

Hebbian Learning and Negative Feedback Neural Networks

Professor Colin Fyfe

Applied Computational Intelligence Research Unit

The University of Paisley, Scotland

colin.fyfe@paisley.ac.uk

The central idea of this presentation is that artificial neural networks which use negative feedback of activation can use simple Hebbian learning to self-organise in such a way that they uncover interesting structure in data sets. The network in its simplest form performs a Principal Component Analysis. Extensions to the network are shown to perform Exploratory Projection Pursuit: they find low-dimensional filters of the data which reveal interesting structure in the data. For example, we might search for outliers from the main body of the data or clusters within the data set and this search can be performed in a hierarchical manner – we find one cluster in the midst of many and then re-project the data from this one cluster to attempt to find sub-clusters. There are two main ways of performing these searches and these are contrasted and compared and a composite method created which exhibits useful properties from the two underlying methods.

The network can also be used to find independent components of a data set in a number of different ways. For example, one extension to the basic network is shown to perform a type of Factor Analysis – it identifies a set of factors which when OR-ed together will construct the data set. Other methods are used to perform Independent Component Analysis which is extensively used in blind source separation – extracting one signal from a linear mixture of signals. The network can also be used for clustering in a topology preserving manner: there are several ways of clustering using this network in such a way that similar data points are clustered close to one another and only similar data points are treated this way.

In part 2 of the talk, twinned networks are introduced: these networks have two input data streams on which they self-organise using simple Hebbian learning with negative feedback again. In their basic form, the networks are shown to perform Canonical Correlation Analysis, the statistical technique which finds those filters onto which projections of the two data streams have greatest correlation. Various extensions of the basic methods are devised in order to create methods which react to more than two data streams at a time or which deal with problems such as multicollinearity. A further extension is the twinning of the Exploratory Projection Pursuit methods from the first part of the book so that the new network identifies shared structure across two data streams. This new network is also shown to perform Independent Component Analysis. A final chapter deviates somewhat from the rest of the book since its emphasis is on an extension of the Principal Curve algorithm so that we now have two curves learning on two data streams simultaneously.

Since the scope of the talk is the development of new algorithms, all algorithms which are derived analytically, are illustrated on artificial data before being used on real data sets. Where it is of interest, the results are compared with those from standard statistics or from alternative artificial neural networks.