COMPUTER SCIENCE TRIPOS Part IB – 2015 – Paper 4

1 Artificial Intelligence I (SBH)

We aim to solve a *supervised learning* problem using a simple neural network taking input vectors $\mathbf{x}^T = (x_1, x_2, \ldots, x_n)$ of features and computing the function

$$h(\mathbf{x}, \mathbf{w}) = \sigma\left(w_0 + \sum_{i=1}^{n} w_i x_i\right)$$

where $\mathbf{w}^T = (w_0, w_1, \ldots, w_n)$ is a vector of weights and σ is an activation function. We have a set $\mathbf{s}^T = ((\mathbf{x}_1, y_1), \ldots, (\mathbf{x}_m, y_m))$ of m labelled training examples and seek to minimize the error

$$E(\mathbf{w}) = \sum_{i=1}^{m} (y_i - h(\mathbf{x}_i, \mathbf{w}))^2.$$

(*a*) Derive the gradient descent training algorithm for this problem. [5 marks]

(*b*) We now notice that for the particular problem of interest a solution will only make sense if a specific subset S of the weights is positive. Devise a modified version of the training algorithm that enforces this. [8 marks]

(*c*) A colleague is attempting to solve a heuristic search problem using the $A^\star$ algorithm, but is unable to decide which of a number of heuristics to use. Your colleague has a large collection of test problems, and believes that the best heuristic to use might depend on particular characteristics of the problem being solved. Explain in detail how you might apply machine learning to help your colleague. [7 marks]