

Psych 258
Exercise 7
Due Nov 28

In this model you will create a neural network simulation that performs a parity (odd/even) judgment on the input units: if the # of input units turned on is even then the output is 1, otherwise the output is zero.

You can also use `rp_xor.m` (on the class web site) as a basis for your solution for this exercise.

Your program should be called `XX_parity.m` where `XX` is your last name. The steps that you should take are as follows.

1. set up the parameters of the model:
learning rate (`lrate`): for this problem, a learning rate of around 1 should be fine
number of input units: 8
number of hidden units: 3
number of output units: 1
2. Set up the patterns that the network will be trained on. This should be a matrix of ones and zeros, with the number of columns equal to the number of input units. The number of rows represents the number of patterns that the network will be trained on; for this exercise, 8 patterns should be sufficient. You can create these patterns by rounding random uniform variates created using `rand()`.

3. Determine the desired output for each pattern, which is 1 if the number of ones in the input is even and zero if the number is odd. Hint: you can use the `mod()` function to determine whether the number of on-elements is odd or even.
4. Create two matrices for the weights that connect the input to hidden units (`w_ih`), and hidden units to the output unit (`w_ho`). Fill these matrices with uniform random numbers between -0.5 and 0.5 .
5. Create a loop that will continue to iterate the model until one of two conditions is met:
 - the sum of squared error between the observed and desired output is less than $.01$
 - the number of epochs is greater than 1000

Each loop through the model should perform the following operations (I provide pseudocode for each of the equations):

- a) Pass the activation from the input units (which is simply the input pattern) to the hidden units:

```
input_to_hidden = patterns * w_ih
```

- b) determine hidden unit activation by passing the input to the hidden units through the activation function:

```
hidden_activation =  
activation_fn(input_to_hidden)
```

you should create a separate function (in its own file) called `activation_fn` that takes in a matrix of activation values and returns a sigmoid on those values:

```
function f=activation_fn(i)
    f=1./(1+exp(-i));
```

c) pass activation from the hidden units to the output units:

```
input_to_output =
hidden_activation*w_ho
```

d) determine output activity by passing the input to the output units through the activation function:

```
output_activation =
activation_fn(input_to_output)
```

e) compute the output error:

```
output_error = desired_output -
output_activation
```

f) determine the error signal (δ) for each layer of connections

```
delta_ho = output_error .* output_act
.* (1- output_act)
delta_ih = delta_ho*w_ho' .* hidden_act
.* (1- hidden_act)
```

g) determine the weight changes (dw) for each layer of connections

```
dw_ih = lrate * patterns' * delta_ih
dw_ho = lrate * hidden_act' * delta_ho
```

h) apply the weight changes

```
w_ih = w_ih + dw_ih
w_ho = w_ho + dw_ho
```

i) compute the sum of squared error using the matrix method:

```
sse = trace(output_error' * output_error)
```

trace() takes the sum of the diagonal elements in a square matrix; if you are not familiar with this, you may want to confirm for yourself that this operation does indeed compute the sum of squares. also, create a variable that will save the sse value for each epoch, so that these values can be plotted later.

j) print a brief report every 10 epochs, listing the epoch number and the SSE value

Once the model has converged ($SSE < .01$), print two figures:

- a plot of SSE over epochs

- a display of the matrices of desired and obtained outputs as images, using the `imagesc()` function. Show the images next to one another in a 1x2 subplot.

You will notice that on some occasions the model may not converge upon a solution to the problem, but instead will get stuck in a local minimum; if the model goes through 1000 epochs without converging (i.e., the SSE is still $> .01$), then the model should exit with a warning that it did not converge. When this happens, re-run the model (with a new set of random weights) until it converges.

After the model has converged, test the model's generalization abilities by creating a new set of patterns (again using rounded random numbers), and running those patterns through the model by performing steps a-e using the new patterns (note that you should not apply weight changes during the test). Does the model generalize well to the new stimuli?