

## CS51A - Neural Network Lab



IN CS, IT CAN BE HARD TO EXPLAIN THE DIFFERENCE BETWEEN THE EASY AND THE VIRTUALLY IMPOSSIBLE.

<http://www.xkcd.com/1425/>

In this lab, we will experiment with a few simple neural networks. The idea is to obtain preliminary answers to questions like:

- "How quickly can a network be trained?"
- "How reliable (or repeatable) is the training?"
- "How accurate is the resulting network?"

For the last half an hour of lab (or, when it seems like people have mostly wrapped up their experiments), each pair/group will give a short 1-2 minute, informal presentation on what they found. Although there is nothing to submit for this lab, you should keep notes and document your results for the presentation/discussion.

### **Intro to the neural net package**

To start with, read this section and play with a few of the examples to familiarize yourself with the neural network package we'll be using.

- Create a directory called `nnlab` somewhere on your computer.
- Download the following zip file into your `nnlab` directory, unzip it, and then delete the `.zip` file.

`http://www.cs.pomona.edu/classes/cs51a/assignments/neural.zip`

- Open up Wing101 and create a new file. Add the following import statement at the top

```
from cs51neural import *
```

and then save the file as `nnlab.py` in your `nnlab` directory.

## Data

The neural network package trains on labeled training data. Each example is a tuple consisting of the data and the label, both of which are represented as lists. For example, we saw the `xor` operator in class:

data		label
0	0	0
0	1	1
1	0	1
1	1	0

This data would be represented as:

```
xor = [([0, 0], [0]),
       ([0, 1], [1]),
       ([1, 0], [1]),
       ([1, 1], [0])]
```

I've written it across multiple lines to highlight the structure, but it could also have been written as:

```
xor = [([0, 0], [0]), ([0, 1], [1]), ([1, 0], [1]), ([1, 1], [0])]
```

## Training the model

To use the neural network model, you first have to create a new instance of the neural network class. To create a new neural network you need to specify the configuration of the network. Our neural network package implements a two-layer neural network, so you need to specify three parameters: the number of inputs into the network (this should be the number of features/dimensions your data has), the number of hidden nodes, and the number of output nodes (one, for most of our data).

For example, to create a network with two inputs, three hidden nodes, and one output node you would write:

```
nn = NeuralNet(2, 3, 1)
```

(Make sure you've imported the neural network package by either running your `nnlab.py` file or typing the import statement above.)

Calling `NeuralNet` creates a new neural network object and we save it into a variable called `nn`.

Now, just like other objects that we've used (e.g., lists, strings, dictionaries), we can call the methods on that object. For example, to train the model on the xor data (assuming you've defined the `xor` variable already) we can do the following:

```
>>> nn.train(xor)
error 0.490028
error 0.388845
error 0.212502
error 0.088414
error 0.045714
error 0.028693
error 0.020248
error 0.015382
error 0.012277
error 0.010150
```

(Note that your numbers will be slightly different since the network starts off with random weights, though it should be roughly in the same relative range.)

The default is that the program will train for 1000 iterations and it will output the training error (i.e., the error on the data that the network is training on) after every 100 iterations. As the model trains, you'd likely see the training errors generally getting smaller.

Once you've trained a model, you can apply it to data using the `test` method. For example, we could test on the training data (though, in practice, train and test should be different, it's a simple way to do a sanity check on your model).

```
>>> nn.test(xor)
([(0, 0), [0], [0.06555671027717622]], ([0, 1], [1], [0.9497613925132822]), ([1, 0], [1],
 ([1, 1], [0], [0.06977380989233718]))]
```

`test` returns the triplets of: data, label/correct answer, model output. You can process them in a more friendly way by unpacking the tuple:

```
for triple in nn.test(xor):
    (input, label, prediction) = triple # unpack the triple
    # remember outputs are lists, so we need to get the first item from the list
    print(str(input) + "\t" + str(label[0]) + "\t" + str(prediction[0]))
```

If you ran this, you'd see the predictions are starting to get pretty close to the real values:

```
[0, 0] 0 0.06555671027717622
[0, 1] 1 0.9497613925132822
[1, 0] 1 0.9082703507697025
[1, 1] 0 0.06977380989233718
```

If you want to evaluate just one input, you may do it with the `evaluate` method. For example, if we wanted to see what the answer was for `([0,0])`:

```
>>> print(nn.evaluate([0, 0]))
[0.06555671027717622]
```

This is the basic functionality of the neural network class. See the Appendix at the end of this document for more information about the different methods. In particular, take a look at the documentation for the `train` method since it has a number of optional parameters that you can specify to change how the training works. For example, to train the model for 2000 iterations we can specify the `iterations` optional parameter:

```
>>> nn.train(xor, iterations=2000)
error 0.003163
error 0.002973
error 0.002803
error 0.002651
error 0.002514
error 0.002389
error 0.002276
error 0.002173
error 0.002078
error 0.001991
error 0.001911
error 0.001837
error 0.001768
error 0.001704
error 0.001644
error 0.001588
error 0.001536
error 0.001487
error 0.001440
error 0.001397
```

It prints out 20 errors now, one every 100 iterations. If I wanted to do this same thing, but only print out the errors every 200 iterations:

```

>>> nn.train(xor, iterations=2000, print_interval=200)
error 0.001317
error 0.001246
error 0.001181
error 0.001123
error 0.001070
error 0.001022
error 0.000978
error 0.000937
error 0.000899
error 0.000865

```

## Experiments

Now that you're comfortable using the neural network class, let's run a few experiments. Jot down notes as you go for discussion later.

1. Construct a network with two hidden nodes and train it on the XOR data below. Notice that we have complete information about the function we are trying to approximate—an unusual situation for a neural network. How quickly do the weights/errors converge, i.e., to the point where the difference from iteration to iteration is small? How well does the resulting network perform? Try training several different networks to see the variation in the convergence rate. You may also want to reduce `print_interval` to get finer information about the changes in the error.

inputs		output
0	0	0
0	1	1
1	0	1
1	1	0

2. Repeat part 1 with a network with eight hidden nodes. Does the convergence go faster? Is the resulting network a better approximation to the XOR function?

Notice that “faster” can mean “fewer iterations” or “less clock time.” A network with more nodes will have more weights to adjust and will take more clock time for each training cycle—but it *may* require fewer iterations to perform well and the total time *may* be shorter.

3. See what happens when you repeat part 1 using a network with just one hidden node. The XOR function cannot be computed with a single node. Why is a network with one hidden node equivalent to a single node?
4. Table 1 contains a sampling of voter opinions. The idea is to deduce voters party affiliations from their views on the importance of various topics. Six voters were asked to rate the

importance of five issues on a scale from 0.0 to 1.0 and to identify themselves as Democrat (0.0) or Republican (1.0).

Write a neural network and train it on the data in Table 1. Then try it on the samples from Table 2 or other cases of your own creation. Can you explain the conclusions drawn by the network?

5. Once you've finished these experiments, look at the three questions at the beginning of the handout and decide which one is the most interesting to you. Spend a bit more time playing with examples and coming up with concrete data/examples that illustrate any observations that you have. Be prepared to share these at the end of lab.

budget	defense	crime	environment	social security	party
0.9	0.6	0.8	0.3	0.1	1.0
0.8	0.8	0.4	0.6	0.4	1.0
0.7	0.2	0.4	0.6	0.3	1.0
0.5	0.5	0.8	0.4	0.8	0.0
0.3	0.1	0.6	0.8	0.8	0.0
0.6	0.3	0.4	0.3	0.6	0.0

Table 1: A sampling of voter opinions. This example is taken from notes by Dave Reed of Creighton University.

budget	defense	crime	environment	social security	party
1.0	1.0	1.0	0.1	0.1	?
0.5	0.2	0.2	0.7	0.7	?
0.8	0.3	0.3	0.3	0.8	?
0.8	0.3	0.3	0.8	0.3	?
0.9	0.8	0.8	0.3	0.6	?

Table 2: Some test cases for the voter network.

And in a format that's friendly for our package:

```
voter_train = [(0.9,0.6,0.8,0.3,0.1),[1.0]),
               (0.7,0.2,0.4,0.6,0.3),[1.0]),
               (0.5,0.5,0.8,0.4,0.8),[0.0]),
               (0.3,0.1,0.6,0.8,0.8),[0.0]),
               (0.6,0.3,0.4,0.3,0.6),[0.0))]
```

## Appendix: The Neural Network Class

The software we're playing with represents a neural network as an *object* of the `NeuralNet` class. We've already seen instances of this, for example, *lists* are a class of objects.

The main difference between a `list` object and our neural network object is that there is no special way for constructing neural network objects (to create lists we use the square brackets, `[]`). Instead, there is a special function called a *constructor* that takes some number of arguments (in our case, three) and creates a neural network object. Once you have that object, you can call methods on it just like you would any other object.

In the following summary, the neural network object is named `nn`. You may, of course, give your networks any name you like.

`nn = NeuralNet(num_input, num_hidden, num_output)` this calls the constructor and creates a network with the specified number of nodes in each layer. The initial weights are random values between  $-2.0$  and  $2.0$ . Notice that all of the other methods below are called *on* a neural network object that has been created.

`nn.evaluate(input)` returns the output of the neural network when it is presented with the given input. Remember that the input and output are *lists*.

`nn.train(training_data)` carries out a training cycle. As specified earlier, the training data is a list of input-output pairs. There are four optional arguments to the `train` function:

`learning_rate` defaults to 0.5.

`momentum_factor` defaults to 0.1. The idea of momentum is discussed in the next section. Set it to 0 to suppress the affect of the momentum in the calculation.

`iterations` defaults to 1000. It specifies the number of passes over the training data.

`print_interval` defaults to 100. The value of the error is displayed after `print_interval` passes over the data; we hope to see the value decreasing. Set the value to 0 if you do not want to see the error values.

You may specify some, or all, of the optional parameters by name in the following format.

```
nn.train(training_data,
         learning_rate=0.8,
         momentum_factor=0.0,
         iterations=100,
         print_interval=5)
```

`nn.test(testing_data)` evaluates the network on a list of examples. As mentioned above, the testing data is presented in the same format as the training data. The result is a list of triples which can be used in further evaluation.

`nn.get_IH_weights()` returns a list of lists representing the weights between the input and hidden layers. If there are  $t$  input nodes and  $u$  hidden nodes, the result will be a list containing  $t$  lists of length  $u$ .



`nn.get_H0_weights()` returns a list of lists representing the weights between the input and hidden layers. If there are  $u$  hidden nodes and  $v$  output nodes, the result will be a list containing  $u$  lists of length  $v$ .

`nn.use_tanh()` changes the activation function from the sigmoid function to a commonly used alternative, the hyperbolic tangent function.