

Neural Networks -- two inputs, one hidden layer

```
In[1333]:= x = 0; Remove["Global`*"]; DateList[Date[]] // Most
```

```
Out[1333]= {2020, 5, 11, 23, 2}
```

We now examine a simple two-dimensional, single hidden layer, neural network.

Define neural net nodes and the approximation

G will be our neuron. Make G evaluate vectors elementwise

```
In[1334]:= G[list_List] := G /@ list
```

We will first create a small neural net so that you can see the structure

```
In[1335]:= numnodes = 3; dim = 2;
```

```
In[1336]:= amat = Table[wi,j, {i, 1, numnodes}, {j, 1, dim}];
bvec = Table[bi, {i, 1, numnodes}];
basis = G /@ (amat . {x, y} + bvec)
```

```
Out[1338]:= {G[b1 + x w1,1 + y w1,2], G[b2 + x w2,1 + y w2,2], G[b3 + x w3,1 + y w3,2] }
```

```
In[1339]:= cvec = Table[ci, {i, Length[basis]}]
```

```
Out[1339]:= {c1, c2, c3}
```

```
In[1340]:= model = Sum[ci basis[[i]], {i, Length[basis]}] + d0
```

```
Out[1340]:= G[b1 + x w1,1 + y w1,2] c1 + G[b2 + x w2,1 + y w2,2] c2 + G[b3 + x w3,1 + y w3,2] c3 + d0
```

Math time

Mathematically, a neuron is equivalent to the function:

$$Y = \theta \left(\sum_{i=1}^n W_i X_i + b \right),$$

which can be conveniently modeled, using a matrix form,

$$Y = \theta(W \cdot X + b),$$

where $W = [W_1 \quad W_2 \quad \dots \quad W_n]$, and $X = \begin{bmatrix} X_1 \\ X_2 \\ \vdots \\ X_n \end{bmatrix}$.

A layer of neurons can be conveniently represented, using matrix notation, as follows:

$$W = \begin{bmatrix} W_{11} & \dots & W_{1M} \\ \vdots & \vdots & \vdots \\ W_{N1} & \dots & W_{NM} \end{bmatrix}.$$

The row index in each element of this matrix represents the destination neuron of the corresponding connection, whereas the column index refers to the input source of the connection.

Designating by Y the output of the layer, you can write

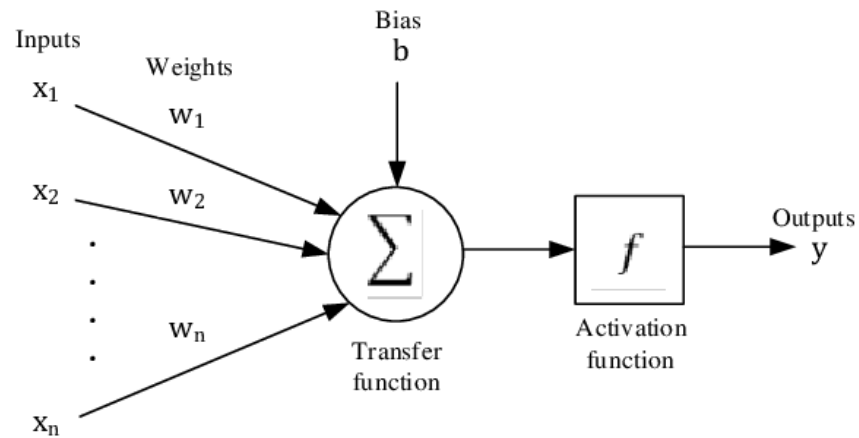
$$Y = \begin{bmatrix} Y_1 \\ \vdots \\ Y_i \\ \vdots \\ Y_N \end{bmatrix} = \begin{bmatrix} \theta \left(\sum_{j=1}^M W_{1j} X_j + b_1 \right) \\ \vdots \\ \theta \left(\sum_{j=1}^M W_{ij} X_j + b_i \right) \\ \vdots \\ \theta \left(\sum_{j=1}^M W_{Nj} X_j + b_N \right) \end{bmatrix} = \theta(W \cdot X + b),$$

The function achieved by this network is

$$Y^3 = \begin{bmatrix} Y_1^3 \\ \vdots \\ Y_i^3 \\ \vdots \\ Y_{N3}^3 \end{bmatrix} = \theta(W^3 Y^2 + b^3) = \theta(W^3 \theta(W^2 Y^1 + b^2) + b^3) = \theta(W^3 \theta(W^2 (\theta(W^1 X + b^1)) + b^2) + b^3).$$

Picture time

The neural net literature is filled with pictures, leaving the reader to decode the hieroglyphics in order to understand the underlying math. Here we see a common picture representing a single neuron with multiple inputs and a scalar output



Note the jargon:

Transfer function is nothing more than the weighted sum of the inputs plus the bias.


Bias is just an additive term

Activation function is just a function of one variable

“Outputs” should be “output” since it is a scalar for an individual neuron.

Fortunately, they got “inputs” right.

The picture above represents a single neuron. Here is what a full single hidden layer neural network looks like where each circle represents a neuron



We now specify a more serious net

```
In[1341]:= numnodes = 9; dim = 2;  
amat = Table[wi,j, {i, 1, numnodes}, {j, 1, dim}];  
bvec = Table[bi, {i, 1, numnodes}];  
basis = G /@ (amat. {x, y} + bvec);  
cvec = Table[ci, {i, Length[basis]}];  
model = Sum[ci basis[[i]], {i, Length[basis]}] + d0;
```


Define node function

We use sigmoid function EXCEPT we avoid overflow and under flow

```
In[1347]:= G[x_] = If[-20 ≤ x ≤ 20, 1 / (1 + Exp[-x]), If[x < -20, 0, 1]]
```

```
Out[1347]= If[-20 ≤ x ≤ 20,  $\frac{1}{1 + \text{Exp}[-x]}$ , If[x < -20, 0, 1]]
```

```
In[1348]:= Plot[G[x], {x, -10, 10}]
```



Set up the optimization

```
In[1349]:= vars = Join[amat, bvec, cvec, {d0}] // Flatten
```


```
Out[1349]= {w1,1, w1,2, w2,1, w2,2, w3,1, w3,2, w4,1, w4,2, w5,1, w5,2, w6,1, w6,2, w7,1, w7,2,  
w8,1, w8,2, w9,1, w9,2, b1, b2, b3, b4, b5, b6, b7, b8, b9, c1, c2, c3, c4, c5, c6, c7, c8, c9, d0}
```

```
In[1350]:= init = vars - vars + 1;
```

```
varsin = {vars, init} // Transpose;
```

Fit a 2D example function


```
In[1352]:= f[x_, y_] = Exp[-(x2/10 + 2 y2/8)];  
Plot3D[f[x, y], {x, -5, 5}, {y, -5, 5}, PlotRange -> All]
```



```
In[1354]:= mpts = 10;  
           npts = 2 mpts + 1;  
           xpts = (Range[npts] - mpts - 1) / 2.
```

```
Out[1356]= {-5., -4.5, -4., -3.5, -3., -2.5, -2., -1.5, -1., -0.5, 0., 0.5, 1., 1.5, 2., 2.5, 3., 3.5, 4., 4.5, 5.}
```

```
In[1357]:= inputs = Flatten[Outer[List, xpts, xpts], 1];  
           inputsT = inputs // Transpose;  
           zpts = f@@@ inputs;  
           data = Join[inputsT, {zpts}];  
           dataT = data // Transpose;  
           ListPointPlot3D[dataT]
```



Let's use a global optimizer with multiple initial guesses to solve the least squares problem

```
In[1363]:= fit = FindFit[dataT, model, varsin, {x, y}, MaxIterations -> 100, Method -> "NMinimize"]
```

 **NMinimize:** Failed to converge to the requested accuracy or precision within 100 iterations.

```
Out[1363]= {w1,1 -> -0.19068, w1,2 -> 0.790665, w2,1 -> 0.221148, w2,2 -> -0.929863, w3,1 -> 0.0568638, w3,2 -> 0.970384, w4,1 -> 0.222165,
w4,2 -> 1.18064, w5,1 -> 1.0348, w5,2 -> 0.69876, w6,1 -> 0.520327, w6,2 -> 0.921683, w7,1 -> 1.07408, w7,2 -> 0.0869792,
w8,1 -> 0.713561, w8,2 -> -0.911572, w9,1 -> -1.15364, w9,2 -> 0.506211, b1 -> 0.341216, b2 -> -0.644581, b3 -> -3.02443,
b4 -> 1.21553, b5 -> 1.38668, b6 -> -1.22828, b7 -> -1.49811, b8 -> -1.50161, b9 -> -1.46302, c1 -> -9.2217, c2 -> -7.21993,
c3 -> 0.956138, c4 -> 1.25642, c5 -> 0.436296, c6 -> -1.07132, c7 -> -0.452691, c8 -> -0.738582, c9 -> -0.375862, d0 -> 8.05017}
```

Note:

Mathematica said it could not solve the problem to the standard 8-digit precision.

These problems are ill-conditioned, so we should use sloppy stopping rules.

Mathematica is using double precision arithmetic.


Most neural network software uses single precision.

If you want better, you need to set the options

“Better” costs you time.


Plot the neural net result

```
In[1364]:= modelf = Function[{x, y}, Evaluate[model /. fit]];  
Plot3D[modelf[x, y], {x, -5, 5}, {y, -5, 5}]
```




Plot the error. Note that the error is in $[-0.02, 0.02]$.

```
In[1366]:= Plot3D[f[x, y] - modelf[x, y], {x, -5, 5}, {y, -5, 5}]
```




```
In[1367]= ContourPlot[f[x, y] - modelf[x, y], {x, -5, 5}, {y, -5, 5}]
```




Let's look at the contours of the true function

```
In[1368]:= ContourPlot[f[x, y], {x, -5, 5}, {y, -5, 5}]
```



Here are the contours of the neural network

```
In[1369]:= ContourPlot[modelf[x, y], {x, -5, 5}, {y, -5, 5}]
```



Comments on the optimization step

The optimization problems are hideously large.

Standard methods are not used.

The dominant method is stochastic gradient and its refinements.

Stochastic gradient

Take steepest descent direction method

Need the gradient: NNs are designed to make this easy!

They call it backpropagation; we call it automatic differentiation

Computing gradient means evaluating NN for all input data

Too expensive!

Instead use a random sample

Use this stochastically approximated gradient to adjust all weights and biases

Software and Hardware

TensorFlow is dominant neural net package


The language looks horrible to me

There are Python front ends to make it easier

Challenge: You need to fight through the NN jargon

Massive parallelization is used.

It is the future



Becoming Human

Exploring Artificial Intelligence & What it Means to be Human