

Neural Networks¹

Neural networks have proved very effective for learning to classify images, speech, text and other patterns, but also to transduce such patterns into other patterns, for example mapping one language to another, mapping images to captions, generating possible continuations of given partial stories, etc. However, they require massive training data, whereas people learn from far fewer examples.

1 Evolution of Neural Net Technology towards “Deep Learning”

What are neural nets?

(Artificial) Neural Networks (ANNs, or just NNs), as understood in AI, are networks of computational units that resemble networks of neurons in the brain. I.e., the units compute a weighted sum of their inputs, and apply a “threshold function” (also called an “activation function”) to produce a single output. This output is 0 or low if the weighted sum of inputs is below threshold, while it is high (typically close to 1) if the weighted sum of inputs is above threshold. The weights applied to the inputs received by a unit can be altered in the course of learning. They can be thought of either as (modifiable) properties of the unit receiving the inputs, or of the unit-to-unit connections.

Each neuron-like unit has a number of input connections to it and a number of output connections leading away from it. Typically, units are arranged into layers, and function either as *input units*, *hidden units*, or *output units*. Input units have a single input or multiple inputs (which could be discrete or continuous) supplied from the outside, where each input represents an element of the pattern to be processed; input units (typically) transmit their output to multiple hidden units in the next layer via connections to those units. Hidden units receive inputs from the previous layer, and transmit their output through multiple connections to the next layer. Output units receive their inputs from the previous layer, and their output represents an element of the NN’s final output. For simple binary classification tasks, there might be just one output unit, supplying a choice between the two classes.

A note about counting layers: There can be some ambiguity about how many layers a NN has. Suppose that each unit of the initial “layer” sees the value of just one element of the input pattern and neither applies a weight to it nor does any thresholding; i.e., it

¹Caveat: This an overview by a non-expert on this topic; For an up-to-date reference see Jurafsky & Martin’s *Speech and Language Processing*, ch. 7, <https://web.stanford.edu/~jurafsky/slp3/7.pdf>

simply transmits the given value to some number of units in the next layer. But then that next layer is not hidden – its inputs are known; so we may well discount the layer of transmission points as a layer of NN units. If, on the other hand, the input units do weight and threshold their inputs, then even if each input unit just sees one element of the input pattern and transmits its output to just one successor unit, it would be counted as an input layer. In general, we can properly count any layer of units that perform computations (other than identity) on their inputs as NN layers, even if the weights applied to inputs are not adjustable, or no thresholding is done.

Some types of NNs

Networks in which signal propagation is unidirectional from layer to layer (in effect, an acyclic directed graph) are called *feed-forward* NNs. Networks allowing cyclic signal pathways are called *recurrent* networks. Important kinds of recurrent networks include Hopfield nets and Boltzman machines, in which connections are bidirectional (and thus certainly “loopy”). In Boltzman machines activation of units is probabilistic, rather than rigidly determined by the inputs. Such networks can model associative access, in the sense of approximately reproducing a complete learned pattern based on inputting parts of the pattern. In effect they sparsely encode the learned patterns in the hidden layer(s). In natural language processing, recurrent neural nets with loops serving as memory, particularly LSTMs (Long Short-Term Memory neural networks) play a key role.

The simplest kind of NN, a *perceptron*, has just one unit, accepting some set of inputs and thresholding their weighted sum. Perceptrons were shown to be able to learn to classify inputs (e.g., handwritten characters) by training them on labeled examples. Training involves use of classification successes and errors to modify the weights of the connections from the inputs to the thresholding function, “rewarding” those connections that contributed to correct outputs and “punishing” those that contributed to incorrect outputs.

Perceptrons can learn to compute classifications only where the classes of inputs are *linearly separable*. For example, in two dimensions (say, 2 input units with continuous inputs), this implies straight-line separability of clusters of points in the Cartesian plane corresponding to distinct classes; in three dimensions, it implies separability by a plane, and in higher dimensions it implies separability by hyperplanes. Examples of classification problems unsolvable by perceptrons are input parity (whether there is an even or odd number of binary inputs) and connectedness (whether or not the “black” pixels in a square binary array of black and white pixels are connected pixel-to-pixel or not). The discovery of these limitations by Minsky & Papert in the late 1960s put a damper on the initial hype that had accompanied the prior literature on perceptrons. Remarkably, however, it was shown in later work that NNs with just one hidden layer (thus, 3 layers) could in principle compute any boolean function, could approximate any bounded, continuous function, and implement any decision boundary; and with two hidden layers,

any function could be approximated. However, these results are not necessarily practical, because an exponential number of hidden units may be required for a given level of approximation. In practice, adding more layers (possibly dozens) has been found to improve NN performance. Also, note that function computation is intuitively inadequate for reasoning over a knowledge base, which may involve a number of steps with no fixed upper limit, and the knowledge base may be continually added to.

Types of learning methods in NNs

In the 1960s and 70s, an obstacle to effective use of multilayer NNs was a lack of good methods for training such networks using labeled examples: How can one assign “credit” or “blame” to connection strengths of inputs to hidden units? This problem was solved by using *differentiable* threshold functions and *backpropagation*. Essentially, this involves taking the derivative of the (squared) output error (for a given NN input pattern) with respect to each connection weight in the NN, and adjusting the weight in a positive or negative direction depending on whether that derivative is negative or positive respectively. In the original conception of perceptrons, threshold functions were discontinuous step functions, and as such not differentiable. It was the introduction of continuous *sigmoid* functions or *hyperbolic tangent* functions as threshold functions (see below) that made backpropagation possible. The mathematics of backpropagation was already developed in the 1960s, but it was not until the 1970s that the applicability to NN learning began to be appreciated, and it was only in 1986 that an experimental demonstration of the technique by Rumelhart, Hinton, and Williams firmly established the utility of the method.

Even with backpropagation, successes in training multilayer NNs were limited by the need for huge numbers of training examples, and by the tendency of iterative learning to settle into local minima, instead of finding weights that globally minimize errors. NN developers generally had to be quite clever in engineering good “features” of an input, for a given classification task, that would ease the learning problem for the NN. This detracted from the idea of NNs as a means of learning arbitrary classification tasks. A technique apparently due to Bourlard and Kamp (1988) was “auto-association” or “autoencoding”. A basic autoencoder has at least an input layer, called an encoder (think of it as having a large number of units, allowing for complex input patterns), a hidden layer with fewer units (think of it as representing the input in an efficient “code”), and an output layer, called the decoder (with the same number of units as the input layer), which is intended to reproduce the input as nearly as possible. To learn to reproduce its inputs (based on numerous input examples), the system is forced to learn an efficient encoding (one of “lower dimensionality”) of its inputs. The important point is that these efficient encodings are learned in a *unsupervised* manner, i.e., the inputs needn’t be labeled with desired outputs, because the desired outputs *are* the inputs! The encoding/decoding idea is key to many other tasks. For example, in language modeling a NN may learn to encode fixed-length segments of text in such a way as to optimize

its prediction of the following word or words. In machine translation, an NN learns to encode source language sentences in such a way as to optimize generation of the corresponding sentence in the target language. (Of course the latter is an example of supervised learning – we need to have the example translations from which to learn.)

In 2006, Hinton *et al.* showed how to make the autoencoding technique practical. In effect they trained layers of restricted Boltzman machines so that each layer (with fewer units than the previous layer) learns the inputs to the previous layer (thus the first hidden layer learns the external inputs, and its outputs serve as inputs to be learned by the next, sparser, layer, etc.). The final (output) layer consists of units fully connected to all units of the previous layer, and only the last one or two layers of the NN are trained in supervised fashion, i.e., the NN seeks to obtain the given, correct outputs for a set of training inputs. In this training phase, the weights of the previously trained “autoencoder” levels are held fixed. In effect this kind of network discovers for itself what features of an input to derive in order to be able to compactly encode the inputs; the successive levels learn more and more abstract features, and these generally serve well for miscellaneous learning tasks, as inputs to the last layer or two.

This work again caused much excitement, and the term “deep learning”, was increasingly used in the literature about such multilayer NNs. Ultimately, it turned out that the availability of web-scale training data made Hinton *et al.*’s autoencoder approach unnecessary for learning from “big data” (though not for modest-size datasets). For example, a multi-layer feed-forward network trained in supervised fashion on many millions of images labeled with objects occurring in those images could be used almost as-is for many other tasks. One just needed to retrain the last couple of layers on labeled examples for the task at hand (e.g., skeletal joint positions on people in an image – as recently implemented by a local grad student, Iftekar Tanveer). Interestingly, in an NN trained on “big data” for a sufficiently diverse task the earlier hidden layers automatically learn to encode the inputs efficiently in a more or less task-independent way, at successive greater levels of abstraction, just as if autoencoding had been used for training.

Convolutional neural networks

It should be added that the successes of the above deep-learning systems have also been aided by a particular way the earlier layers are organized: Each of several initial hidden layers consist of identical units with identical input weights, each taking inputs from a small (say, 9-element) local “patch” of the input (where patches slightly overlap). Computing the weighted sum of the elements of each input patch, performed uniformly over identical, slightly overlapping patches that “tile” (cover) the entire input pattern, amounts to mathematical *convolution*. Another way to think about it is as matching a template to each patch of the input, and measuring the quality of the match; the match results are then the outputs of the convolutional layer. The use of convolutional templates was inspired by the “receptive fields” that play a key role in mammalian (including human) vision. Often the convolutional layers are followed by “pooling”

layers, which again take their inputs from "patches" of the outputs of a convolutional layer, but in this case they use adjacent, nonoverlapping patches and each computes the maximum value in a patch. This ensures a degree of insensitivity to exactly where the convolution values (template matches) were highest in a local region.

The uniformity of the convolutional and pooling layers in convolutional NNs (CNNs) significantly eases the learning problem, since only one set of weights, used by each unit in a convolutional layer, needs to be learned for such layers, and pooling layers implement fixed max-functions. CNNs have been successfully applied to image and video processing, various NLP tasks (speech recognition, query-based document retrieval, syntactic and semantic parsing, machine translation, etc.), drug discovery, and the game of Go, in particular Google's AlphaGo (this defeated a 9-dan Go player 4 games to 1; the program also used special move prediction networks, board-value networks, and Monte Carlo game-tree search); the later AlphaGo Zero program was even stronger, playing well above human levels.

These recent successes (along with the "Watson" Jeopardy win and many news items about self-driving cars) have led to much hype about imminent human-level AI. However, NNs don't yet understand language, or learn to reason, plan, or engage in dialogue the way people do. They are still totally dependent on huge amounts of data for any particular task in order to learn that task, whereas people learn from relatively few examples. Currently the most advanced systems for reasoning, planning, or engaging in dialogue are still ones based on symbolic representations, rather than NNs.

2 Some Specifics

The behavior of individual units

As noted, standard NN units are thought of as "threshold units" because, in analogy with organic neurons, their output may be close to 0 until the net value of the inputs reaches a certain threshold, and will then rather abruptly become much higher. More precisely, suppose that unit i receives inputs from units $1, 2, \dots, n$; (normally unit i itself won't be one of these, unless it feeds its own output back to itself as an input). Suppose further that the outputs of those units are s_1, s_2, \dots, s_n respectively; (below we'll write them as $s_{i1}, s_{i2}, \dots, s_{in}$ to make clear that units $1, 2, \dots, n$ are the ones with connections to unit i ; the point is that each unit in general receives inputs from a distinct set of prior units). The connections from units $1, 2, \dots, n$ that transmit their outputs to unit i are assumed to have *weights* (connection strengths) $w_{i1}, w_{i2}, \dots, w_{in}$. These can be positive or negative, and we can think of a signal transmitted over a connection with a positive weight as "excitatory" (it tends to make the target unit "fire") and one transmitted over a connection with negative weight as "inhibitory" (it tends to prevent the target unit from firing). The net input net_i to unit i is taken to be the weighted sum of the signals received from units $1, 2, \dots, n$:

$$net_i = \sum_{j=1}^n w_{ij}s_{ij}.$$

This is the value that feeds into the threshold function (activation function) f of unit i , causing the unit to produce a 0 output or very low output when net_i is below threshold, and an output near 1 when net_i is above threshold:

$$s_i = f(net_i), \quad \text{where, e.g.,}$$

$$f(x) = \frac{1}{1 + e^{-x}} \quad (\text{a sigmoid function}).$$

You can see that for very negative x , $f(x)$ is close to 0, for $x = 0$ it is 0.5, and for $x > 0$, it rises to 1. Thus 0 is in effect the threshold value. There are other sigmoid functions, i.e., differentiable functions with this sort of shape. An advantage of the above version, called the *logistic* function, is that its derivative is expressible in terms of itself:

$$f'(x) = f(x)(1 - f(x)).$$

This makes the derivation of the backpropagation rules easier.

Backpropagation

First you should note that since we have a mathematical expression for the input-output behavior of each unit of an NN, we can also derive a mathematical expression for the output of each unit (including the NN output units) of the NN for any given input. Thus we can conceptualize the problem of finding the optimal weights w_{ij} , given a set of training examples, as a problem of minimizing the average output error over all the examples. This is a mathematical problem that can be tackled, for example, by gradient descent methods, since our functions are all differentiable.

However, from a learning perspective it is more convenient to use a method that learns example-by-example, and that is what backpropagation enables. We won't go through that derivation (which is based on the chain rule of differentiation, whose general form is $\delta f(g(\dots))/\delta w = f'(g(\dots))\delta g(\dots)/\delta w$, where w is some parameter whose influence on $f(g(\dots))$ we are trying to determine). Backpropagation makes weight adjustments in proportion to the derivatives of output error terms at each layer, working backward from the final layer. The algorithm is simple enough to be stated here for a feed-forward network:

Consider a particular training example;

1. Compute the output s_i for each unit i , in a forward sweep from the input layer to the output layer;
2. For each output unit i , compute its *error term*,
 $\delta_i \leftarrow s_i(1 - s_i)(t_i - s_i)$,
 where t_i is the correct (desired) output value at output unit i . Note the resemblance

of $s_i(1 - s_i)$ to the derivative of the logistic sigmoid function; and of course $(t_i - s_i)$ is the error at output unit i ;

3. For each layer whose successor layer has already been processed:

For each unit i in the layer, compute its error term,

$$\delta_i \leftarrow s_i(1 - s_i) \sum_{k \in \text{outputs}} (w_{ki} \delta_k);$$

note that w_{ki} is the weight of the connection *from* unit i to successor unit k – whose error term δ_k we have already computed;

4. Increment each network weight

$$w_{ij} \leftarrow w_{ij} + \eta \delta_i s_{ij},$$

where η is a constant called the *learning rate* (which should be small enough not to give excessive importance to a single input-output example); recall that s_{ij} is the output of unit j , but indicating that it is connected to unit i ; (we could have just written s_j).

We can't expect to get an optimal set of weights after a single pass through the training set; rather, we iterate until weights no longer change significantly. Often a *momentum term* is added to the increment in step 4, which is just a constant ≤ 1 times the increment added at the previous iteration. This can accelerate convergence and help avoid local minima.

As noted earlier, in applications to language modeling and other “sequence modeling” tasks, we generally use *recurrent* NNs (RNNs), such as LSTMs, which contain feedback loops that enable “remembering” and using features of input segments seen earlier in the processing. We'll discuss such RNNs separately.