This is the best proof that neural networks can evolve and develop much like a person does. Time will show how far this development can reach.

References:

1. Neural Networks [Electronic resource]. – Mode of access: http:// www.explainthatstuff.com/introduction -to-neural-networks. html. – Date of access: 05.03.2017.

2. Google Translate AI invents its Own Language to Translate with [Electronic resource]. – Mode of access: https://www.newscientist.com/article/2114748-google-translate -ai-invents-its-own-language-to-translate-with/. – Date of access: 04.03.2017.