Learning from Data, Tutorial Sheet for week 7

School of Informatics, University of Edinburgh

Instructor: Amos Storkey

1. Given training data $D = \{(x^\mu, y^\mu), \mu = 1, \ldots, P\}$, you decide to fit a regression model $y = mx + c$ to this data. Derive an expression for m and c in terms of D using the minimum sum squared error criterion.

2. The MATLAB code below implements a Radial Basis Function LPM on some artificial training data:

```
x = 0:0.05:1; y = sin(20*x); % training data
alpha = 0.1;  % basis function width
m = -1.5:0.1:1.5; % basis function centres

phi = phi_rbfn(x,alpha,m);
w=(phi*phi'+10^(-7)*eye(size(phi,1)))\sum(repmat(y,size(phi,1),1).*phi,2);

ytrain = w'*phi; xpred = -0.5:0.01:1.5;
ypred=w'*phi_rbfn(xpred,alpha,m); plot(x,y,'rx'); hold on;
plot(xpred,sin(20*xpred)); plot(xpred,ypred,'--');
```

Write a routine `function phi=phi_rbfn(x,alpha,m)` to complete the above code where $\phi_i(x) = \exp(-0.5(x - m_i)^2/\alpha^2)$, and run the completed code. Comment on your predictions.

3. Given training data $D = \{(x^\mu, c^\mu), \mu = 1, \ldots, P\}$, $c^\mu \in \{0, 1\}$, you decide to make a classifier

$$p(c = 1|\mathbf{x}) = \sigma\left(\mathbf{w}^T \boldsymbol{\phi}(\mathbf{x})\right)$$

where $\sigma(x) = e^x/(1 + e^x)$ and $\boldsymbol{\phi}(\mathbf{x})$ is a chosen vector of basis functions.

- Calculate the log likelihood $L(\mathbf{w})$ of the training data.
- Calculate the derivative $\nabla_{\mathbf{w}} L$ and suggest a training algorithm to find the maximum likelihood solution for $\mathbf{w}$.
- Comment on the relationship between this model and logistic regression.
- Comment on the decision boundary of this model.