

Chapter 13

Neural Networks

To give you an idea of how new this material is, let's do a little history lesson. The origins of neural nets are typically dated back to the early 1940's and work by two physiologists, McCulloch and Pitts. Hebb's work came later, when he formulated his rule for learning. In the 1950's came the *perceptron*. *The perceptron* is what we called a linear neural network- it became clear that the perceptron could only classify certain types of data (what we now call linearly separable data). This shortcoming led to a stop in neural net research for many years- It was not until the 1980's that neural net research really took off from a coming together of many disparate strands of research- There was a group in Finland led by T. Kohonen that was working with self-organizing maps (SOM); there was the work from the 70s with Stephen Grossberg, and finally several researchers were discovering methods for training new networks that could be put in parallel (back-propagation).

It was in the 80's and 90's that researchers were able to prove that a neural network was a **universal function approximator**- That is, for any function with certain nice properties, we are able to find a neural network that will approximate that function with arbitrarily small error. It is interesting to note that the use of a neural network could be used to solve Hilbert's 13th Problem¹:

“Every continuous function of two or more real variables can be written as the superposition of continuous functions of one real variable along with addition.”

Let's look at some highlights of more recent developments in neural nets:

- 1989: George Cybenko proves that neural nets with certain types of activation functions are universal function approximators.
- 1989: Yann LeCun used convolutional neural networks to read handwritten digits. This led to many rapid developments, and Dr. LeCun shared the Turin prize in 2018 for this work.
- 1989: Q-Learning is something that takes machine learning to a different branching- This greatly improves the area of reinforcement learning.
- 1993: Support Vector Machines (SVM), around for a long time but brought into machine learning, designed by C. Cortes and V. Vapnik.
- 1998: Yann LeCun brings in Stochastic Gradient Descent as a practical solution to large problems.
- 2006 or so: Approximately the beginning of “deep learning”, as coined by Geoffrey Hinton (a key player in machine learning for a very long time, also given the Turin award in 2018).
- 2009: The launch of ImageNet (Fei-Fei Li at Stanford). As of 2017, it contains 14 million **labeled** images that are available to researchers. The casual reader may not appreciate how key this is-

¹ In 1900, mathematician David Hilbert listed what he viewed as the 23 greatest unsolved problems. This list has been worked on ever since, and as of 2021, problems 8, 12, 13 and 16 remain unsolved.

“Deep learning” requires vast amounts of data, and this kind of data was not readily available to the community earlier.

- 2011: IBM’s Watson computer wins *Jeopardy!*.
- 2011: The use of the Rectified Linear Unit (or ReLU) as activation function. This function is used in deep networks to counter what is called “the vanishing gradient problem”.
- 2012: Creation of AlexNet - Its success kicked off a spike in researching convolutional neural nets (that continues to this day).
- 2014:

Generative Adversarial Networks (GAN). Basically, a second neural net is built as an “adversary”, where it builds data to try to fool the neural net that is learning a particular task. The basic idea has expanded, and these networks can now generate realistic images of human beings. The person to the right does not exist. (Photo from Wikipedia Commons)



Coming to present day, research is aimed at something called **deep neural nets**- in summary, the difference now is that we have a LOT of training data, and new algorithms and architectures that make it feasible to train complex associations. A nice summary of “The Decade of Deep Learning” is given at the site below (accessed Apr 2021): <https://bmk.sh/2019/12/31/The-Decade-of-Deep-Learning/>

Back to feed forward nets


The term “neural network” has come to be an umbrella term covering a broad range of different function approximation or pattern classification methods. The classic type of neural network can be defined simply as a connected graph with weighted edges and computational nodes- We’ve seen a linear neural network (using Widrow-Hoff training rule). We now turn to the workhorse of the neural network community: The feed forward neural network.

13.1 From Biology to Construction

We’ve seen the basic model before in the linear neural networks. We include some of it again for clarity. Information flows from the dendrites to the cell body through the axon to a synaptic junction connecting to the dendrite to the next neuron.

The synaptic junction is made up of the presynaptic node (the end of an axon), the postsynaptic node (the beginning of a dendrite), and the “empty space”, which is the synapse.

A diagram of a neuron is shown to the left.



As information flows across a synapse, information can be amplified, inhibited, or re-polarized. We’ll discuss our model of cell body processing and flow over the axon, but this is all the biology we use. It’s best to think of neural nets as being inspired by biology, although there are researchers interested specifically in working with biological neural nets.

Before we look at the neural net on the level of a layer of neurons, let’s look at how a single neuron (or “node”) processes information:

Let x_1, \dots, x_n denote the signals incoming along the dendrites. We model the processing done here by multiplying the signal by a real number which we call a **weight**.


At the cell body, these signals are first collected in what we'll call the **prestate** of the cell. They are added to a "bias term", b , which loosely represents the resting state of the cell.

$$p = \sum_{j=1}^n w_j x_j + b = \mathbf{w} \cdot \mathbf{x} + b$$

An activation function is applied, we'll just call it f for now. Applying f , we get what is called the **state** of the cell, or the cell's **activation**:


$$s = f(p)$$

And this value is then sent along the axon. Here is a typical way we might make a diagram of a single neuron, where the cell body has been divided into the sum stage, then the activation.



Putting lots of cells together is what a neural network is. Below on the left, we show a three layer network. Data is input at the first layer, the second layer is really the layer of cells, and the last layer is the output layer. In finding a mapping from \mathbb{R}^n to \mathbb{R}^m with k cells, we would construct a $n - k - m$ network.

Below and to the right, we see networks don't have to have a single layer; we can apply lots of layers to get a multilayer network. Training these can be difficult, and is the subject of our next chapter in deep learning. For the time being, we'll keep with a single hidden layer.



Next we'll look at how information flows through the network. It is easiest to describe using layers and linear algebra.

13.2 How Networks Compute

At the input layer, we're presented with a real value at each input node, which we'll denote as a vector $\mathbf{x} \in \mathbb{R}^n$. Next, we move from \mathbb{R}^n to \mathbb{R}^k . The affine mapping that takes us from the input layer to the hidden layer is defined as

$$P^{(2)} = W^{(1)}\mathbf{x} + \mathbf{b}^{(i)}$$

where $P^{(2)}$ is the vector in \mathbb{R}^k representing the prestate in each node (or cell). Since the weight matrix $W^{(1)}$ is $k \times n$, we make an important observation:

$W_{jk}^{(i)}$ is the weight connecting cell k from layer i to cell j on layer $i + 1$.

It is tempting to say this connection backwards (“ j connected to k ”), so please be careful with that- It is the opposite of what you might think it is.

Now continuing, we apply function f to the prestate vector P (element-wise) to get the state vector S .

$$S^{(2)} = f\left(P^{(2)}\right) = f\left(W^{(1)}\mathbf{x} + \mathbf{b}^{(i)}\right)$$

Now we apply an affine transformation from \mathbb{R}^k to the m dimensional output layer. This means that the second weight matrix has dimensions $m \times k$:

$$P^{(3)} = W^{(2)}S^{(2)} + \mathbf{b}^{(2)}$$

Typically, we don't do any more processing, so we can think of the activation function f on the final layer as the identity function, $f(x) = x$.

$$S^{(3)} = P^{(3)} = W^{(2)}S^{(2)} + \mathbf{b}^{(2)} = W^{(2)}\left(f\left(W^{(1)}\mathbf{x} + \mathbf{b}^{(i)}\right)\right) + \mathbf{b}^{(2)}$$

To be consistent with all layers, we usually define the prestate at the input layer to be just \mathbf{x} , and the activation function as the identity, so the layer doesn't do anything but output the same as what is input. With that, the computations a neural network takes can be described by the chain:

$$\text{input} = (P^{(1)} \rightarrow S^{(1)}) \rightarrow (P^{(2)} \rightarrow S^{(2)}) \rightarrow (P^{(3)} \rightarrow S^{(3)}) = \text{output}$$

It is simple now to see how to add more layers. We can keep extending this chain almost indefinitely.

The Network Produces Composition(s)

There is one important way to look at this sequence of operations that has some implications when we compute derivatives: This sequence of operations is function composition, which is clear when we write the output layer in terms of the input layer:

$$S^{(3)} = W^{(2)}\left(f\left(W^{(1)}\mathbf{x} + \mathbf{b}^{(i)}\right)\right) + \mathbf{b}^{(2)}$$

If we add more layers, then the state at the third layer should be written again as $f(P^{(3)})$, because presumably, f would not now be the identity. In that case, the output at the added fourth layer would be:

$$S^{(4)} = W^{(3)}\left[f\left(W^{(2)}\left(f\left(W^{(1)}\mathbf{x} + \mathbf{b}^{(i)}\right)\right) + \mathbf{b}^{(2)}\right)\right] + \mathbf{b}^{(3)}$$

and so on.

13.3 The Activation Function

The function f here is key. If f is linear, then the neural net is a linear network (we've already discussed those). Importantly, f on the hidden layer must be nonlinear. And importantly, in 1989, Cybenko proved the result about universal function approximation assuming that f is a **sigmoidal function**.

Definition: Suppose a function $\sigma(x)$ has the following properties:

- It is monotonically increasing.

- $\lim_{x \rightarrow -\infty} \sigma(x) = A$
- $\lim_{x \rightarrow \infty} \sigma(x) = B$

Then $\sigma(x)$ is called a **sigmoidal function**.

Common choices for the sigmoidal function include:

1. The log sigmoidal function (in Matlab, this is `logsig`)

$$\text{logsig}(x) = \frac{1}{1 + e^{-x}}$$

2. The hyperbolic tangent (in Matlab, this is `tansig`)

$$\tanh(x) = \frac{e^{2x} - 1}{e^{2x} + 1}$$


Typically, we will use the first option, since its derivative is very easy to compute (see the exercises), and typically the so-called “linear activation function”, $\sigma(x) = x$ is used at the input and output layer.

13.3.1 Recent Developments for the Activation Function

While the sigmoidal function was used by Cybenko in his proof, it has some shortcomings that didn’t become clear until very large networks were being trained. The two main shortcomings:

- The “squashing” function becomes saturated.
The sigmoidal only has a small interval on which it can keep input points apart (see the graph below). If $|x|$ is greater than about 2, then all positive points start to map to 1, and all negative points to 0 (this is explored more in the exercises).
- The vanishing gradient problem.

From the graph, see if you can estimate the derivative. For most x , the derivative is less than 1. Recall that a neural net is function composition so that the derivative is multiplication. What happens when you multiply a bunch of numbers that are less than 1? The values go to zero- This is the vanishing gradient problem (we’ll discuss this more in the section on training).



The Rise of ReLU

The activation function that most use now is the “rectified linear unit”, or **ReLU**. Mathematically, this is defined as:

$$\text{ReLU}(x) = \begin{cases} 0 & \text{if } x < 0 \\ x & \text{if } x \geq 0 \end{cases}$$

The perhaps obvious issues with using this function are (1) nondifferentiable at 0, and (2) it becomes unbounded. Further, if a learning rate is set too high, it has been observed that the ReLU’s can be “pushed” out so far that they are basically zero- this is the “Dying ReLU problem”.

In spite of the difficulties, in 2011 it was found that the ReLU function is a viable alternative to the sigmoidal function, even though it does not have all the properties of a sigmoidal. A few short years later, in 2017 it was found to be the most popular activation function for deep networks.

Exercises

1. Sketch the ReLU transfer function and its derivative.
2. If $\mathbf{x} \in \mathbb{R}^n$ and our targets $\mathbf{t} \in \mathbb{R}^m$, and we use k nodes in the hidden layer, how many unknown parameters do we have to find?

As you are constructing your network, keep this number in mind. In particular, you should have at least several data points for each unknown parameter that you are looking for.

3. We have to be somewhat careful when our data is badly scaled. For example, complete this table of values:

x	0	0.5	1	10	40	100
$\tanh(x)$						
$\text{logsig}(x)$						

You should see that, while the change in x between neighboring points is getting very large, that the corresponding changes in y are going to zero. This is problematic- If your data is large, or the weights are large, then the sigmoidal may simply start to output a constant.

This phenomenon goes by the name of *saturation*.

4. Some people like to scale the sigmoidal function by an extra parameter, β , that is $\sigma(\beta x)$. Show by sketching what happens to the graph of the sigmoidal (either the **tansig** or **logsig**) as you change β .

It is not necessary to scale the sigmoidal, as this is equivalent to scaling the data instead (via the weights).

5. Show that if $\sigma(x)$ is the **logsig** function, then

$$\sigma'(\beta x) = \beta \sigma(\beta x)(1 - \sigma(\beta x))$$

This is the reason this function is so popular- It is very easy to compute its derivative.

6. Show, using the definition of the hyperbolic sine and cosine, that the hyperbolic tangent can be written as:

$$\tanh(x) = \frac{\sinh(x)}{\cosh(x)}$$

7. Show that the hyperbolic tangent can be computed as:

$$\tanh(x) = \frac{2}{1 + e^{-2x}} - 1$$

(Matlab claims that this version is faster, but warns about possible numerical error)

8. Other Extensions of the activation function

Some other interesting activation functions can be used at the nodes. Here are a couple of unique ones- They are used to encode circular or spherical information:

(a) The Circular Node (two inputs, two outputs per node):

$$\sigma(x, y) = \left(\frac{x}{\sqrt{x^2 + y^2}}, \frac{y}{\sqrt{x^2 + y^2}} \right)$$

(b) The Spherical Node (three inputs, three outputs):

$$\sigma(x, y, z) = \left(\frac{x}{\sqrt{x^2 + y^2 + z^2}}, \frac{y}{\sqrt{x^2 + y^2 + z^2}}, \frac{z}{\sqrt{x^2 + y^2 + z^2}} \right)$$

See [19, 20] for examples of how to implement the last two transfer function types.

13.4 Training and Error

To define a three layer neural network in the form $n - k - m$, we should first define the activation function f . Although we could define a different function for every neuron, we typically will use the same transfer function for all the neurons in a single layer.

Once that is done, then we have to find matrices $W^{(1)}, W^{(2)}$ and the bias vectors $\mathbf{b}^{(1)}, \mathbf{b}^{(2)}$. Altogether, this makes $(nk + k) + (mk + m)$ parameters. Ideally, we would have much more data than that in order to get good estimates. In any case, we want to minimize the usual sum of squared error:

$$E(W^{(1)}, W^{(2)}, \mathbf{b}^{(1)}, \mathbf{b}^{(2)}) = \frac{1}{2} \sum_{i=1}^p \|\mathbf{t}^{(i)} - \mathbf{y}^{(i)}\|^2$$

where $\mathbf{y}^{(i)}$ is the output of the neural net using the i^{th} input. In the case of a single hidden layer, we have:

$$\mathbf{y}^{(i)} = W^{(2)} \left(\sigma \left(W^{(1)} \mathbf{x}^{(i)} + \mathbf{b}^{(1)} \right) \right) + \mathbf{b}^{(2)}$$

As usual, “training” means finding the weights and biases that minimize the error function, and now we’re back at function optimization. We’ve discussed these methods before, but here is a short list:

- Method of Steepest Descent (or Gradient Descent).
- Stochastic Gradient Descent is an excellent one to use when we’re programming it ourselves, since we only have to deal with the derivative one data point at a time.
- Newton’s Method (an indirect method, solving for where the derivative of the error is 0).
- Conjuage Gradient (Search along the eigenvectors of the Hessian of the error)
- Levenburg-Marquardt (A combination of the techniques above).