

Introduction to programming in R: neural networks

Erik Spence

SciNet HPC Consortium

5 April 2023

To find today's slides, go to the "Introduction to programming in R" page, under Lectures, "Neural networks".


`https://scinet.courses/1277`

Neural networks are particularly good at detecting patterns, and for certain problems perform better than any other known class of algorithm. Neural networks are used for

- Image recognition, object detection (pneumonia, cancer).
- Medical diagnosis.
- Natural language processing (voice recognition).
- Novelty detection (detection of outliers).
- Next-word predictions.
- Text sentiment analysis.
- System control (self-driving cars).

Neural networks are finding their way into everything.

Consider the problem of hand-written digit recognition:



How would you go about writing a program which can tell you what digits are displayed?


- All the algorithms you might use to describe what a given number "looks like" are extremely difficult to implement in code. Where do you even start?
- And yet humans can easily tell what these digits are.
- Neural networks are based on a "biologically inspired" approach to solving such classification problems.
- This is one of the classic problems which have been solved using neural networks.

Rather than focus on the details of what individual numbers look like, we will instead ignore those details altogether. We will use a completely different approach:

- Break the data set of numbers into two or three groups: training, testing, and optionally validation.
- As with other supervised machine-learning algorithms, feed the training data to the neural network and train it to recognize one number from another.
- Rather than focus on details of the numbers, let the neural network figure out the details for itself.

This is the goal of this class.

Neural networks are built upon "neurons". This is just a fancy way of saying a "function that takes multiple inputs and returns a single output".



The function which the neuron implements is up to the programmer, but it must contain free parameters so that the network can be trained. These functions usually take the form

$$f(x_1, x_2, x_3) = f\left(\sum_{i=1}^3 w_i x_i + b\right) = f(\mathbf{w} \cdot \mathbf{x} + b)$$

Where \mathbf{w} are the 'weights' and b is the 'bias'. These are the trainable parameters.

What function should we use for f ? One which is usually used, at least when initially teaching about neural networks, is the "sigmoid function" (also called the "logistic function").


$$\sigma(z) = \frac{1}{1 + e^{-z}}$$

And so our neuron function becomes

$$f(x_1, x_2, x_3) = f(w \cdot x + b) = \frac{1}{1 + e^{-(w \cdot x + b)}}$$


Where again w are the 'weights' and b is the 'bias'.

Why the sigmoid function?



Because it ranges from 0 to 1 smoothly.

Suppose we combine many neurons together, into a proper network, consisting of "layers".



Some details about the graphic on the previous slide:

- The input neurons do not contain any functions. They merely represent the input data being fed into the network.
- Each neuron in the 'hidden' layers and the output layer all contain functions with their own free parameters, \mathbf{w} and \mathbf{b} .
- Each neuron outputs a single value. This output is passed to all of the neurons in the subsequent layer. This type of layer is known as a "fully-connected", or "dense", layer.
- The number of free parameters in the neurons in any given layer depends upon the number of neurons in the previous layer.
- The output from the output layer is aggregated into the desired form to calculate the cost function.

You might legitimately wonder why on Earth we would think this would lead anywhere.

- As it happens, this topology is similar to some simple biological neural networks.
- Each layer takes the output of the previous layer as its input.
- Each layer makes "decisions" about the information that it receives.
- In this way the later layers are able to make more complex and abstract decisions than the earlier layers.
- A many-layered network can potentially make sophisticated decisions.

However, there are subtleties in training such a network.

How do we optimize the weights and biases? We need to define some sort of "cost function" (sometimes called "loss" or "objective" function):

$$C = \frac{1}{2} \sum_i (g(\mathbf{x}_i) - \mathbf{y}_i)^2$$

where g is our neural network, and \mathbf{y}_i are the correct answers associated with each \mathbf{x}_i . Here we are using the "quadratic" cost function.

We then use an optimization algorithm to search for the values of \mathbf{w} and \mathbf{b} which generate the minimum of C , given the data \mathbf{x} and \mathbf{y} . In our example we will use the Gradient Descent algorithm to find this minimum.

Suppose the function we want to minimize has only one parameter.

$$C = w^2$$


Supposed we've guessed that the minimum of C is w_i , and we wish to improve the guess.

Gradient descent says to move according to the formula:

$$w_{i+1} = w_i - \eta \frac{\partial C}{\partial w_i}$$

where η is called the step size. We then repeat until some stopping criterion is satisfied.

If we have multiple parameters, we step them all.



How do we apply Gradient Descent to a neural network?

- Suppose that we decide to try to use gradient descent to train the network from five slides ago (slide 9).
- Each of the neurons has its own set of free parameters, \mathbf{w} and \mathbf{b} . There are lots of free parameters!
- To update the parameters we need to calculate every $\frac{\partial C}{\partial w_i}$ and $\frac{\partial C}{\partial b}$ for every neuron!
- But how do we calculate those derivatives, especially for the parameters associated with the neurons that are several layers away from the output?

Actually, as it happens, this is a solved problem. The algorithm is called Backpropagation, but we won't cover it today.



One of the classic data sets on which to test neural-network techniques is the MNIST data set.


- A database of handwritten digits, compiled by NIST.
- Contains 60000 training, and 10000 test examples.
- The training digits were written by 250 different people; the test data by 250 different people.
- The digits have been size-normalized and centred.
- Each image is grey scale, 28×28 pixels.

We can create a neural network to classify these digits.

How would we design a network to analyze this data?

- Each image is $28 \times 28 = 784$ pixels. Let the input layer consist of 784 input nodes. Each node will consist of the grey value for that pixel.
- The output will consist of a one-hot-encoding of the networks analysis of the input data. This means that, if the input image depicts a '7', the output vector should be $[0, 0, 0, 0, 0, 0, 0, 1, 0, 0]$.
- Thus, let there be 10 output nodes, one for each possible digit.
- To start, let's just use a single hidden layer.

Fortunately, packages exist which make coding such a network quite easy.



Now that we have a plan for our network, how are we going to code it? The standard way is to use a neural network 'framework'. Why would you do that?

- Coding your own networks from scratch can be a bit of work.
- Neural network (NN) frameworks have been specifically designed to solve NN problems.
- R, of course, is not a high-performance language.
- The NN frameworks which have been developed are compiled before being used, thus being much faster than interpreted R.
- The NN frameworks are also designed to use GPUs, which make things significantly faster than just using CPUs.
- Standard NN frameworks include TensorFlow, Torch, MXNet, Caffe and many many others.

We will use Keras, on a TensorFlow backend.

We will use Keras on top of TensorFlow.

- Keras is a NN framework, but it's only the top-most level.
- More accurately, it's an API standard for creating neural networks.
- Designed for fast development of networks.
- The original version ran on top of a 'back end', which by default is now TensorFlow, as Keras is being absorbed into TensorFlow.
- Historically it ran on top of many other backends also: Theano, CNTK, MXNet, TypeScript, JavaScript, PlaidML, Scala, CoreML, and others.
- Because it's a proper framework, all of the NN goodies you need are already built into it.
- Has an R interface available for it.

Keras in R can be a pain to install, because it will want to be run on top a of Python, using TensorFlow. If you don't have Python installed already it will probably want to install it.

Let us implement our neural network using Keras. First let us get the MNIST data.

- Keras now requires a non-system version of Python to be installed on your machine. When you install Keras it will ask you to install Miniconda, which is a distribution of Python.
- The data can be automatically downloaded by Keras. But when it does so it will install TensorFlow into your Miniconda distribution.
- The data comes pre-split into training and testing data sets.

```
>
> library(keras)
>
> mnist <- dataset_mnist()
>
> x.train <- mnist$train$x
> y.train <- mnist$train$y
> x.test <- mnist$test$x
> y.test <- mnist$test$y
>
> str(x.train)
int [1:60000, 1:28, 1:28] 0 0 0 0 0 0 ...
>
```

We now need to pre-process the data into the form we want.

- We use the "array_reshape" command to change the data to 1-D.
- We use this, rather than "dim(x.train) <- c(..)" to maintain the row-major nature of the data (R is column-major).
- We need to one-hot encode the target data. We use the "to_categorical" function to do this.

```
>
> x.train <- array_reshape(x.train,
+   c(nrow(x.train), 784))
> x.test <- array_reshape(x.test,
+   c(nrow(x.test), 784))
>
> str(y.train)
int [1:60000(1d)] 5 0 4 1 9 2 1 3 4 ...
>
> y.train <- to_categorical(y.train, 10)
> y.test <- to_categorical(y.test, 10)
>
> str(y.train)
int [1:60000, 1:10] 0 1 0 0 0 0 0 0 0 ...
>
```

We can now implement our neural network using Keras, with 30 neurons in the hidden layer.

- A "Sequential" model means the layers are stacked on one another in a linear fashion.
- A "Dense" ("fully-connected") layer is the layer we've already discussed.
- Use "input_dim" in the first layer to indicate the shape of the incoming data.
- The "activation" is the output function of the neuron.

```
> model <- keras_model_sequential()
> model %>% layer_dense(units = 30,
+   activation = 'sigmoid', input_shape = c(784))
> model %>%
+   layer_dense(units = 10, activation = 'sigmoid')
>
> summary(model)
-----
Layer (type) Output Shape Param #
-----
dense_1 (Dense) (None, 30) 23550
-----
dense_2 (Dense) (None, 10) 310
-----
Total params: 23,860
Trainable params: 23,860
Non-trainable params: 0
```

Now that the network is constructed, it must be compiled.

- The loss (cost) function must be specified.
- The optimizer indicates what minimization algorithm to use. Here we use Stochastic Gradient Descent (SGD), which is a variation on regular Gradient Descent.
- The 'metrics' flag indicates what to print out during the training of the network.
- The 'fit' command is used to execute the training.
- The number of epochs, and the batch size, are parameters which apply to the optimization algorithm.

```
>
> model %>% compile(loss = 'mean_squared_error', optimizer = optimizer_sgd(),
+   metrics = c('accuracy'))
>
> fit <- model %>% fit(x.train, y.train, epochs = 200, batch_size = 128)
Epoch 1/200
60000/60000 [=====] - 0s - loss: 0.1368 - acc: 0.1933
Epoch 2/200
60000/60000 [=====] - 0s - loss: 0.0923 - acc: 0.3126
:
Epoch 199/200
60000/60000 [=====] - 0s - loss: 0.0225 - acc: 0.8947
Epoch 200/200
60000/60000 [=====] - 0s - loss: 0.0225 - acc: 0.8947
>
```


Now check against the test data.
88.6%!

This isn't great. We can do better, but the techniques to do so are beyond the scope of today's class.

```
>
> pred <- model %>% evaluate(x.test, y.test)
5 7616/10000 [=====>.....] - ETA: 0s
>
> pred[['accuracy']]
[1] 0.8859
>
```

We can do better. What might we do? There are a few simple approaches that might be explored.

- Change the activation function.
- Change the cost function.
- Change the optimization algorithm.
- Change the way things are initialized.
- Add regularization, to try to deal with over-fitting.

The most important technique, however:

- Completely overhaul the network strategy.

This field is huge; we've barely scratched the surface.

You've probably heard the term. What is Deep Learning?

- Quite simply: a neural network with many hidden layers.
- Up until the mid-2000s neural network research was dominated by "shallow" networks, networks with only 1 or 2 hidden layers.
- The breakthrough came in discovering that it was practical to train networks with a larger number of hidden layers.
- But it only became practical with the advent of sufficient computing power (GPUs) and easily-accessible huge data sets.
- State-of-the-art networks today can contain dozens of layers.

Neural network classes:

- <http://neuralnetworksanddeeplearning.com>
- <http://www.cs.utoronto.ca/~fidler/teaching/2015/CSC2523.html>
- <http://cs231n.stanford.edu>

I am offering a SciNet neural network programming class starting in April. The class is taught in Python.