Your job is to implement a multilayer perceptron (MLP) in Matlab and test its performance on three data sets provided along with this assignment.

The MLP will be implemented with two functions:

```
M = trainMLP(X, d, layers, phi, phiprime, eta, batch, alpha)
```

where

`X` is an m X N matrix, each of whose columns is a training vector
`d` is a kXN matrix, each of whose columns is an output vector
`layers` is a vector whose elements are the numbers of units in the hidden layers
`phi` is a Matlab function pointer to the activation function, a function of v
`phiprime` is a pointer to the derivative of `phi`, a function of `phi(v)`
`eta` is a pointer to the eta function, a function of the epoch number
`batch` takes value 1 if batch training is to be done, and 0 otherwise
`alpha` is a constant momentum term

`M` encodes the resulting MLP, including the weights and the activation function

and

```
y = evalMLP(x, M)
```

where

`x` is an input column vector of length m
`M` is the result of training an MLP

`y` is a result vector of length k

You should train your MLP for each input set using hyperbolic tangent activiation function with constants a and b as suggested by LeCun and report in section 4.6.3 of Haykin 3/e. In addition the weights should be initialized in accordance with the suggestions in section 4.6.6.

Your MLP implementation should be able to support arbitrary numbers of hidden layers. Use a configuration involving at least a two-hidden-layer network to solve one of the problems.

Your MLP should be able to use momentum for weight updates.

You shall use your MLP to solve the three problems by employing N-fold crossvalidation where $5 \leq N \leq 10$. Use some form of early-stopping method to terminate training. You can vary all the parameters of the network and training to solve these problems.

You shall write a brief report (no more than 7 pages) with the following sections:

1. Overview (describing the contents of the report).
2. Methods (describing what your system does and how it is trained).
3. Experiments (telling all the experiments you performed to figure out how to get a solution for each of the data sets).
4. Experimental results (explaining how well your work solved the problems provided).
5. Conclusion (telling what you learned).