

Reconstruction of Tokamak Density Profiles Using Feedforward Networks

Chris Bishop,¹ Iain Strachan,¹ John O'Rourke,² Geoff Maddison³ and Paul Thomas²

¹AEA Technology, Harwell Laboratory, Oxon OX11 0RA, UK, ²JET Joint European Torus, Oxon OX14 3EA, UK, ³AEA Technology, Culham Laboratory, Oxon OX14 3DB, UK.

The tokamak is currently the principal magnetic confinement system for controlled fusion research. In seeking to understand the physics of the high temperature plasma inside the tokamak, it is important to have detailed information on the spatial distribution of electron density. One technique for density measurement uses laser interferometry, which gives line-integral information along chords through the plasma. This requires an inversion procedure to extract spatially local density information. In this paper we make use of feedforward networks to extract local density profiles from the line-integral data obtained from the multichannel interferometer on the JET (Joint European Torus) tokamak. An important feature of our approach is the use of profile data from a second high resolution diagnostic system, called LIDAR, to train the network. The LIDAR system provides data at high spatial resolution but with a low repetition rate, and therefore has a complementary rôle to interferometry which operates at a high sampling rate but with much lower spatial resolution. Results show that the neural network is able to extract significantly more detailed profile information than the conventional Abel inversion method currently used on JET.

Keywords: Neural network; Ill-posed; Inverse problem; Principal components analysis; Interferometer

1. Introduction

Measurements of the spatial distribution of electron density in a high temperature tokamak plasma experiment require non-invasive diagnostic techniques. One approach is to use laser interferometry to obtain line-integral measurements of density along a number of chords through the plasma. The extraction of spatially local density information from this line-integral data is usually performed by generalised Abel inversion. In this paper we describe a novel approach to this problem using feedforward neural networks. A significant feature of our approach is the use of profile information obtained from a second diagnostic, known as LIDAR, to train the network. The LIDAR system provides data at high spatial resolution but with a low repetition rate, and therefore has a complementary rôle to interferometry which operates at a high sampling rate but with much lower spatial resolution. As we shall show, the neural network approach allows significantly more detailed density profiles to be obtained from the interferometer data than is possible with conventional methods.

Many plasma phenomena can occur on timescales of microseconds, whereas a plasma pulse can last for many seconds. Since modern tokamak experiments have numerous plasma diagnostics, many of which are capable of high sampling rates, the volumes of data which can potentially be generated, are very large. Furthermore, there is often only a few minutes available between plasma pulses (or 'shots') in which to analyse the data if it is to be available in time to influence the choice of operating parameters for the next pulse. Neural

Original manuscript received 15 January 1992

Correspondence and offprint requests to: Dr C. M. Bishop, AEA Technology, Harwell Laboratory, Oxon OX11 0RA, UK.

network techniques offer a number of features, including the capability to perform high speed non-linear transformations, which make them of considerable interest for data analysis purposes in the context of tokamak physics. Other tokamak applications of neural networks, for data analysis and control, can be found in [1–3].

A brief overview of the tokamak magnetic confinement system is given in Sect. 2, and is followed by a description of the LIDAR and multichannel interferometer diagnostics on the JET tokamak in Sect. 3. The general problem of profile reconstruction from interferometer line-integral data is discussed in Sect. 4. In Sect. 5 we describe the neural network approach to profile reconstruction and present results obtained from JET data. Finally, a summary and discussion are given in Sect. 6.

2. An Overview of the Tokamak

The tokamak is currently the favoured experimental system for research into the magnetic confinement approach to controlled fusion. The photograph shows the Joint European Torus (JET) experiment, currently the world's largest tokamak, at Culham Laboratory in Oxfordshire. A tokamak (from the Russian for 'toroidal magnetic chamber') consists of a toroidal vacuum vessel in to which a small quantity of gas (usually isotopes of hydrogen) is introduced. The gas is ionised and raised to a very high temperature, typically a few 10^7K , by a large toroidal electric current (up to 7×10^6 amps in JET) which is induced by transformer action from a set of primary poloidal field coils. The plasma current also interacts with its own magnetic field in such a way as to confine the high pressure of the plasma. Additional magnetic fields, produced by currents flowing through external poloidal and toroidal field coils, act to stabilise and shape the plasma. The superposition of the various magnetic fields leads to a complex field structure, with individual field lines forming spirals which live on closed nested toroidal flux surfaces. Figure 1 shows a schematic illustration of the flux surfaces in a typical tokamak.

In seeking to understand the physics of the magnetically confined plasma, it is of considerable interest to be able to measure the spatial distributions of plasma density and temperature, and to investigate how these distributions evolve with time. In this paper we are concerned with the determination of electron density profiles. To a good approximation it can be assumed that the tokamak has rotational symmetry around the central vertical axis (the Z-axis in Fig. 1). It is therefore sufficient to consider




Fig. 1. Schematic illustration of a tokamak vacuum vessel, shown in cross-section, together with the nested toroidal flux surfaces. In many instances these flux surfaces are also the surfaces of constant electron density.

the density distribution in a cross-sectional plane at one toroidal location. The contours of constant density generally form closed nested surfaces, with the density going to zero at the walls of the vacuum vessel. The spatial distribution of density is usually expressed in terms of a density profile as a function of major radius R (see Fig. 1), along the toroidal mid-plane. This density profile then specifies the density at all points within the tokamak, provided the surfaces of constant density are known. In many circumstances, the surfaces of constant density coincide with the flux surfaces. This fact is used explicitly in the conventional methods for profile reconstruction. Note, however, that in some situations (in particular, when the plasma is strongly rotating) the density surfaces and flux surfaces need not coincide. Unlike conventional approaches, our neural network method does not require the density to be constant on the magnetic surfaces.

3. Density Profile Measurements on JET

The two principal diagnostics on JET for measurement of the electron density profile are the multichannel interferometer and the LIDAR Thomson scattering system. Here we give a brief overview of the main characteristics of these two diagnostics, with emphasis on those features which are of particular relevance to the current study.

The JET far infrared interferometer can be seen on the left side of Fig. 2, and a detailed description of the system can be found in [4]. The instrument is based on a Mach-Zender interferometer driven by a $195 \mu\text{m}$ DCN laser. Six vertical lines of sight, shown schematically in Fig. 3, provide phase shift measurements to $1/20$ of a fringe, permitting line-integral densities in the range 5×10^{17} to 10^{21}m^{-2} to be measured. (The dashed contours in Fig. 3 represent the magnetic flux surfaces). The spatial




Fig. 2. JET tokamak showing the multichannel interferometer (the tall box-section structure on the left).

resolution of the interferometer is, at best, of order 20 cm, while the sampling rate is currently around 10^4 Hz, although the interferometer itself is capable of much finer time resolution.

To determine the density profile from the six line-integral measurements, an inversion procedure is needed. This will typically take into account information on the geometry of the flux surfaces obtained from other diagnostics, with the assumption that the density is constant on a flux surface. In this study we shall compare the results of a neural network inversion technique with those obtained routinely on JET using a conventional Abel inversion procedure.

The LIDAR (light detection and ranging) diagnostic is described in detail in [5]. It uses a ruby laser to obtain density profile information across the mid




Fig. 3. Schematic illustration of the six vertical chords of the multichannel interferometer on JET, together with the LIDAR line-of-sight. The dashed contours represent the flux surfaces.

plane, as shown schematically in Fig. 3. Thompson scattering of the laser light, together with time of flight measurements, allows the electron density profile to be measured with a spatial resolution of order 9 cm. The laser repetition rate is typically 1–2 Hz.

Thus, the two density profile diagnostics play complementary rôles. The interferometer has low spatial resolution (> 20 cm) with a high repetition rate (10^4 Hz), while the LIDAR gives relatively high spatial resolution (9 cm) but at the expense of a very low repetition rate (1–2 Hz). As we shall show, our neural network approach allows the high resolution of the LIDAR profiles to be exploited in the interpretation of interferometer data, to give an improved reconstruction of the electron density profile. In addition, the trained neural network provides a non-linear procedure for extracting the density profiles which is sufficiently fast to allow inter-shot analysis of the large volumes of interferometer data.

4. Interferometer Profile Reconstruction

In this section we discuss the general problem of extracting an electron density profile from the line-integral data generated by the JET interferometer. We also outline the conventional Abel inversion technique traditionally used on JET. A detailed discussion of our neural network approach, however, is deferred to Sect. 5.

Reconstruction of a density profile from a set of line-integral measurements represents a form of inverse problem. The forward problem, illustrated in Fig. 4a, involves the calculation of the expected

values of the line-integrals, given both the density profile and the contours of constant density (taken to be the flux surfaces). Since this involves simply the evaluation of integrals in a geometrical problem, it is well defined and gives a unique result. However, the 'inverse' problem, illustrated in Fig. 4b, requires the determination of the density profile from the line-integral data, and is an example of an ill-posed problem. (Note that this is not quite the inverse of the forward problem, since in both cases it is assumed that the geometry of the flux surfaces is known.) A problem is said to be ill-posed [6,7] when one or more of the following conditions is not met: (i) there exists a solution; (ii) the solution is unique; and (iii) the solution is stable (so that small variations in the data produce small changes in the solution). Ill-posed problems occur in many areas of physics and engineering, and have been extensively studied. The problem of finding the density profile from a finite number of line-integral measurements is ill-posed, since there are infinitely many density profiles which give rise to the same values for the line-integrals.

There are two basic approaches to the solution of ill-posed problems. The first, known as structural stabilisation, involves limiting the number of degrees of freedom in the function being fitted. In the case of density profile reconstruction, this would require fitting a given functional form for the density profile which contains a limited number of adjustable parameters. Thus the density profile might, for instance, be represented by a finite order polynomial, with the various coefficients being determined by a least-squares fitting procedure. The principal disadvantage of such an approach is that a specific functional form must be selected, and the choice made may not be optimal for the particular problem.

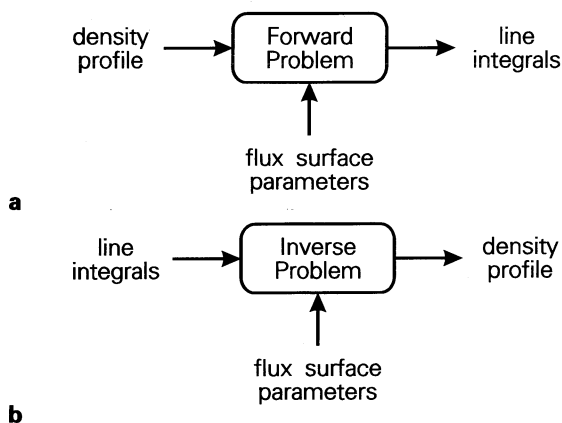


Fig. 4. **a** Illustration of the forward problem of calculating the expected values for the density line-integrals for a given density profile; **b** illustration of the 'inverse' problem of extracting a density profile from a set of line-integral measurements.

The second approach to the solution of ill-posed problems involves the use of regularisation terms to bias the form of the solution, and thereby render the solution unique. The theory of regularisation methods is described in [8]. They have been applied to the problem of neural network learning in [9,10]. The regularisation terms correspond to prior knowledge concerning the expected form of the solutions. For instance, if the solution is expected to be smooth then the regularisation term may be taken to be the mean square curvature of the function. In the context of density profile reconstruction, the use of LIDAR profiles represents prior knowledge concerning the expected form of the density profiles, although this information will be used in Sect. 5 to provide target outputs for network training rather than in the form of a regularisation term.

The conventional method used on JET to reconstruct the density profiles, known as generalised Abel inversion, adopts the former approach of limiting the number of degrees of freedom in an assumed functional form for the density profile. An indication of the appropriate number of degrees of freedom can be obtained as follows. For the JET interferometer, there are six independent line-integral measurements. If information on the shape of the flux surfaces is also available, and we assume that the density is constant on a flux surface, then this effectively doubles the number of independent items of information to 12. This can be seen intuitively from the fact that knowledge of the density at one point on the mid-plane gives the value of the density at a second point on the mid-plane, where the two points lie on the same constant density surface. Furthermore, we know that the density should vanish at the walls of the tokamak, and this provides two further independent pieces of information. Thus there are in total 14 independent items of information available from the line-integral measurements, plus knowledge of the constant density surfaces. This is sufficient to determine the values of the parameters in a 14 parameter functional form for the density profile. This does not, however, take account of the presence of noise on the data. To ensure the results are not strongly sensitive to noise, it is necessary to restrict further the number of degrees of freedom. On JET, the line-integral of density along a vertical chord is represented as polynomial function of major radius R (where R is defined in Fig. 1), and it is usual to consider 6–8 degrees of freedom, (where the number of degrees of freedom is given by the number of polynomial coefficients to be determined from the data). The values of the coefficients are then found by a linear matrix inversion algorithm.

The neural network approach set out below also makes use of functions having a restricted number of degrees of freedom to ensure that the profile reconstruction is well-posed, although the functions are defined in a quite different way. In Sect. 5 we shall compare the results from the neural network approach directly with those obtained using Abel inversion.

5. Neural Network Approach to Profile Reconstruction

The neural network approach to profile reconstruction described in this paper can be broken down into two distinct parts. The first consists of the determination of an appropriate reduced-dimensionality representation for the profiles. As with the Abel inversion method, the approach used in this paper also involves the use of a functional form for the density profile having a restricted number of degrees of freedom. However, unlike the polynomial used in Abel inversion, no explicit choice of functional form for the profiles will be made. Instead, we shall use the high-resolution LIDAR profile information to generate an appropriate functional form. This will involve finding a reduced dimensionality representation of the LIDAR profiles using principal component analysis, and will be discussed in Sect. 5.2. To facilitate comparison with the profiles obtained by Abel inversion (which uses polynomials of order 6–8), we shall consider eight degree-of-freedom representations.

The second stage of the neural network approach then makes use of networks of the multilayer perceptron type to map the interferometer data (together with information concerning the geometry of the flux surfaces) onto the reduced dimensionality representation of the density profile. Combining the two parts together we arrive at the equivalent of a single network which maps the line-integral and flux surface data directly onto a density profile. The second stage will be discussed in more detail in Sect. 5.3. We begin, however, with a brief description of the dataset used to train and test the network.

5.1. Training Data

A dataset for training and testing the networks was assembled from the JET database. Each entry in the dataset corresponds to a particular time on a particular shot, and includes the values of the six line-integrals, together with the values for the LIDAR density profiles which are stored in the

dataset at 35 equally spaced values of major radius R . (Although the spacing of the data point is approximately 5 cm, the true resolution of the LIDAR system is around 9 cm.)

As well as the values of the line-integrals of density from the interferometer, the network must also be provided with information about the geometry of the constant density contours. Unfortunately, this information is not available directly. We do know, however, that the density is often constant on the magnetic flux surfaces, and even when it is not, there is expected to be a significant correlation between the flux surfaces and the constant density surfaces. Information on the geometry of the flux surfaces can be obtained from the large number of magnetic field measurements taken around the perimeter of the vacuum vessel. For the purposes of this study we have made use of the output produced by the FAST code [11], which provides the values of a number of parameters (such as major radius, elongation, triangularity, and so on), known as Lao–Hirshman coefficients, based on measurements of the magnetic fields at a number of points around the perimeter of the vacuum vessel. This code has been chosen because it is sufficiently fast that the information will be available for all required time slices on an inter-shot basis, allowing the network to be used for inter-shot profile reconstruction.

The data points were selected so as to give a wide range of profile shapes, covering both broad and centrally peaked LIDAR profiles. Also included in the dataset were the density profiles obtained using the standard Abel inversion technique employed on JET, for the corresponding time slices. These provided a point of comparison for the neural network results. The dataset contains a total of 1777 entries each corresponding to a particular shot and time, and was divided randomly into a training set of 893 points and a test set of 884 points. Abel profiles were available for 620 of the training points and 618 of the test points. Considerable care and effort was expended in setting up the dataset. In particular, the interferometer traces were examined individually so as to discard those which suffer from the problem of ‘fringe jumps’, which correspond to an error of 2π in the phase measurement for a particular chord, and which lead to significant errors in the density line-integral measurements.

An important point concerns the normalisation of the profile data. We first note that the profile reconstruction problem is essentially geometrical in nature. If all line-integral values are scaled by a constant factor, then the reconstructed profile should also be scaled by that same factor, and any reconstruction method should reflect this exact

linear invariance property. In the case of the neural network approach, we avoid having to train the network to reconstruct self similar profiles of different amplitude by normalising both the line-integral data and the corresponding LIDAR target data. In effect, we are using the network to solve the non-linear problem of determining the shape of the density profile, with the overall scale recovered by subsequent post-processing. Thus, for any given training pattern p , the line-integrals λ_k^p (where $k = 1, \dots, 6$ labels the interferometer channels) are normalised to give a vector of unit length with components:

$$\bar{\lambda}_k^p = \lambda_k^p / \bar{\lambda}^p \quad (1)$$

where

$$\bar{\lambda}^p = \left\{ \sum_{k=1}^6 (\lambda_k^p)^2 \right\}^{1/2} \quad (2)$$

Similarly, the LIDAR density values are normalised to give:

$$\bar{n}_i^p = n_i^p / \bar{\lambda}^p \quad (3)$$

where $i = 1, \dots, 35$ labels the values of major radius R at which the LIDAR density values are specified. The normalised \bar{n}_i^p will be used to provide target values for training the network. When testing the network on new data, the line-integrals are similarly normalised, and the network outputs are post-processed by multiplying by $\bar{\lambda}^p$ to recover the physical density values. As will be discussed in more detail later, this normalisation effectively uses up a degree of freedom, and so within the network architecture we shall be seeking seven degree-of-freedom representations, in keeping with the desire to have eight degrees of freedom overall.

The Lao-Hirshman coefficients are normalised to zero mean and standard deviation = $\frac{1}{3}$ (with respect to the training set), thereby ensuring that all network inputs are 0(1).

5.2. Dimensionality Reduction

The first stage of the neural network approach to profile reconstruction involves finding a representation for the density profiles which has a restricted number of degrees of freedom (i.e. adjustable parameters). To achieve this we shall make use of the LIDAR profiles from the training part of the dataset. We shall consider both the conventional technique of principal component analysis (PCA), and a non-linear neural network approach.

Principal component analysis makes use of the covariance matrix for the LIDAR training data, defined by:

$$C_{ij} = \sum_{p=1}^P (\bar{n}_i^p - \bar{n}_i)(\bar{n}_j^p - \bar{n}_j) \quad (4)$$

where \bar{n}_i^p denotes the (normalised) LIDAR electron density at spatial point i ($i = 1, \dots, N$, where $N = 35$) for pattern p , P is the total number of patterns, and \bar{n}_i is defined by:

$$\bar{n}_i = \frac{1}{P} \sum_{p=1}^P \bar{n}_i^p \quad (5)$$

where the sum runs over all patterns p in the training set. The principal components are given by the eigenvectors of the matrix C :

$$\sum_{j=1}^N C_{ij} u_j^{(m)} = \mu^{(m)} u_i^{(m)} \quad m = 1, \dots, N \quad (6)$$

Since C is, by construction, a real symmetric matrix, it has eigenvectors which form an orthonormal, complete set:

$$\sum_{i=1}^N u_i^{(m)} u_i^{(n)} = \delta_{mn} \quad (7)$$

$$\sum_{m=1}^N u_i^{(m)} u_j^{(m)} = \delta_{ij} \quad (8)$$

where δ is the Kronecker symbol. A reduced dimensionality representation of the density profiles is then obtained simply by retaining the $M < N$ eigenvectors corresponding to the M largest eigenvalues. Any new density profile \bar{n}_i^p can then be projected down onto the M -dimensional subspace using:

$$\eta_m^p = \sum_{i=1}^N (\bar{n}_i^p - \bar{n}_i) u_i^{(m)} \quad m = 1, \dots, M \quad (9)$$

to give a set of M coefficients η_m^p . These coefficients represent an approximation to \bar{n}_i^p obtained by projecting back into the N -dimensional profile space using:

$$\hat{n}_i^p = \bar{n}_i + \sum_{m=1}^M \eta_m^p u_i^{(m)} \quad i = 1, \dots, N \quad (10)$$

If $M = N$ then the original profile is recovered exactly. This can be seen by substituting Eq. (9) into Eq. (10), and using Eq. (8) to give $\hat{n}_i^p = \bar{n}_i^p \forall i, p$. However, if $M < N$ then there is some error involved in the mapping down to M -dimensions and back to N -dimensions. Principal component analysis is designed so that the mean square error over the whole dataset is as small as possible, subject to the limitations of using a linear procedure. We can define a reconstruction error for a particular pattern p by:

$$\bar{E}_{\text{RMS}}^p = \left\{ \frac{1}{35} \sum_{i=1}^{35} (\bar{n}_i^p - \hat{n}_i^p)^2 \right\}^{1/2} \quad (11)$$




Fig. 5. Compression error for LIDAR profiles as a function of the number of retained principal component eigenvectors.

A plot of this error (for the LIDAR profiles in the training set), resulting from projection onto successively increasing numbers of eigenvectors, is shown in Fig. 5. For a seven degree-of-freedom representation of the normalised density profiles we should retain seven principal components. When the reconstructed 35-dimensional profile, obtained from Eq. (10), is post multiplied by the normalisation factor λ^p , we recover the full eight degree-of-freedom representation.

Neural networks offer a non-linear dimensionality reduction technique [12] which includes principal component analysis as a special case. Consider the 5-layer network architecture shown in Fig. 6. The input and output layers contain the same number N of units, while the second hidden layer contains




Fig. 6. Compression network used for non-linear dimensionality reduction.

$M < N$ units. The first and third hidden layers are often chosen to have the same number of units, although this is not essential. Training data for the network consists of the (training data set) LIDAR profiles, and the network is trained to map the density profiles onto themselves. This mapping is unlikely to be exact, however, since there are fewer units in the second hidden layer than there are input units. The network therefore seeks a reduced dimensionality representation which minimises the profile reconstruction error. The first hidden layer acts to map the N -dimensional input vector on to the M -dimensional representation of the second hidden layer, in a manner analogous to Eq. (9). Similarly, the third hidden layer maps from the M -dimensional representation of the second hidden layer back into the N -dimensional profile space, in an analogous way to Eq. (10). Unlike principal component analysis, however, these mappings can be non-linear. In fact, using the universality theorem for single-hidden-layer networks given in [13], these mappings can approximate any non-linear mapping of the same dimensionality to arbitrary accuracy, provided there are sufficiently many non-linear hidden units. Furthermore, it can be shown [14] that if the hidden units in the network operate on the linear part of the sigmoidal curve, then the subspace spanned by the activations of the M middle layer hidden units corresponds to the space of the first M principal component eigenvectors. In this sense, the neural approach contains principal component analysis as a special case.

Networks having five layers with 35 inputs and 35 (linear) outputs, and various numbers of units in the first and third hidden layers (always taken to be equal), were trained using the normalised LIDAR profiles from the training dataset. In all cases there were seven units in the second hidden layer to force the network to find an optimal seven degree-of-freedom representation for the profiles. All network training described in this paper was performed using the BFGS memoryless quasi-Newton algorithm [15], which is found to be much faster and more robust than gradient descent based methods. The best results, i.e. the smallest error with respect to the test set, were obtained for a network having ten units in each of the first and third hidden layers, and are compared with the corresponding results from PCA in Table 1. The root-mean-square (RMS) errors of the network outputs for the trained network are calculated separately for the training and test sets, using the following definition:

$$\tilde{E}_{\text{RMS}} = \left\{ \frac{1}{35 P} \sum_{p=1}^P \sum_{i=1}^{35} (\tilde{n}_i^p - \tilde{i}_i^p)^2 \right\}^{1/2} \quad (12)$$

Table 1. Comparison of RMS errors due to dimensionality reduction by both neural network and PCA methods, for the training and test sets.

Method	RMS (train) $\times 10^3$	RMS (test) $\times 10^3$
Neural network	4.44	4.67
PCA	4.56	4.67

where \tilde{t}_i^p denotes the i th component of the normalised target vector for pattern p . Note that normalised output data is used in defining Eq. (12) to ensure that the results are not dominated by a small number of high density shots.

It is clear that there is no significant improvement resulting from the non-linearity offered by the neural network approach to dimensionality reduction, and so from now on we shall use the PCA results. This has a number of practical advantages resulting from the linear nature of principal component analysis.

5.3. Network Mapping

So far we have set up a 7-dimensional representation of the density profiles given by Eq. (10) with $M = 7$. The profiles themselves are described by the values of electron density at 35 points along the major radius, and can therefore be considered as points in a 35-dimensional ‘profile space’. In our representation, each profile will be described by the values of the seven parameters η_m . The additional global rescaling of the profiles resulting from the normalisation used in Eq. (3) then leads to profiles having eight degrees of freedom. The profiles are now represented by a general linear expansion of the same form as Eq. (10) replace with the profiles now occupy an 8-dimensional hyperplanar sub-space of the 35-dimensional profile space.

The second step in the neural network approach to profile reconstruction involves setting up a mapping which will generate the best example of the reduced dimensionality representation corresponding to particular values of the line-integrals and flux surface parameters. In a conventional analysis (such as the Abel inversion method) this would typically be done using a least-squares calculation to select values for the adjustable parameters such as to minimise the mean square error between the measured line-integrals and those predicted by the reconstructed profile (using the forward calculation illustrated in Fig. 4a). With the neural network approach, however, we make use of the training dataset to determine a direct mapping from the interferometer and flux surface data onto

the reduced-dimensionality representation.

The complete neural network system is shown schematically in Fig. 7. The network has 13 inputs corresponding to the six line-integrals and the seven Lao–Hirshman coefficients, and 35 linear outputs corresponding to the values of (normalised) electron density at the same 35 radial points as used to represent the LIDAR profiles. The activations of the seven units in the layer labelled ‘principal components’ correspond to the parameters η_m^p of Eq. (10). The first hidden layer permits a general mapping onto the principal directions (again using the universality theorem for networks having a single non-linear hidden layer, described in [13]). Thus the only adaptive weights in the network are now those either side of the first hidden layer.

It is important to note that a feedforward neural network can only generate one–one or many–one mappings (functional mappings), since for any given set of input values the output values are uniquely defined. The forward problem illustrated in Fig. 4a is clearly a many–one mapping, since the line-integrals can be calculated uniquely once the profile is specified, but many profiles (specified by a 35-dimensional vector) can give rise to the same line-integral values. The inverse problem of Fig. 4b is therefore potentially a one–many mapping, and we cannot expect to train a neural network to perform this mapping for arbitrary profiles. We are not, however, interested in generating all possible 35-dimensional output vectors, but only those which represent realistic tokamak density profiles. By representing the normalised profiles by just seven principal component eigenvector coefficients, there are now more degrees of freedom in the input space than there are in the output space. In addition,

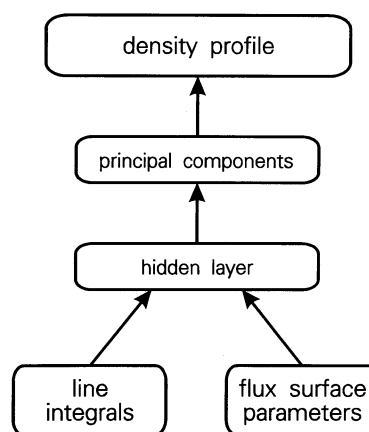


Fig. 7. Neural network architecture for profile reconstruction using principal components to define the reduced dimensionality representation of the density profiles.

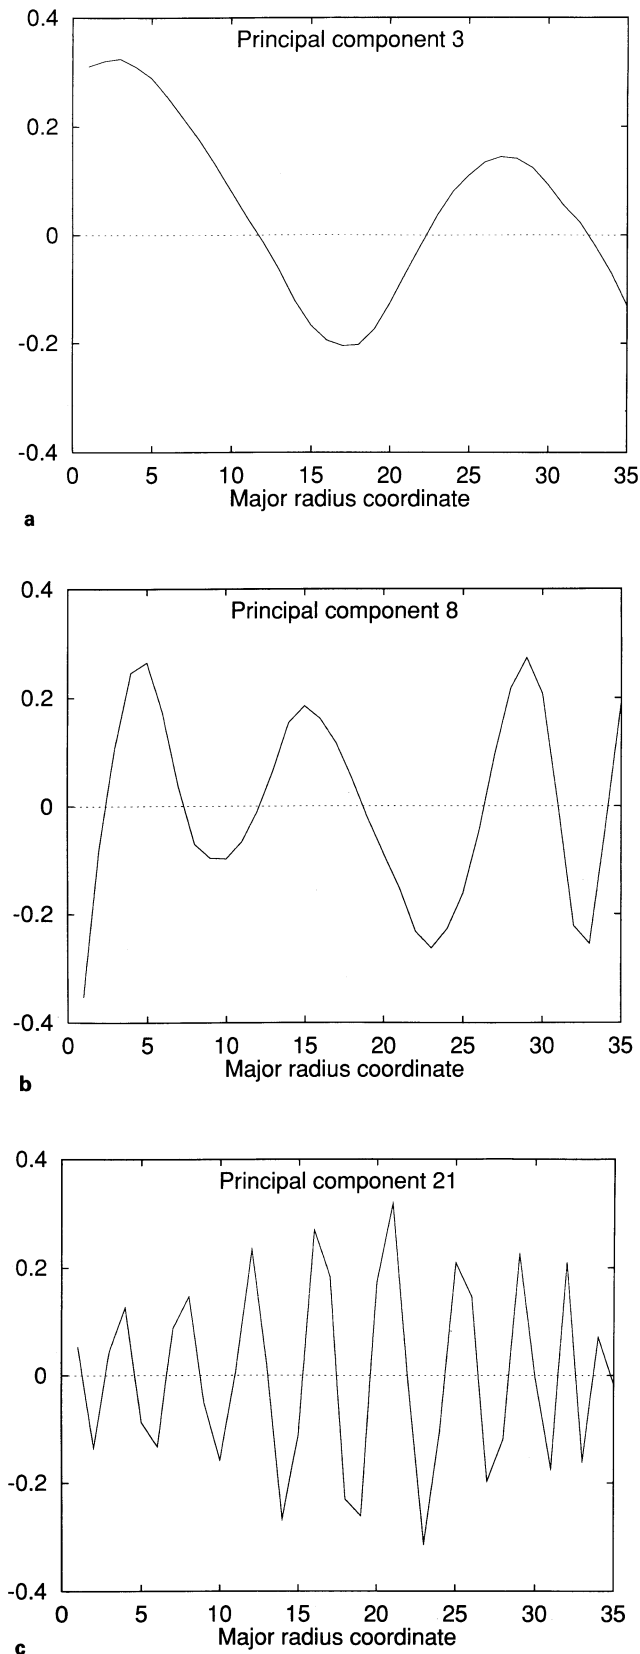


Fig. 8. Plots of the principal component eigenvectors corresponding to eigenvalues **a** 3, **b** 8 and **c** 21, showing the steadily increasing spatial frequency.

we note that the successive principal component eigenvectors, corresponding to increasing eigenvalues, have steadily increasing spatial frequency. This is illustrated in Figs. 8a–c by the eigenvectors corresponding to the 3rd, 8th and 21st eigenvalues. Our approach of using the first seven eigenvectors as the basis representation for the density profiles is therefore consistent, since we expect the line-integral information to determine the low spatial frequencies in the profiles. Thus we expect that the mapping to be learned by the network can be closely approximated by a functional mapping.

Instead of defining the profile targets directly in the 35-unit output layer and backpropagating errors to the principal components layer, it is computationally more efficient to generate a modified training data set by mapping the LIDAR training profiles onto the principal vectors using Eq. (9). The resulting parameter values act as targets for training what is now a single-hidden-layer network having seven output units. The network is then trained by minimising a mean-square error defined with respect to the principal component coefficients. These two approaches are formally equivalent, as is demonstrated in the Appendix.

5.4. Results

Results for the case of networks using seven principal components are given, for various numbers of units in the first hidden layer, in Table 2. Here the RMS errors are again defined, for both training and test sets, by Eq. (12). Also shown in Table 2 are the corresponding RMS errors between the Abel profiles and the LIDAR profiles. These results are also plotted in Fig. 9.

Table 2 (Fig. 9) exhibits the classic form showing a reduction in the training error as the number of hidden units is increased, while the test error decreases at first and reaches some optimum value, and then increases as the number of hidden units is increased further (corresponding to over-fitting).

We see from Table 2 that the RMS error between the network profiles and the corresponding LIDAR profiles for the test data (i.e. data not used in the training of the network) is smaller than the corresponding error between the Abel profiles and the LIDAR. This implies that the network is producing reconstructed profiles which are closer to LIDAR than are the corresponding Abel profiles. As we shall see later, the neural network approach is able to extract detailed features from the interferometer data which are also seen in the LIDAR profiles, but which are not seen in the Abel results.

Table 2. Results for neural network reconstruction of density profiles for various numbers of hidden units showing the RMS errors between the network profiles and those from LIDAR, together with the corresponding RMS errors between the Abel and LIDAR profiles

N_{hidden}	RMS (train) $\times 10^2$	RMS (test) $\times 10^2$
9	1.650	1.641
12	1.592	1.640
15	1.556	1.578
18	1.523	1.541
21	1.580	1.603
24	1.489	1.598
27	1.432	1.738
30	1.450	1.966
Abel profiles	2.29	2.36

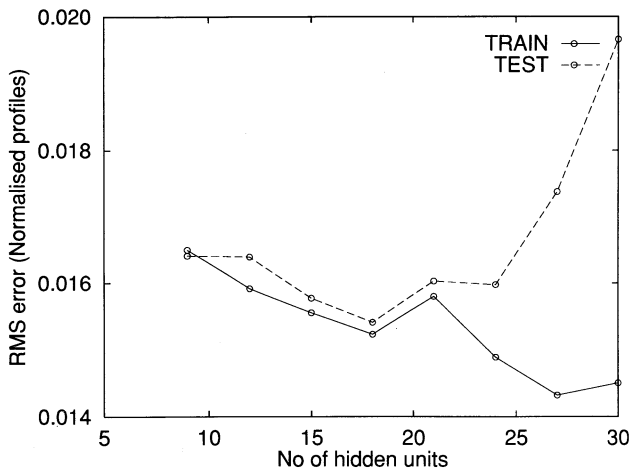


Fig. 9. Plot of the RMS error for both training and test data for the PCA approach to profile reconstruction (using seven principal components) versus the number of hidden units in the network.

Note also, by comparison with the results in Table 1, that only a small part of the overall error can be attributed to the use of a reduced dimensionality representation for the profiles.

In Figs. 10a,b we plot two examples of density profiles reconstructed using the neural network of Table 2 with 18 hidden units (since this has the smallest test error). The LIDAR profiles are shown by the circles, the Abel profiles are shown by the dashed curves, and the neural network profiles by the solid curves. The network error appearing on these graphs is the RMS error between the profile predicted by the neural network (after post multiplication by the corresponding normalisation factor $\bar{\lambda}^p$ given by Eq. (2)) and the corresponding LIDAR profile, and is defined by:

$$E_{\text{RMS}} = \left\{ \frac{1}{35} \sum_{i=1}^{35} (n_i^p - t_i^p)^2 \right\}^{1/2} \quad (13)$$

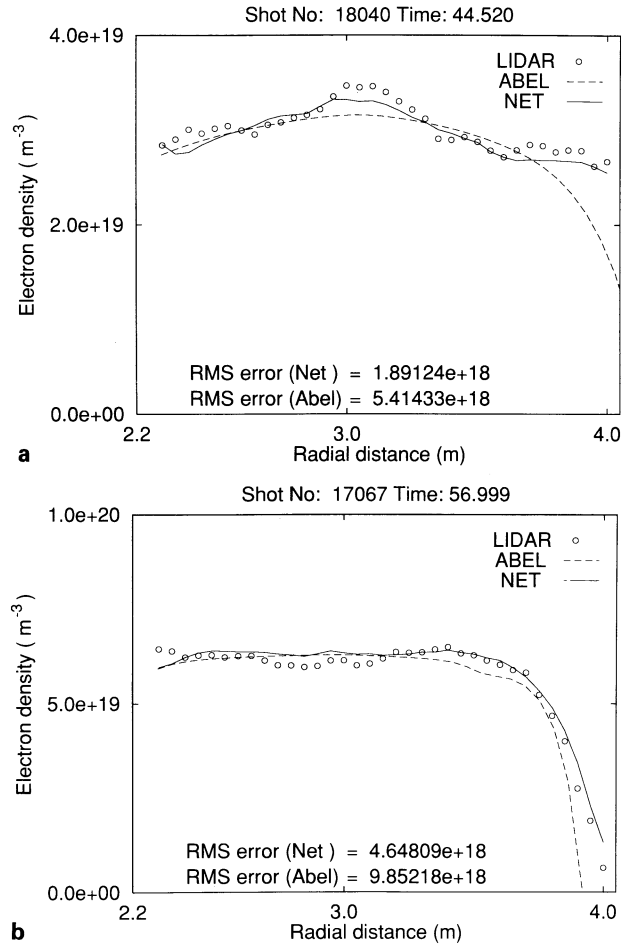


Fig. 10. a,b. Two examples of density profiles (from the test data set) reconstructed using the neural network approach (solid curves), together with the corresponding LIDAR profiles (circles) and Abel inversion profiles (dashed curves). The network chosen had 18 hidden units and used seven principal components.

The RMS error between the Abel profile and the LIDAR profile is defined in a similar way. We see that the neural network can reconstruct significantly more detail in the density profiles than can the Abel method.

In Fig. 11 we show a scatter plot of the network RMS error versus the Abel RMS error, for the profiles in the test set, again using the network with 18 hidden units. It is clear that for the majority of these profiles, the neural network gives a significantly closer fit to the LIDAR profiles than does the Abel reconstruction.

6. Discussion

In this paper we have developed a neural network approach to the problem of reconstructing profile

- approach to tokamak equilibrium control. Proceedings of the NCM'91 Conference: Applications of Neural Networks. London: Springer-Verlag, 1991
2. Allen L, Bishop CM. Neural network approach to energy confinement scaling in tokamaks. AEA Technology, UK, Report AEA-FUS-120, 1991 accepted for publication in plasma
 3. Bishop CM, Roach CM. Fast curve fitting using neural networks. (submitted to Review of Scientific Instruments)
 4. Braithwaite G, Gottardi N, Magayar G, O'Rourke J, Ryan J, Veron D. The JET Polari-interferometer. *Review of Scientific Instruments* 1989; **60**: 2825
 5. Salzmann H, Bundgaard J, Gadd C, Gowers C, Hansen KB, Hirsch K, Nielsen P, Reed K, Schrodtter C, Weisberg K. The LIDAR Thomson Scattering Diagnostic on JET. *Review of Scientific Instruments* 1989; **59**: 1451
 6. Hadamard J. Lectures on the Cauchy Problem in Linear Partial Differential Equations. Yale University Press, 1923
 7. Morozov VA. Methods for Solving Incorrectly Posed Problems. New York: Springer-Verlag, 1984
 8. Tikhonov AN, Arsenin VY. Solution of Ill Posed Problems. New York: Wiley, 1977
 9. Bishop CM. Curvature-Driven Smoothing in Backpropagation Neural Networks. *Proceedings of the International Neural Network Conference*, Paris, 1990, Vol 2, 749-752; Taylor JG and Manion CLT, editors, Theory and Applications of Neural Networks, London: Springer-Verlag, 1990, 139-148 (submitted to *Neural Networks*)
 10. Bishop CM. Improving the Generalisation Properties of Radial Basis Function Neural Networks. *Neural Computation* 1991 **2**(4): 579-588
 11. Christiansen JP. Integrated Analysis of Data from JET. *J Computat Phys* 1987; **73**: 85
 12. Cottrell GW, Munro P, Zipser D. Learning Internal Representations from Gray-Scale Images. *Proceedings of the 9th Annual Conference of the Cognitive Science Society*, 1987, 461-473
 13. Hornik K, Stinchcombe M, White H. Multilayer Feedforward Networks are Universal Approximators. *Neural Networks* 1989; **2**(5): 359-366
 14. Baldi P, Hornik K. Neural Networks and Principal Component Analysis: Learning from Examples Without Local Minima. *Neural Networks* 1989; **2**: 53-58
 15. Battiti R, Masulli F. BFGS Optimization for Faster and Automated Supervised Learning. *Proceedings International Neural Network Conference*, Paris, Vol 2, 1990, 757-760