Development of Neural Network Techniques for the Analysis of JET ECE Data

D V Bartlett, C M Bishop[1].

JET Joint Undertaking, Abingdon, Oxon, OX14 3EA.
[1] Neural Networks Group, AEA Technology,
Harwell Lab, Abingdon, Oxon. OX11 0RA.

DEVELOPMENT OF NEURAL NETWORK TECHNIQUES FOR THE ANALYSIS OF JET ECE DATA

D V Bartlett, C M Bishop*

JET Joint Undertaking, Abingdon, Oxon. OX14 3EA, UK

** Neural Networks Group, AEA Technology, Harwell Lab, Abingdon, Oxon. OX11 0RA, UK*

ABSTRACT

This paper reports on a project currently in progress to develop neural network techniques for the conversion of JET ECE spectra to electron temperature profiles. The aim is to obtain profiles with reduced measurement uncertainties by incorporating data from the LIDAR Thomson scattering diagnostic in the analysis, while retaining the faster time resolution of the ECE measurements. The properties of neural networks are briefly reviewed, and the reasons for using them in this application are explained. Some preliminary results are presented and the direction of future work is outlined.

(1) INTRODUCTION

Neural networks provide a powerful class of algorithms for non-linear data processing. This makes them well suited to a variety of data analysis tasks in plasma diagnostics. They have already been used successfully in the reconstruction of density profiles from line integral data [1], the fitting of complex impurity emission spectra [2,3] and the extraction of plasma equilibrium parameters from boundary magnetic measurements [4]. In all of these cases the neural networks have been able to achieve very rapid processing of the data, with accuracy comparable to or better than that of the standard analysis techniques.

For the present application, a neural network approach is being used because the network can learn the non-linear transformation required to correct for the combination of systematic errors present in the ECE T_e profiles, without needing any independent formulation of their behaviour. This is discussed further in Section 3.

The LIDAR profiles are used as the "target data" for training the neural network. The network learns the systematic non-linear transformations needed to convert measured ECE spectra into temperature profiles, with various magnetics parameters used as additional inputs. During the training process the network tends to average over the random LIDAR errors. Once trained, the network should therefore produce profiles in which both systematic and random errors are relatively small, using only the ECE spectra and information from the magnetics measurements as input data. The time resolution would then be determined only by the ECE measurement.

The particular goals of the present study are: (i) to examine different possible neural network structures and develop a network based on the most promising of these, and (ii) to assess the suitability of this analysis for routine use at JET.

Following a brief review of the JET ECE and LIDAR diagnostics, the paper explains why neural networks are well suited to this problem, outlines how they operate and are trained, and then describes how they can be applied to the analysis of JET ECE data. The current status of development of this work is described and possible future enhancements and applications are discussed.

(2) THE JET ECE AND LIDAR DIAGNOSTICS

The JET ECE measurement system has been described at a previous conference in this series [5]. The emission along a major radius 0.13 m below the plasma mid-plane is collected by in-vessel antennas and transported by oversized waveguides to the measurement area outside the biological shielding wall. A variety of measurement instruments is in use, but we restrict discussion here to the Michelson interferometer whose spectra have been employed in the present work. The characteristics of this instrument are listed in Table I.

The JET LIDAR diagnostic [6] measures electron temperature profiles by time-of-flight analysis of the Thomson back-scattered light from a very short pulse ruby laser. The laser pulse passes through the plasma along a mid-plane major radius and light which is 180° back-scattered is spectrally analysed in the conventional way to give electron temperature and density profiles. The essential characteristics of the LIDAR diagnostic are also summarized in Table I.

TABLE I: Characteristics of the Two Electron Temperature Profile Diagnostics

	LIDAR	ECE (Michelson)
Radial spatial resolution	≈ 0.10 m	≈ 0.15 m
Temporal resolution	few ns	15 ms
Repetition rate of measurement	0.5 to 1 Hz	60 Hz
Number of profiles per plasma pulse	5 to 10	≈ 300 *

* Limited by data acquisition memory.

The dominant sources of uncertainty in these two temperature measurements are quite different. The purpose of the present work is to exploit these differences to obtain improved determination of the electron temperature profile.

As is well known, the uncertainties in the electron temperature profiles obtained from ECE measurements are dominantly systematic. They arise principally from the spectral calibration of the instrument, which gives an uncertainty on the temperature values, and the limitations on the accuracy of the calculations of the magnetic field profile, which lead to a radial uncertainty via the spectral frequency to radius transformation. The JET magnetic equilibrium code [7] is believed to give generally accurate magnetic field profiles, but it is difficult to estimate how large the errors might be. From an examination of individual LIDAR and ECE profiles, it appears that under some plasma conditions the total field error can correspond to a radial error in the ECE T_e profile location of up to 0.15 m near the plasma centre. The random error in the T_e values, due to noise in the measurement of plasma emission, is generally small.

The uncertainties in the LIDAR temperature profiles are primarily random and due to the statistics of the limited number of photoelectrons in each spectrometer channel. The

systematic uncertainties, from factors such as the calibration of the relative response of the spectrometer channels, are believed to be small.

The estimated values of these uncertainties are summarized for the two diagnostics in Table II.

TABLE II: Sources of Uncertainty in the Two Electron Temperature Measurements

	LIDAR	ECE
Random error in profile location	< 0.05 m	< 0.02 m
Systematic error in profile location	< 0.05 m	up to 0.15 m
Random error in temperature values	Typically ≤ 7% of T_{eo}	≤ 2%
Systematic error in temperature values	≤ 3%	± 10% (absolute), ± 5% (relative)

(3) OVERVIEW OF NEURAL NETWORKS AND THEIR APPLICATION TO ECE

The discussion of neural networks in this paper is directed specifically at the present application. More general overviews can be found in [8] and [9].

Large numbers of simultaneous measurements are available from the two T_e diagnostics, typically four or five profiles suitable for this analysis per JET plasma pulse. By comparing many pairs of simultaneous measurements, it should be possible to average out the statistical fluctuations in the LIDAR data. Similarly, any small differences between individual profiles due to the different integration times of the two measurements (for example if the profile is perturbed by a rotating MHD structure) should also be suppressed by the averaging. However, the difficulty which arises in making these comparisons is that the two sources of systematic error in the ECE temperature profiles cannot be separated. To estimate the ECE calibration errors by a direct comparison with LIDAR requires an accurate knowledge of the magnetic field profile, so that ECE frequencies can be associated with the correct LIDAR spatial points. Alternatively, the total field profile could be estimated by associating the ECE frequencies with the LIDAR radii which have the same temperature, provided that the uncertainties in both temperature measurements are small. However, when both sources of uncertainty are present, the relationship between ECE spectra and LIDAR profiles is complex and non-linear. A neural network can learn the required transformation from ECE spectrum to T_e profile from the available data, without the dependence of the field on the magnetics parameters, or the form of the ECE calibration errors, being specified independently.

The power of the neural network technique arises from a property known as the "universality" of the network which expresses its ability to learn **any** non-linear transformation. This can be translated into the present context as follows. If the LIDAR errors are randomly distributed around the true temperature, if the ECE random errors are small, and if a large number of training examples covering the whole range of the input parameter space

are used, then a suitably constructed and trained network will reproduce from a given ECE measurement the "average" profile which LIDAR would have measured. Hence, it is in principle possible to obtain profiles from the ECE which have low random error, the low systematic error of the LIDAR and the time resolution of the ECE measurements. In practice, we must determine how well the network does learn the transform from the available data, and how the network can be structured to optimize its performance. These questions are discussed further in the later sections.

Many different neural network architectures have been developed in recent years for a variety of applications. In the present case we use a particular class of network known as the multilayer perceptron (MLP). This is a relatively simple class of network, which has nevertheless found application to a vast range of practical problems. It can be regarded as a class of non-linear functions which map a set of input variables to a set of output variables. Any particular transformation is governed by a set of network parameters whose values can be chosen with the aid of a set of examples of the desired mapping. This is the process referred to above as network training.

The structure of a multilayer perceptron is illustrated diagramatically in Figure 1. In this figure, the small circles represent nodes, each of which is associated with a numerical value, and the lines represent the parameters which control the calculation of the node values. The network consists of a number of layers: an input layer, an output layer and one or more hidden layers, each containing a number of nodes. The input layer has a node for each of the input data points in a given measurement, while the output layer has a node for each of the output data values. The optimum selection of the number of hidden layers and the number of nodes in these layers is determined empirically during the network training. For this application, it is anticipated that reasonable results will be obtained using only one or two hidden layers, each having a comparable number of nodes to the input and output layers.

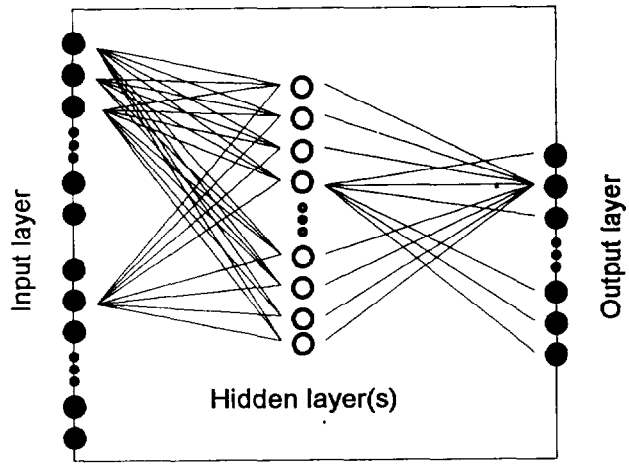

Figure 1: *A sketch illustrating the internal structure of a multilayer perceptron neural network. The small circles represent the network nodes, while the lines are the parameters controlling the calculations which link the nodes. In the present case, each node in the input layer corresponds to an input data value (ECE data point or magnetics parameter) while the output nodes represent the predicted T_e profile values.*

The network processing represented in Figure 1 consists of calculating successively the values at the nodes in each layer (starting at the first hidden layer) by a weighted sum over the values in the immediately preceding layer. Defining the value at node n in layer l as $V_{l,n}$:

$$V_{l,n} = f\left(\sum_{n'}\left(W_{l,n,n'}.V_{l-1,n'}\right) + B_{l,n}\right) \qquad (1)$$

where n' runs over the nodes in layer $l-1$, $W_{l,n,n'}$ are the weights linking each pair of nodes in the two layers, and the $B_{l,n}$ are biases which can be added to each node. The function $f()$, which is applied during the calculation of the hidden node values, is typically chosen to be the *tanh* function, plotted in Figure 2. For the calculation of the output layer values, the function $f()$ is not applied, so that the final layer of weights simply perform a linear transformation. It is the inclusion of this function which gives the network the ability to produce non-linear transformations, rather than simple matrix multiplications. The exact form of the function is not critical, but for applications of the present type it is desirable to have a linear central region with a smooth transition to non-linear (saturation) behaviour at the extremes. The *tanh* function possesses these properties, and has the advantage that during network training its derivatives can be calculated very efficiently [9].

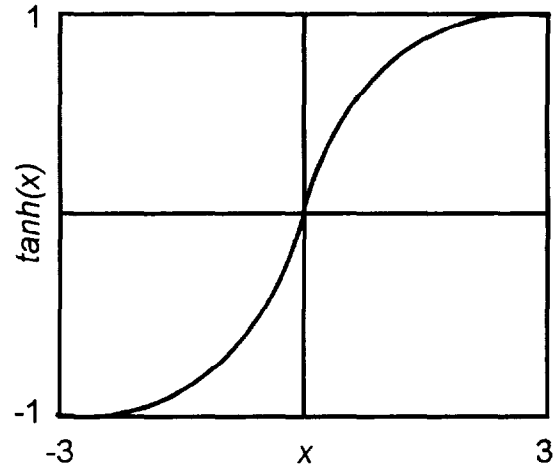

Figure 2: *A sketch of the tanh(x) function. This function is applied to the weighted sums used to calculate the value at each hidden node. It gives the network the ability to generate arbitrary non-linear mappings.*

Network training consists of minimizing a suitably chosen error function such as:

$$E^S = \sum_p \sum_j \kappa_j^p \left\{ T^{Np}(R_j) - T^{Lp}(R_j) \right\}^2 \qquad (2)$$

where the T^{Np} represent the network output temperature values, the T^{Lp} are the LIDAR profile points, each of these at the radii R_j, and the κ are weights which reflect the estimated size of the LIDAR statistical errors at each point. The sums are made over all the radius values (j index) and all the measured profiles in the training data set (p index).

The mathematical technique used to achieve this minimization is based on a procedure known as error backpropagation. This is an efficient algorithm for evaluating the derivatives of the error function with respect to the weights and biases. These derivatives can then be used in one of a variety of standard non-linear optimization schemes, such as gradient descent, conjugate gradients etc. We have chosen to use the limited memory BFGS quasi-Newton method which is particularly well suited to this type of problem [9].

Several aspects of the training process are important in determining the behaviour of the network. The two most crucial of these are:

(i) It is essential that the training data covers a range of the input parameter space at least as large as that for which the network will be used. If the n input parameters are regarded as defining an n-dimensional space, then the training data are points in this space which determine an n-dimensional volume. The network must only be used on data which lies within this

volume. Otherwise the results are unpredictable since it has no basis on which to make its transformation. A simple example of this would be using the network on data taken at a value of the toroidal magnetic field outside the range in the training data set. Since the network would have not learnt to correct the ECE calibration errors at the required frequencies, its output would be meaningless. More subtle problems could arise if the network was used, for example, at a low value of toroidal field (which had been included in the training), but at a plasma current level above those found in the training data set for this field.

(ii) The network will, during training, learn any systematic errors in the training data and they will be transferred to the network outputs when it is used.

(4) IMPLEMENTATION FOR JET ECE AND LIDAR DATA

To train and test the networks being constructed in this first phase of development, approximately 2000 ECE and LIDAR profiles have been selected, from about two months of continuous JET operation. The data have been screened to remove any defective points, and to ensure that the ECE and LIDAR measurements are in all cases simultaneous to within 50 ms. A number of one dimensional magnetics parameters have also been selected. These are: the vacuum toroidal field, the plasma current, the toroidal beta, the plasma inductance, the plasma elongation at centre and edge, the plasma triangularity, the plasma vertical displacement, and the plasma major and minor radii. Some of these quantities are not directly measured, but are outputs from the magnetic equilibrium calculations. This should not affect the analysis, since the network will automatically learn any systematic errors in these quantities during training, and compensate for them when it is used. In principle, the raw magnetics measurements from which these parameters are calculated could be used directly as the training inputs, but the physical quantities are preferred in this first investigation so that some insight into the relative importance of the different parameters can be obtained.

The data set has been divided into three groups. The first group is used for training various network configurations, with different numbers of hidden nodes. The second, validation, set is used to assess the performance of these various configurations, and the third set is used to test the network which is finally selected.

Some pre-processing is required before the input data is fed to the neural network. The purposes of this pre-processing are to eliminate any bias which might arise during network training due to very different numerical values for the various input parameters, and to simplify the transform which the network must learn. The pre-processing applied is as follows:

(i) The magnetics parameters are scaled to zero mean and unit standard deviation, the scale factors being determined from the training data. This prevents the very different numerical values arising from the various units of measurement (eg tesla for magnetic field, amperes for plasma current) from influencing the network.

(ii) The ECE spectra are pre-processed by being converted to approximate T_e profiles within fixed major radius limits using the vacuum magnetic field profile. This simplifies significantly the complexity of the transform which the network must learn. The network will just make a radial shift of the profile points, rather than needing to learn how to select the second harmonic region of the spectrum and then normalize the frequency scale according to the magnetic field. To facilitate comparisons, all the T_e profiles (LIDAR, ECE input to the network and the

network output) are transferred to a fixed grid of 11 points uniformly spaced in major radius between 2.8 m and 3.8 m.

(iii) All T_e profiles are scaled by the ECE profile average of that measurement. This scaling is applied to both the ECE temperatures and, during training, to the LIDAR temperatures, in order to have values ~1. The correct absolute temperatures are recovered by re-scaling the network outputs. This scaling means that the network does not need to learn during training that its transformation should not depend on absolute temperature values (to within a multiplicative constant), that is, that the transformation of any profile is invariant under a linear scaling of temperature values.

The computations reported here are made on IBM RS/6000 workstations which have a peak performance of 50 Mflops. Even with this computing power, and using efficient learning algorithms, many hours of execution time are generally required to train a network. When the network is used however, very little execution time is used since there is only a small number of floating point operations in the forward pass calculation.

(5) PRELIMINARY RESULTS

For the first examination of the performance of neural networks on this dataset, we have omitted the κ weighting factors from Equation (2), and considered only networks having two layers of adaptive weights, that is, a single hidden layer. In order to perform some optimization of the network architecture, networks having from 1 to 20 hidden nodes have been trained using the training dataset. The performance of these networks was then compared using the validation set, and the network with 7 hidden nodes in one layer was found to give the best results. Finally, the performance of the selected network was checked using the test set.

The quantities used to assess the performance of the networks and to compare the results from the network with those from the conventional analysis are rms errors, defined in similar fashion to the error function in Equation (2):

$$E_p^{RMS} = \sqrt{\frac{1}{11}\sum_{j=1}^{11}\left\{T^{Test}(R_j) - T^{Lp}(R_j)\right\}^2} \tag{3}$$

$$E^{RMS} = \sqrt{\frac{1}{P}\sum_{p}\frac{1}{11}\sum_{j=1}^{11}\left\{T^{Test}(R_j) - T^{Lp}(R_j)\right\}^2} \tag{4}$$

where P is the number of instances in the data set, and T^{Test} can represent either the network output T_e profile, or the profile obtained from the conventional analysis (that is, the profile obtained from the ECE spectrum using the magnetic field profile calculated by the equilibrium code). This is certainly not the only possible choice for measures of network performance, but has been adopted for the first analysis because of its simplicity.

For each of the networks which has been trained, the value of E^{RMS} has been calculated for the validation data set, and the network with 7 hidden nodes gives the smallest

result. The corresponding values of Ep^{RMS} have also been evaluated for each instance in the test data set using both the profiles calculated by the 7 hidden node network and the conventional analysis as T^{Test} . Figure 3 shows a scatter plot of these two error measures for all instances in the test set. The oblique line corresponds to equal errors, so that the points which lie above this line represent instances for which the network generated a lower RMS error than the conventional analysis.

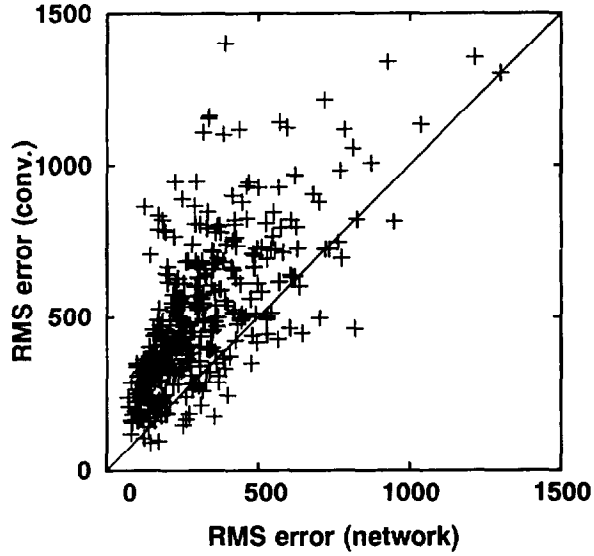

Figure 3: A scatter plot comparing rms errors (as defined in Equation 3) of the conventional analysis and the network output, for the network with 7 hidden units. All the data in the test set have been included. The bias towards values above the 45° line indicates that the network outputs tend to have lower rms errors.

Figure 4a shows a typical example of the profiles generated by the network, together with the corresponding LIDAR profiles and ECE profiles obtained from the conventional analysis. In Figure 4b, the equivalent profiles are shown for the measurement which gave the smallest rms error between network and LIDAR profiles of all the data in the test set.

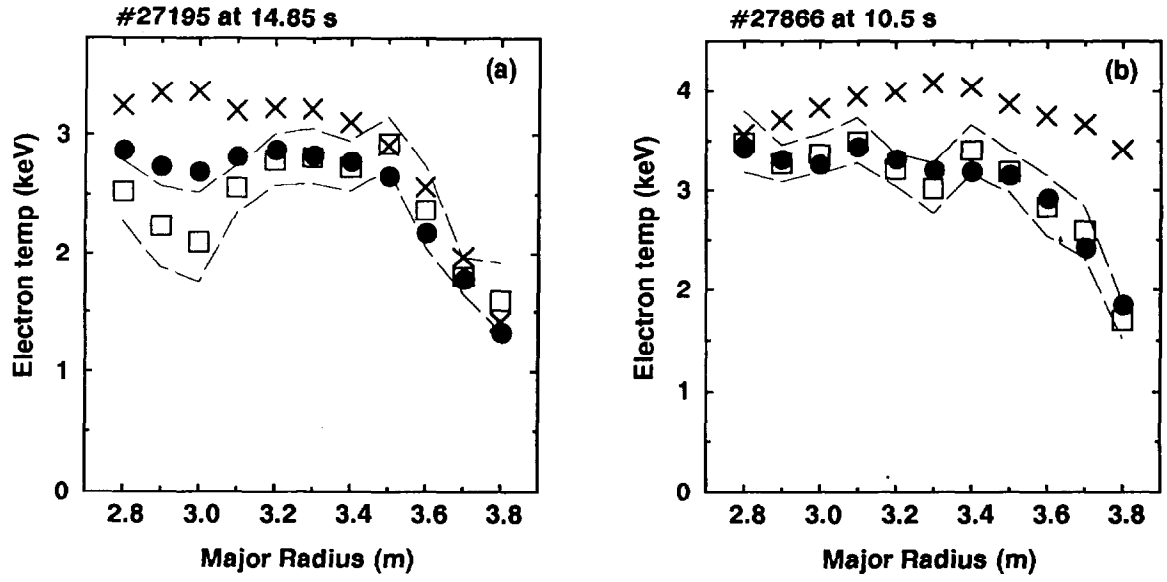

Figure 4: A comparison of T_e profiles from the conventional ECE analysis (crosses), LIDAR T_e profiles (squares) and the network outputs (circles). The estimated statistical error of the LIDAR profiles is shown by the dashed lines. Part (a) is a typical example, and part (b) is the measurement with the lowest rms error of all the data in the test set.

These results are very encouraging and suggest that worthwhile improvements over the conventional analysis may be obtained from the neural network. However, it must be stressed that these are very preliminary results and that many issues remain to be addressed before a firm conclusion can be reached. In particular, the κ factors must be reinstated, and much more careful optimization of the network architecture is needed. Also, it should be noted that when the number of units in the hidden layer is less than the number of output units then the outputs are not independent. An optimized network structure should be able to achieve better results with more hidden units, and so remove this limitation.

(6) POSSIBLE FUTURE DEVELOPMENTS

If the results of this initial investigation are satisfactory then a number of further developments can be foreseen, both short term (to optimize the present application) and long term (possibly for application in Next Step tokamaks).

The most important extension of the present work which is immediately required is a technique for estimating how well the network is performing for any given set of input data. Since the network has been trained to produce outputs which look like T_e profiles, it may not be easy to determine, simply by examining the outputs, whether or not the network is behaving correctly, within its trained parameter space. The question of networks providing such reliability estimates is a general one in neural network research. A variety of techniques is being actively pursued [10], some of which should be applicable to the present analysis.

The network structure described in the preceding sections makes a minimum of assumptions about the physical behaviour of the spectrum-to-profile transformation and the associated errors. This is considered the best choice for the first attempt at constructing a network for two reasons. First, it minimizes the likelihood of placing constraints on the network which may be invalid or unnecessary, and secondly it generates a standard network structure to which many previously developed algorithms can be applied with little or no modification. However, it may be possible to construct a network which can be trained to produce more accurate results if we make use of our physical understanding of the problem. The network would be structured in such a way that it can only produce mappings from ECE spectrum to T_e profile which are consistent with this understanding. In this case the training process should be much more efficient, since the number of mappings which the network is able to generate (and hence the number of incorrect mappings which the training must exclude) has been greatly reduced. A network which separates the correction of the ECE spectral calibration and the magnetic field profile calculation is illustrated in Figure 5. The two halves of the network would operate independently on separate sets of input data (ECE and magnetics), and their outputs would be combined in the same way that spectra are normally converted to profiles. This network structure would however be more complex to train, since the two halves would have to be trained simultaneously using the same set of training data. An advantage of this approach would be that the intermediate outputs, from the two halves of the network, would be accessible. The calculated total field profile would have to be treated with caution since it would have limited accuracy, particularly when the T_e profiles are flat. On the other hand, the deduced ECE calibration correction (the output from the upper half of the network) might be used to determine a more accurate calibration, and also could be a very useful aid in monitoring the stability of the ECE system's spectral response.

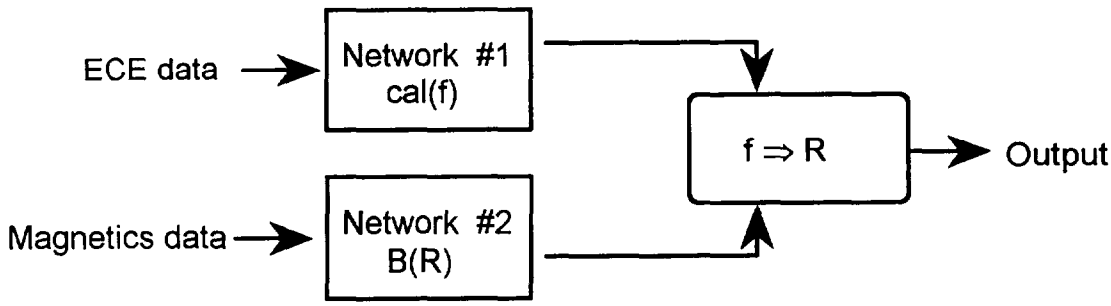

Figure 5: *A possible network structure which reflects our knowledge of the transformation from ECE spectra to T_e profiles. The correction for the ECE calibration errors and the determination of the total magnetic field profile are made separately, and the two results are combined in the same way as for the conventional analysis.*

Finally, there are two issues which may determine the longer term future of this type of analysis for ECE data, in particular its relevance for Next Step tokamaks. For larger and hotter plasmas the difficulties of calculating the plasma internal magnetic fields will become greater. Similarly, the problems of routinely making accurate spectral calibration measurements will also become greater on future machines. Hence, if suitable Thomson scattering measurements are available, then a neural network technique may provide the most accurate interpretation for ECE measurements. Alternatively, neural networks may be useful on both present and future machines because of the speed with which they operate. Dedicated hardware for implementing neural networks already exists, eg [4], and there is potential for very high speed processing for real time applications. There is therefore the possibility that if the network is robust enough and reliable enough that it could be used to process ECE data in real time control applications.

(7) SUMMARY

This paper has shown that neural networks are potentially a valuable tool in the interpretation of ECE data. It may be possible to use them to reduce considerably the systematic uncertainties in ECE derived T_e profiles. This would be achieved by training a network to convert ECE spectra to profiles, using LIDAR T_e profiles as the target. The result should be T_e profiles with the low systematic uncertainty of the LIDAR measurement, the low random uncertainty of the ECE measurement, but the same time resolution as the ECE data. The only other input data required are various one-dimensional magnetics parameters.

The results obtained so far, while still preliminary, indicate that even a simple network structure with limited training, can produce profiles which are closer on average to the LIDAR profiles than the conventional analysis. It is anticipated that optimization of the network structure and more extensive training will improve on these results.

In the future, it is planned to investigate network structures which are more closely adapted to this problem. It is hoped that this will yield networks which can be more efficiently trained to give more accurate results.

(8) ACKNOWLEDGEMENTS

The authors wish to thank C Gowers and P Nielsen for helpful discussions about the nature of the LIDAR errors, P Stott for originally suggesting the investigation of this topic, and D Holden for his contribution to the development of the neural network software.

(9) REFERENCES

[1] C M Bishop, I G D Strachan, J O'Rourke, G P Maddison and P R Thomas, (1993) "Reconstruction of tokamak density profiles using feedforward networks" Neural Computing and Applications (Springer Verlag) Vol. 1, No. 1 pp 4 - 16.

[2] C M Bishop, C M Roach and M von Hellerman (1992) "Automatic Analysis of JET Charge Exchange Spectra using Neural Networks", submitted to Nuclear Fusion.

[3] C M Bishop and C M Roach "Fast curve fitting using neural networks" Rev. Sci. Instrum. **63** (10) 4450 (1992).

[4] C M Bishop, P Cox, P Haynes, C M Roach, M E U Smith, T N Todd and D L Trotman "A neural network approach to tokamak equilibrium control" in Neural Network Applications, (Springer Verlag) Ed. J G Taylor, pp 114 - 128 (1992) .

[5] D V Bartlett, D J Campbell, A E Costley, S Kissel, N Lopes Cardozo, C W Gowers, S Nowak, T Oyevaar, N Salmon, B Tubbing "Overview of JET ECE measurements", Proceedings of EC-6, the 6th Joint Workshop on ECE and ECRH (Culham Laboratory Report CLM-ECR 1987).

[6] H Salzmann, J Bundgaard, A Gadd, C Gowers, K B Hansen, K Hirsch, P Nielsen, K Reed, C Schrodter, and K Weisberg "The LIDAR Thomson scattering system on JET" Rev. Sci. Instrum. **59**, 1451 (1988)

[7] E Lazzaro and P Mantica "Experimental identification of tokamak equilibrium using magnetic and diamagnetic signals" Plasma Physics and Contr. Fusion **30** 1735 (1988)

[8] J Hertz, A Krogh and R G Palmer (1991) "Introduction to the theory of neural networks" Addison Wesley.

[9] C M Bishop (1993) "Neural networks and their applications", invited review article, accepted for publication in Reviews of Scientific Instruments.

[10] C M Bishop (1992) "Validation of neural network solutions: an illustration from multiphase flow monitoring", submitted to Neural Computation.