

Fast Feedback Control of a High Temperature Fusion Plasma

Chris M. Bishop¹, Paul S. Haynes², Mike E. U. Smith², Tom N. Todd²,
and David L. Trotman²

¹Neural Computing Research Group, Department of Computer Science, Aston University, Birmingham; ²AEA Technology, Culham Laboratory, Oxfordshire, UK

One of the most promising approaches to achieving fusion of the light elements, as a potential large-scale energy source for the next century, is based on the magnetic confinement of an ionised high temperature plasma. Most of the current research in magnetic confinement makes use of toroidal plasma configurations in experiments known as tokamaks. Theoretical results have predicted that the characteristics of a tokamak plasma can be made more favourable to fusion if the cross-section of the plasma is appropriately shaped. However, the accurate generation of such plasmas, and the real-time control of their position and shape, represents a demanding problem involving the simultaneous adjustment of the currents through several control coils on time scales as short as a few tens of microseconds. In this paper, we present results from the first use of neural networks for the control of the high temperature plasma in a tokamak fusion experiment. This application requires the use of fast hardware, for which we have developed a fully parallel custom implementation of a multilayer perceptron, based on a hybrid of digital and analogue techniques. Our results demonstrate that the network is indeed capable of fast plasma control in accordance with the predictions of software simulations.

Keywords: Neural networks; Feedback control; Real-time; Plasma; Fusion; Tokamak; Hardware implementation

Original manuscript received 2 February 1994

Correspondence and offprint requests to: C.M. Bishop, Department of Computer Science, Aston University, Birmingham B4 7ET, UK.

1. Introduction

When the nuclei of light elements (such as isotopes of hydrogen) fuse they release huge amounts of energy. However, in order to fuse, the nuclei must collide with considerable kinetic energy to overcome their mutual Coulomb repulsion. In the sun, the very high temperatures provide the nuclei with the necessary kinetic energy. Over the last few decades, a considerable research effort worldwide has been aimed at trying to recreate similar conditions in the laboratory, as a potential large-scale energy source for the next century. Despite the enormous technical difficulties which must be overcome, the last few years have seen considerable progress, due in large part to the development of sophisticated devices known as 'tokamaks'. In these experiments, isotopes of hydrogen are raised to very high temperatures (up to a few hundred million degrees K) where they form a highly ionised plasma. Strong magnetic fields are used to confine the plasma against its own internal pressure, and additional magnetic fields are used to control its position and shape.

At the heart of a tokamak experiment is a toroidal vacuum vessel in which the plasma is generated. Early tokamaks produced plasmas whose toroidal cross-sections were circular. However, there is considerable interest in generating plasmas with more complex cross-sectional shapes, as these offer the possibility of significantly improved performance in terms of the physical conditions needed for fusion. However, the accurate generation of such plasmas, and the real-time control of their position and shape, represents a demanding problem involving the simultaneous adjustment of the currents

through several control coils on time scales as short as a few tens of microseconds.

In this paper, we describe an approach to the control of plasma position and shape in a tokamak based on feedforward neural networks. We also present results from the first successful use of neural networks in this application. In the next section we give a brief introduction to the principles of magnetic confinement, and in Sect. 3 we describe the problem of plasma feedback control, and review the principal conventional approaches to its solution. We then introduce an alternative approach based on feedforward neural networks. In Sect. 4 we describe the network architecture and training procedure, and in Sect. 5 we present results from software simulations of the network. The combined requirements of high bandwidth and high precision have led us to develop a fully parallel custom hardware implementation of the multilayer perceptron, described in Sect. 6, based on analogue signal processing, together with digitally stored synaptic weights. First results from real-time feedback control of the plasma in the COMPASS tokamak are presented in Sect. 7, and a brief discussion is given in Sect. 8.

2. Magnetic Confinement

To release the energy available from fusion, the nuclei of light elements, usually isotopes of hydrogen, must be forced together against their electrostatic Coulomb repulsion. This can be achieved by raising appropriate isotopes to very high temperatures so that the kinetic energy of the nuclei, with some help from quantum tunnelling, allows them to fuse. The average kinetic energy per nucleon needs to be of order 10^4 electron volts, corresponding to a temperature of 10^8 K, for the probability of fusion to be significant. Under these conditions, the hydrogen atoms are fully ionised and form an electrically conducting plasma. Economic production of power requires the fusion process to be self-sustaining so that the energy released by the fusion reactions maintains the extremely high temperatures, and the plasma is said to have reached 'ignition'. This in turn requires the density of the plasma to be high (so that the rate of reaction is high), and it requires the heat produced to be well confined (so that it does not leak away too quickly). The quantitative condition for ignition was developed by Lawson [1], who showed that the triple product of the density, temperature and energy confinement time (the characteristic time for energy to be transported out of the plasma)

must exceed a threshold value. Currently, the best plasma conditions achieved are about a factor of 5–10 short of this threshold.

To achieve such conditions in practice, the high temperature plasma must be contained and thermally insulated. Since the plasma is ionised, it can carry electrical currents and can therefore be controlled and shaped by magnetic fields. The high pressure of the plasma (governed by the product of density and temperature) can then be balanced by the pressure of the magnetic field. This is called 'magnetic confinement'. The principal experimental device for research into this approach to controlled fusion is the tokamak, which consists of a toroidal vacuum vessel containing a small quantity of hydrogen which is ionised and heated by a large toroidal electric current. This current is induced by transformer action using a time-varying magnetic field, which is itself generated by currents flowing through external coils. The plasma current also generates a magnetic field which contributes to the confinement of the high pressure plasma. A strong magnetic field, generated by large 'toroidal-field' coils, serves to stabilise the plasma, and additional magnetic fields, generated by currents flowing through 'poloidal-field' coils, control the plasma shape and position. Each plasma pulse (or shot) in COMPASS typically lasts a few hundred ms, while in large tokamaks the plasma can be sustained for several tens of seconds, and one Japanese tokamak is even able to operate in steady state.

Figure 1 shows a cross-section of the COMPASS




Fig. 1. Cutaway drawing of the COMPASS tokamak experiment at Culham Laboratory showing the D-shaped cross-section toroidal vacuum vessel. The overall height of the experiment is around 3 m.

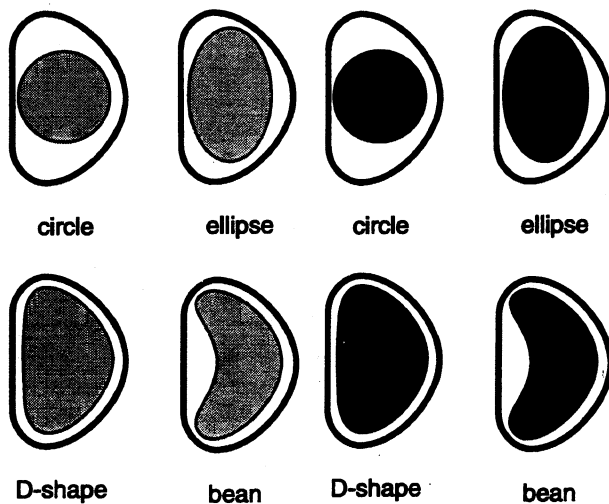


Fig. 2. Cross-sections of the vacuum vessel showing some examples of plasma shapes which COMPASS is designed to explore. The solid curve is the boundary of the vacuum vessel, and the plasma is shown by the shaded regions.

tokamak at Culham Laboratory, showing the toroidal vacuum vessel as well as the support structure and several of the control coils. Note that the cross-section of the vacuum vessel has a D-shape, unlike most earlier tokamak designs which were constructed with circular cross-section vessels. This D-shaped vessel allows non-circular cross-section plasmas to be generated, which are found to have improved energy confinement times and pressure limits, allowing significantly better plasma conditions to be achieved. Figure 2 shows some examples of the kinds of plasma shapes which can be produced in COMPASS. The Joint European Torus (JET) experiment (currently the world's largest tokamak), situated adjacent to Culham Laboratory, as well as nearly all designs for future tokamak experiments, also feature D-shaped plasmas.

3. Plasma Feedback Control

While non-circular cross-section plasmas offer a number of physics advantages over circular plasmas, they also present a number of additional problems. In particular, such plasmas are more difficult to produce and to control accurately, since currents through several control coils must be adjusted simultaneously to generate the desired shapes. Furthermore, during a typical plasma pulse, the shape must evolve, usually from some initial near-circular shape. Due to uncertainties in the current and pressure distributions within the plasma, the desired accuracy for plasma control can only be achieved by making real-time measurements of the

position and shape of the boundary, and using error feedback to adjust the currents in the control coils.

The physics of the plasma equilibrium is determined by force balance between the magnetic pressure and the thermal pressure of the plasma, and is relatively well understood. We shall assume that the tokamak has rotational symmetry about the vertical axis, which is an excellent approximation in the present context. Since variations around the torus are neglected it is sufficient to consider one particular cross-sectional plane. As a consequence of the assumed axi-symmetry, the component B_p of magnetic field in this plane (known as the 'poloidal field') can be written in terms of the gradient of a scalar flux function Ψ :

$$\mathbf{B}_p = \mathbf{e} \times \nabla \Psi(R, Z) \quad (1)$$

where \mathbf{e} is a unit vector in the toroidal direction (i.e. perpendicular to the cross-sectional plane), and R and Z are the radial and vertical coordinates, respectively, such that the Z -axis is the axis of symmetry of the tokamak. These coordinates are illustrated in Fig. 3. The surfaces of constant Ψ are known as 'flux surfaces', and in three dimensions any magnetic field line will spiral around the torus while remaining confined to such a surface. One of the flux surfaces forms the boundary of the plasma. The requirement of force balance then leads to a description of plasma equilibrium configurations in terms of solutions of the Grad-Shafranov equation [2], given by:

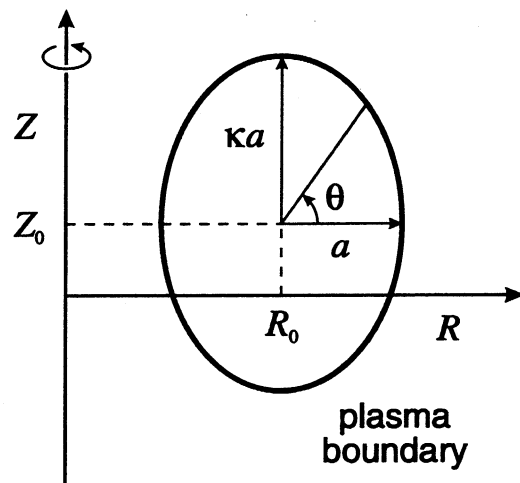


Fig. 3. Schematic illustration of a cross-section of the toroidal vacuum vessel showing the definitions of various coordinates and parameters. The elliptical curve denotes the plasma boundary. The Z -axis is the axis of symmetry of the torus, R is the major radius coordinate and θ is called the poloidal angle. The triangularity δ (not shown) describes the departure of the plasma boundary from a simple ellipse (values of $\kappa = 1$ and $\delta = 0$ correspond to a circular plasma boundary).

$$R \frac{\partial}{\partial R} \left(\frac{1}{R} \frac{\partial \Psi}{\partial R} \right) + \frac{\partial^2 \Psi}{\partial Z^2} = -RJ(\Psi, R) \quad (2)$$

where the function $J(\Psi, R)$ specifies the toroidal plasma current density, and takes the form:

$$J(\Psi, R) = -Rp'(\Psi) - \frac{I(\Psi)I'(\Psi)}{R} \quad (3)$$

Here the function $p(\Psi)$ and $I(\Psi)$ describe the plasma pressure and the toroidal magnetic flux, respectively, and are determined by a variety of competing processes within the plasma such as impurity radiation, and thermal and particle transport. Since prediction of these processes from first principles is exceedingly difficult (due to non-linear turbulence-like processes in the plasma), it is usual to assume simple parameterised functional forms for $p(\Psi)$ and $I(\Psi)$. One specific example of such a parameterisation will be discussed in Sect. 4. Fortunately, the plasma configurations obtained from solution of the Grad-Shafranov equation are relatively insensitive to the precise choice of these functions.

Due to the non-linear nature of the Grad-Shafranov equation, a general analytic solution is not possible. However, for a given current density function $J(\Psi, R)$, the Grad-Shafranov equation can be solved by iterative numerical methods, with boundary conditions determined by currents flowing in the external control coils which surround the vacuum vessel. On the tokamak itself it is changes in these currents which are used to alter the position and cross-sectional shape of the plasma. Numerical solution of the Grad-Shafranov equation represents the standard technique for post-shot analysis of the plasma, and is also the method used to generate the training dataset for the neural network, as described in the next section. However, this approach is computationally very intensive, and is therefore unsuitable for feedback control purposes.

For real-time control it is necessary to have a fast (typically $\leq 100 \mu\text{s}$) determination of the plasma boundary shape. This information can be extracted directly from a variety of diagnostic systems, the most important being local magnetic measurements taken at a number of points around the perimeter of the vacuum vessel. Most tokamaks have several tens or hundreds of small pick up coils located at carefully optimised points around the torus for this purpose. It is convenient to represent these magnetic signals collectively as a vector \mathbf{m} . Other diagnostic information can also be used, for instance information about the spatial profiles of density, temperature and poloidal magnetic field (which together

help to determine the current profile function J) obtained from laser diagnostics. In the work reported here, only magnetic signals have been considered, although it is intended to extend this to include other diagnostic information in the future.

The position and shape of the plasma boundary can be described in terms of a set of geometrical parameters such as vertical position and elongation (the particular parameters used in this work will be discussed in Sect. 4). The basic problem which has to be addressed, therefore, is to find a representation for the mapping from the magnetic signals \mathbf{m} to the values of the geometrical parameters (which we denote by y_k), which can be implemented in suitable hardware for real-time control. This mapping is non-linear, but is continuous and smooth for almost all values of \mathbf{m} . For tokamaks which operate with a small range of shapes (circular cross-sections only for instance) the mapping may be approximated by a linearisation around a fixed operating point. However, if a wide range of shapes must be accommodated, as, for instance, when a strongly shaped plasma is evolved from some initially circular configuration, then the non-linearity of the mapping must be taken into account.

Several conventional approaches have been considered for the solution of this problem, many of which rely on the generation of a large dataset of representative equilibria obtained by numerical solution of the Grad-Shafranov equation. A parameterised functional form can then be fitted to the dataset using, for instance, a least-squares procedure. We can group these approaches into the following three categories:

1. Hand-selected functional forms are prescribed which contain a number of parameters whose values are determined by fitting to the data set. Generally, these functions are designed to describe the mapping in the neighbourhood of a particular operating point, and yet capture some of the non-linearity of the mapping. The functions are hardwired for use in real-time control. This technique has been applied successfully to the DIII-D tokamak [3]. Its main limitation is that it is still limited to a restricted range of equilibria.

2. The dimensionality of the vector \mathbf{m} (which is typically of order 10–100) is first reduced by a principal components projection, based on the dataset measurements, to give a vector $\tilde{\mathbf{m}} = \mathbf{P} \cdot \mathbf{m}$, where the rows of the matrix \mathbf{P} are given by the eigenvectors of the covariance matrix for the vector \mathbf{m} with respect to the dataset. Subsequently, the parameters y_k are expressed as a general quadratic function of $\tilde{\mathbf{m}}$ in the form:

$$y_k = w_k + \mathbf{w}_k \cdot \tilde{\mathbf{m}} + \tilde{\mathbf{m}} \cdot \mathbf{W}_k \cdot \tilde{\mathbf{m}} \quad (4)$$

where the parameters w_k , \mathbf{w}_k and \mathbf{W}_k are determined by least-squares fitting to the dataset. This technique, known as 'function parameterisation', was developed for the interpretation of spark chamber data from high energy physics experiments [4], and subsequently applied to the interpretation of data from the ASDEX tokamak [5,6]. While it can represent some degree of non-linearity, it is clear that the unsupervised nature of the principal components reduction is in general sub-optimal, and that the restrictive choice of a quadratic fit (which is based on computational simplicity and not on any prior knowledge) will also in general be sub-optimal.

3. Since the required mapping is continuous, it can be approximated to arbitrary accuracy by a piecewise-linear function. Each linear segment can be determined by least-squares fitting to an appropriate subset of the database, or can be obtained analytically by expansion of the Grad-Shafranov equation around a given equilibrium [7]. The mapping itself is now represented by a simple matrix, but the elements of the matrix must be switched as the plasma moves from one region of \mathbf{m} -space to another. This requires that the hardware implementation allows the matrix elements of the linear mapping to be switched rapidly many times during the plasma pulse as the plasma evolves from one region of parameter space to another. A hardware system which implements this approach has been developed for control of the TCV and ALCATOR C-MOD tokamaks [8]. The principal drawbacks of this approach are the complications involved in deciding how to partition the input space into suitable regions, and the potential problems arising from discontinuous switching within a control loop.

In this paper, we consider an alternative approach to the problem of tokamak equilibrium feed-back control based on feedforward neural networks [9–11]. As in other approaches, a parameterised functional form is fitted to a large database of numerical equilibria. The functional fit, however, is produced using a multilayer perceptron, which is capable in principle of approximating with arbitrary accuracy an arbitrary continuous mapping from a compact set [12]. The need to switch parameters can therefore be avoided, while no a-priori limitation on the functional form of the mapping is introduced.

Figure 4 shows a block diagram of the control loop for the neural network approach to tokamak equilibrium control. The rôle of the neural network is to map the measured magnetic signals \mathbf{m} onto the values y_k of the geometrical parameters describing the plasma position and shape. These are

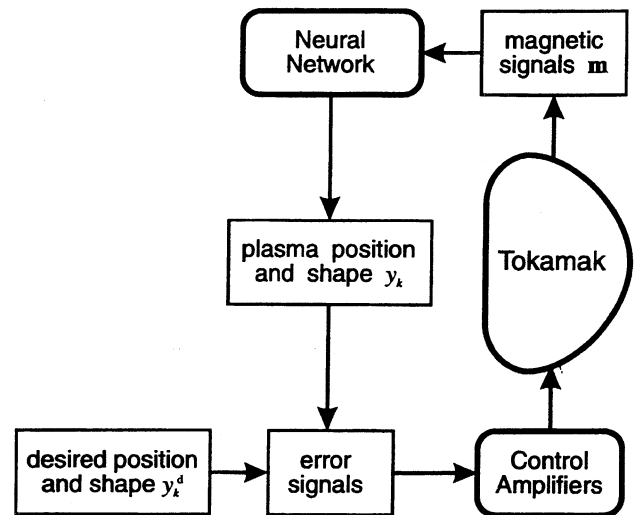


Fig. 4. Block diagram of the control loop used for real-time feedback control of plasma position and shape.

compared with desired values y_k^d for the parameters, which are pre-programmed as functions of time prior to the plasma pulse. The differences $y_k - y_k^d$ constitute error signals which generate corrections to the currents flowing in the control coils, through standard linear proportional-differential (PD) control laws.

4. Dataset and Network Architecture

The dataset for training and testing the network was generated by numerical solution of Eq. (2) using a free-boundary equilibrium code. This code contains a detailed description of the COMPASS hardware configuration, and allows the boundary conditions to be expressed directly in terms of currents in the control coils. The database currently consists of over 2000 equilibria spanning the wide range of plasma positions and shapes available in COMPASS. Each configuration takes several minutes to generate on a fast UNIX workstation. Since the shape of the plasma boundary is not strongly sensitive to the precise form of the current density function J , we have considered the specific form:

$$J(\Psi, R) = b \left\{ \beta R + \frac{(1-\beta)}{R} \right\} (1 - \Psi^{\alpha_1})^{\alpha_2} \quad (5)$$

where b is a constant, β controls the ratio of plasma pressure to magnetic field energy density, and the parameters α_1 and α_2 are numbers ≥ 1 which can be varied to generate a variety of current profiles.

For a large class of equilibria, the plasma boundary can be reasonably well represented in terms of a

simple parameterisation, governed by the poloidal angle θ , given by:

$$\begin{aligned} R(\theta) &= R_0 + a \cos(\theta + \delta \sin \theta) \\ Z(\theta) &= Z_0 + a\kappa \sin \theta \end{aligned} \quad (6)$$

where we have defined the following parameters:

- R_0 radial distance of the plasma centre from the major axis of the torus
- Z_0 vertical distance of the plasma centre from the torus midplane
- a minor radius measured in the plane $Z = Z_0$
- κ elongation
- δ triangularity

which are shown geometrically in Fig. 3. Each of the entries in the database has been fitted using the form in Eq. (6), so that the equilibria are labelled with the appropriate values of the shape parameters. Figure 5 shows scatter plots of pairs of boundary parameters for the whole dataset. These are seen to form distinct regions which correspond to the physical range of the tokamak operating space. For example, in the plot of elongation versus triangularity, there are no points with high triangularity and low elongation, since all of the shaping coil configurations produce elongation as well as triangularity. Within these regions the parameter spaces are seen to be reasonably well populated. The plot of triangularity versus major radius shows stripes, due to the fact that, for any

given configuration of the shaping coil circuits as provided on the COMPASS tokamak, there is a strong correlation between triangularity and major radius. Different configurations produced different correlations, but the configurations themselves can only be changed in discrete steps. Such correlations are also expected in experimentally achieved equilibria. The dataset was randomly partitioned into halves for training and testing.

The results presented in this paper are based on a multilayer perceptron architecture with a single hidden layer. The units in the hidden layer compute a symmetric sigmoid function of the form:

$$f(x) \equiv \frac{1 - e^{-x}}{1 + e^{-x}} \quad (7)$$

where x is the input to the unit and $f(x)$ is the output. Units in the output layer are taken to be linear.

It is important to note that the transformation from magnetic signals to flux surface parameters involves an exact linear invariance. This follows from the fact that, if all of the currents are scaled by a constant factor (so that the pressure profile term in Eq. (3) is implicitly scaled by the same factor), then the magnetic fields will be scaled by this factor, and the geometry of the flux surfaces, and hence of the plasma boundary, will be unchanged. It is important to take advantage of this prior knowledge and to build it into the network structure, rather than force the network to learn it by example.

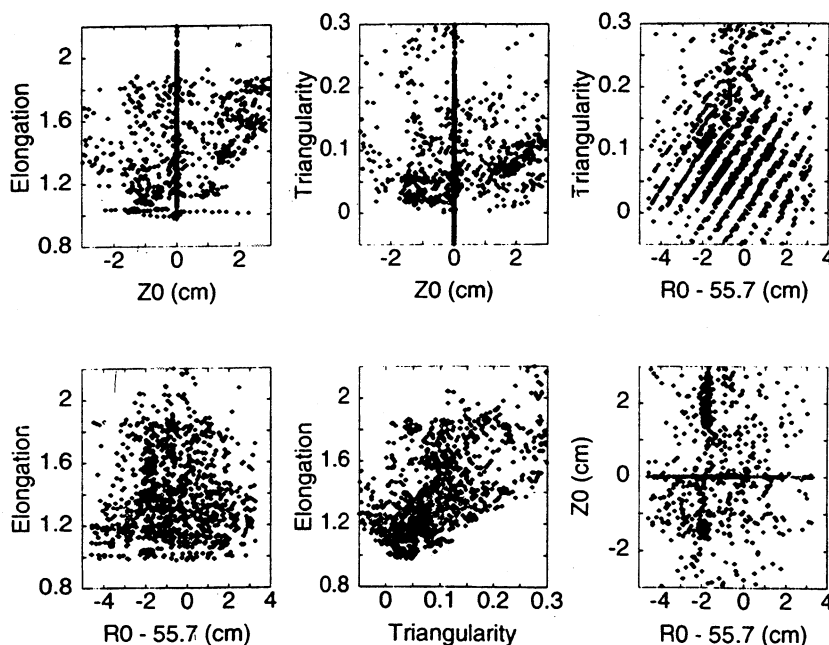


Fig. 5. Scatter plots of plasma boundary parameters from the database.

We therefore normalise the vector \mathbf{m} of input signals to the network by dividing by a common factor proportional to the total plasma current. A scaling of the magnetic signals by a common factor then leaves the network inputs (and hence the network outputs) unchanged. This normalisation of the input vector brings three distinct advantages. First, the network exhibits exact invariance to rescaling of the currents, compared with the approximate invariance which would result if the network had to learn by example. Second, output accuracy can be maintained over a wide range of plasma currents. The input signals have a significant dynamic range (as the toroidal plasma current grows from a few kA to a few 100 kA during the plasma pulse), and without this normalisation we would expect a loss of accuracy at low plasma currents due to the limited precision of the hardware implementation. Finally, the network training can be performed with a smaller dataset than would otherwise be possible, which can be generated for just one value of total plasma current. Note that the normalisation has to be incorporated into the hardware implementation of the network, as will be discussed in Sect. 6.

5. Results from Software Simulations

Networks are trained by minimization of a sum-of-squares error defined over the network outputs. The target data from the training set are first normalised so that each output has zero mean and unit standard deviation. This ensures that the outputs are treated on an equal footing. Note that a weighting of the sum-of-squares error could in principle be used to enforce a higher relative accuracy on some of the outputs at the expense of others, although this has not been explored in the work done to-date. Error backpropagation is used to compute the derivatives of the error with respect to the weights and thresholds in the network. These derivatives are then used in a standard conjugate gradients optimisation algorithm [13] to find a minimum of the error function. In all studies conducted so far, networks having a single hidden layer were used, with full interconnections from inputs to hidden units and full interconnections from hidden units to output units. The number of hidden units was optimised by training networks having various numbers of hidden units and then comparing them using the test set.

The results from the neural network mapping are compared with those for the optimal linear mapping, that is the single linear transformation which minimises the same sum-of-squares error as is used

in the neural network training algorithm. This minimisation can be expressed in terms of a set of linear equations whose solution can be found efficiently and robustly using the technique of singular value decomposition [13]. This comparison also provides some indication of the degree of non-linearity involved in the mapping.

On the COMPASS experiment, there are some 120 magnetic signals which could be used to provide inputs to the network. Since each input could either be included or excluded, there are potentially $2^{120} \approx 10^{36}$ possible sets of inputs which might be considered. In principle, to select the best subset, it is necessary to optimise a network on each subset and compare their performances. This is clearly impractical, and so we have used two simplifications of the selection process. First, for the purposes of input selection, we have replaced the neural network with the simple linear mapping. It is hoped that the set of inputs which are chosen on this basis will also be appropriate for the neural network. This is a reasonable expectation, since it is known that the linear mapping can give an acceptable result provided the range of equilibria is not too great. The second simplification involves the use of forward sequential selection [14] to choose the inputs, rather than exhaustive search. For N possible inputs, this approach starts by considering the N possible choices of just one input, and selecting the one which gives the best result, in the sense of producing the smallest residual error. This input is retained, and all possible sets of two inputs obtained by combining the first selected input with each of the $N - 1$ remaining candidates, are considered, and again the best choice is retained. This process is continued until a sufficient number of inputs has been accumulated. Simulations aimed at finding a network suitable for use in real-time control have so far concentrated on 16 inputs, since this is the number available from the initial hardware configuration (see Sect. 6).

Initial results were obtained on networks having three output units, corresponding to the values of vertical position Z_0 , major radius R_0 and elongation κ , these being parameters which are of interest for real-time feedback control. Results from neural networks with various numbers of hidden units, and from the optimal linear mapping, are shown in Table 1. These results are plotted in Fig. 6

The smallest average test set error is obtained from the network having 16 hidden units. Typically, the residual error on test data is reduced by about 30%–40% as a result of going from the linear system to the best neural network. In the context

Table 1. Comparison of the training and test set errors for networks having various numbers N_H of hidden units. The linear mapping results are shown as $N_H = 0$

N_H	Training				Test			
	ϵ_R	ϵ_Z	ϵ_K	$\bar{\epsilon}$	ϵ_R	ϵ_Z	ϵ_K	$\bar{\epsilon}$
0	15.02	13.34	24.75	18.40	15.40	12.50	24.70	18.30
4	11.70	8.19	16.93	12.79	11.13	9.14	20.87	14.64
6	10.54	8.66	11.53	10.31	10.18	9.76	20.09	14.17
8	8.90	7.07	9.35	8.50	11.45	7.97	21.18	14.64
10	8.36	6.26	8.18	7.66	10.71	7.48	16.47	12.14
12	7.11	6.44	9.26	7.70	12.84	8.53	16.59	13.08
14	7.58	5.96	8.35	7.37	11.61	7.92	19.47	13.86
16	6.80	6.07	7.52	6.82	10.49	7.38	15.66	11.69
20	7.06	5.59	7.15	6.64	11.12	9.04	17.26	12.96
24	5.93	5.34	6.52	5.95	11.15	8.25	18.08	13.16
28	5.82	5.20	6.86	6.00	11.59	9.51	16.99	13.08

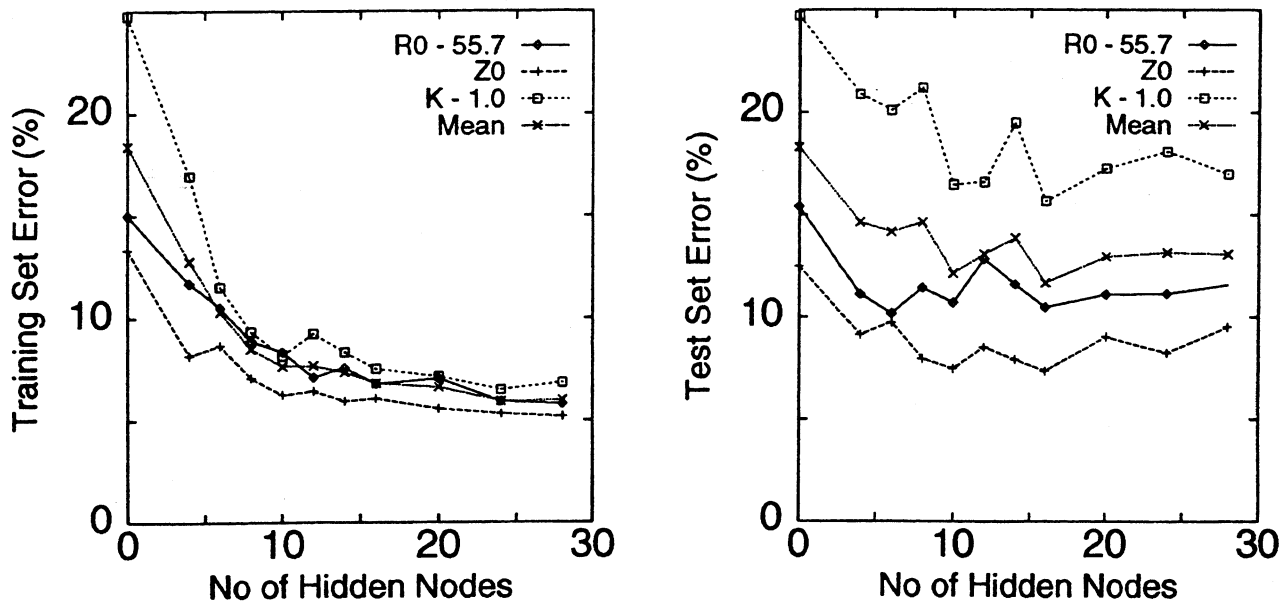


Fig. 6. Plot of the normalised errors for both training and test sets versus the number of hidden units, as given in Table 1. The results from the linear mapping are shown as 0 hidden units.

of this application, such an improvement is very significant.

For the experiments on real-time feedback control described in Sect. 7, the available hardware only permitted networks having four hidden units, and so we consider the results from this network in more detail. Figure 7 shows two examples of reconstructed plasma boundaries. Each example shows the cross-section of the COMPASS tokamak vacuum vessel together with the plasma boundary

from one of the test set equilibria, and the corresponding predictions for the boundary (plotted using the parametric representation in Eq. (6)) from the neural network and from the optimal linear mapping.

Figure 8 shows plots of the network predictions for various parameters versus the corresponding values from the test set portion of the database. Analogous plots for the optimal linear map predictions versus the database values are also shown. Comparison of the corresponding figures show the




Fig. 7. Two examples of the prediction of the plasma boundary shape using the outputs from the neural network (with four hidden units) and from the linear mapping, for equilibria from the test set. In each case, the COMPASS vacuum vessel is given by the outer curve, the solid curve shows the boundary obtained from the equilibrium code, the prediction from the neural network is shown by the heavy dashed curve, and the prediction from the linear mapping is shown by the light dashed curve. In the right-hand diagram, the network prediction and the actual boundary are barely distinguishable.

superior predictive capability of the neural network approach.

6. Hardware Implementation

The hardware implementation of the neural network must have a bandwidth of ≥ 10 kHz to cope with the fast timescales of the plasma evolution, and an output precision of at least 8 bits to ensure that the final accuracy which is attainable will not be limited by the hardware system. It is also desirable, for development work, to have a system in which the network architecture can be readily changed to explore different numbers of inputs, different numbers of hidden units, the effects of a second hidden layer, and so on. We have chosen to develop a full parallel custom implementation of the multilayer perceptron, based on analogue signal paths with digitally stored synaptic weights [15]. While in principle a fully digital solution (based on arrays of digital signal processing chips for instance) is always a possible alternative, the use of analogue signal paths avoids the need for conversions to and from the digital representation, and fits more naturally into the control framework currently used on most tokamaks including COMPASS.

A VME-based modular construction has been chosen as this allows flexibility in changing the network architecture, ease of loading network weights, and simplicity of data acquisition. Three

separate types of card have been developed as follows:

- A combined 16-input buffer and signal normalizer
- A 16×4 matrix multiplier and output driver
- A 4-channel sigmoid module

We now describe each of these cards in more detail.

6.1. Input Normaliser

Due to the significant dynamic range of the measured signals, and the requirement for invariance to rescaling of the input signals (as discussed in Sect. 4), an analogue hardware implementation of the input vector normalisation has been developed. This has been combined with input buffering on a separate VME card, which has been designed to provide independent scaling of groups of eight inputs by an arbitrary function of an external reference signal. In the present application, the reference signal is taken to be the plasma current (determined by a magnetic pick-up loop called a 'Rogowski coil'), and the function is chosen to be a simple inverse proportionality. The reference signal is converted to digital form using an ADC (analogue-to-digital converter), and this is mapped to another digital signal using an EPROM look-up table, which is then used to provide a scaling factor for the main signal inputs via multiplying DACs (digital-to-analogue converters). The normalising system operates in less than $4 \mu\text{s}$, allowing the buffer gain to be adjusted at a bandwidth in excess of 100 kHz.

6.2. Matrix Multiplier

The requirement for 8-bit precision suggests that the synaptic weights should be stored to significantly higher precision than this. Detailed simulations of the hardware network indicate that in feedforward mode there is little degradation in error as the precision of the weights is reduced until about 10-bit precision is reached. The error then starts to increase, becoming unacceptable for precisions less than about 7-bits. The synaptic weights are produced using 12-bit frequency-compensated multiplying DACs, which can be configured to allow 4-quadrant multiplication of analogue signals by a digitally stored number. The weights are obtained as a 12-bit 2's-complement representation from the VME backplane. Note that the DACs are being used here as digitally controlled attenuators, and not in




Fig. 8. Plots of the values from the test set versus the values predicted by the linear mapping for the three equilibrium parameters, together with the corresponding plots for the neural network with four hidden units.

their usual rôle of converting digital signals into analogue signals. Each module contains four separate 16-input 1-output multipliers, on separate daughter boards. The products from each group of 16 weighted inputs are summed internally to provide the final output to the next module, so that each matrix card provides a 16×4 matrix of weights. Synaptic weights are downloaded (prior to the plasma pulse) via the VME backplane from a central control computer, using an addressing technique to label the individual weights. The complete system includes extensive diagnostics, allowing voltages at all key points within the network to be monitored as a function of time via a series of multiplexed output channels.

6.3. Sigmoidal Summer

There are many ways in which to produce a sigmoidal non-linearity, and we have opted for a solution using two transistors configured as a long-tailed-pair, to generate a transfer characteristic of the form given in Eq. (7). The principal drawback of such an approach is the temperature sensitivity due to the appearance of temperature in the denominator of the exponential transistor transfer characteristic. An elegant solution to this problem has been found by exploiting a chip containing five transistors in close thermal contact. Two of the transistors form the long-tailed pair, one of the

transistors is used as a heat source, and the remaining two transistors are used to measure temperature. External circuitry provides active thermal feedback control, and stability to changes in ambient temperature over the range 0°C – 50°C is found to be well within the acceptable range. Noise levels and response times are also well within the required limits. A separate 12-bit DAC system, identical to those used in the matrix multiplier cards, but with a fixed DC input, is used to provide a bias for each sigmoid.

7. Real-Time Feedback Control

Prior to use in closed-loop feedback control, extensive testing of the hardware was performed to ensure its correct operation. A software implementation of the network in feedforward mode was developed, which provides a detailed simulation of the hardware system, and in particular allows voltages at all diagnosed points within the network to be predicted. After carrying out full tests of the various modules to ensure they met the design specification, a complete network system was configured and connected to the tokamak in an open-loop mode in which the network receives the diagnostic signals from the magnetic pick-up coils, but where the network outputs are simply digitised for post-shot analysis. During this time the feedback control was

provided by a linear system based on a combination of four hand-selected magnetic signals. This simple controller is adequate to allow the generation of usable plasmas, but does not allow accurate control of a wide range of shapes. The input signals to the network were digitised to allow for post-shot reconstruction of the plasma configuration using the same free-boundary equilibrium code as was used to generate the training data (see Sect. 4). In addition, the outputs from the network were digitised, as were the voltages at all key points within the network.

The network architecture used had 16 inputs, four hidden units and three output units, and was trained in software to output the values of major radius R_0 , vertical position Z_0 and elongation κ . Initial tests with the network operating open-loop allowed all of the internal and output signals from the network to be monitored, and to be compared with the corresponding signals from the software simulation of the network when loaded with the same weight and bias values and presented with the same input signals. Satisfactory agreement between these was obtained, and ensured that such issues as the sign conventions on the input signals and the addressing of individual weights in the network had been handled correctly.

Finally, the output signals from the network were used for closed-loop real-time feedback control of the plasma elongation κ , using the same network configuration and weight values as had been used during the open-loop evaluation phase. The remaining two outputs were digitised for post-shot diagnosis, but were not used for feedback control. An example of the results is given in Fig. 9, which shows the evolution of the plasma elongation as a function of time during a plasma pulse. Here the desired elongation has been pre-programmed to follow a series of steps as a function of time. Initial interpretation of the data was performed using a simple 'filament' code (which models the current distribution inside the plasma as a small number of filamentary currents). This gives relatively rapid post-shot plasma shape reconstruction but with limited accuracy. A more careful reconstruction using the full free-boundary equilibrium code gives closer agreement with the network predictions. The graph clearly shows the network generating the required elongation signal in close agreement with the reconstructed values. The typical residual error is of order 0.07 on elongation values up to around 1.5. Much of this error is attributable to the restricted number of hidden units available with the initial hardware configuration, and is expected to fall significantly when larger networks become

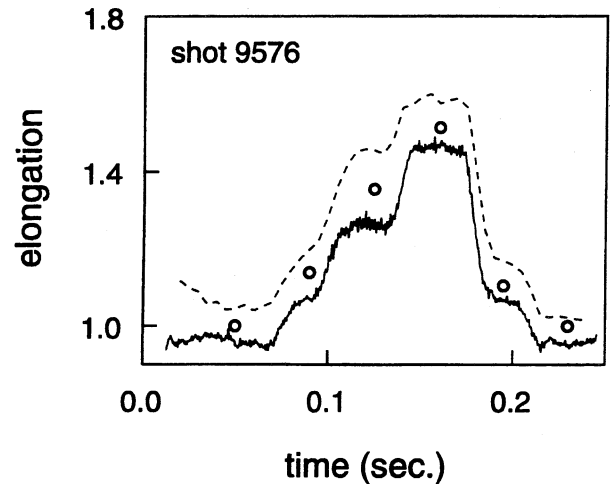


Fig. 9. Plot of the plasma elongation κ as a function of time during shot no. 9576 on the COMPASS tokamak. The solid curve shows the value of elongation produced as output (in real time) by the neural network. The dashed curve shows the post-shot reconstruction of the elongation obtained from a simple filament code, using the same signals as provided to the network, while the circles denote the more accurate reconstruction obtained from the full equilibrium solver.

available. Additional improvements in the accuracy with which κ is determined should be possible by weighting the sum-of-squares error function used during training, to reduce the errors on κ at the expense of increased errors on the other variables. The use of a larger set of input variables is also expected to give further reductions in error. It should be noted that the plasma behaved satisfactorily throughout the shot, without any need for switching of the mapping parameters. While these results represent the first obtained using closed loop control, it is clear from earlier software modelling of larger network architectures (such as 32-16-4) that residual errors of order a few per cent should be attainable. The implementation of such larger networks is being pursued, following the successes with the smaller system.

8. Discussion

We have shown that the multilayer perceptron offers an alternative approach to the problem of real-time feedback control of the plasma in a tokamak experiment. Unlike many conventional approaches, it does not assume a specific functional form for the mapping needed to determine the plasma shape, nor does it involve switching of the mapping in real-time. Simulation results demonstrate that the technique is capable of achieving the desired accuracy over a wide range of equilibrium

configurations. Custom neural network hardware has been designed which combines analogue signal paths with 12-bit digital synaptic weights to give high bandwidth together with high precision. This hardware system is well diagnosed internally, and can readily be reconfigured for a wide variety of network architectures.

We have also presented the first results from the use of neural networks for real-time feedback control of the plasma in a tokamak fusion experiment. These show that the network is indeed capable of performing as expected in controlling the plasma elongation in the range 1.0–1.5 while simultaneously decoding two other equilibrium parameters. Software simulations indicate that the neural network approach should compare favourably with alternative methods being considered for the control of strongly shaped plasmas. While it is too early to form a definitive conclusion as to which technique is the best, it is already clear that the neural network approach must be a strong contender.

Work is currently underway to improve the accuracy of the neural network, and to provide control of several shape parameters simultaneously. In addition, the possibility of extending the network mapping to include information from other diagnostics (such as laser scattering systems for measuring plasma pressure) to provide still greater accuracy is being considered.

Neural networks have already been used with great success for fast interpretation of the data from tokamak plasma diagnostics to determine the spatial and temporal profiles of quantities such as temperature and density [16–18]. There is currently considerable interest in extending these techniques to allow real-time feedback control of the profiles to give more complete determination of the plasma configuration than is possible by boundary shape control alone. For such applications, neural networks appear to offer one of the most promising approaches.

Acknowledgements. We would like to thank Peter Cox, Jo Lister and Colin Roach for many useful discussions and technical contributions. This work was partially supported by the UK Department of Trade and Industry.

References

1. Lawson JD. Some criteria for a power producing thermonuclear reactor. *Proc Phys Soc* 1957; B70 (6).
2. Shafranov VD. On magnetohydrodynamical equilibrium configurations. *Sov Phys JETP* 1958; 8: (710).
3. Osborne TH, Fukumoto H, Hosogane N, et al. Plasma shape and position control on DIII-D. *Bull Am Phys Soc* 1986; 31: 1502.
4. Wind H. Proceedings of the CERN Computing and Data Processing School, 1972, CERN-72-21, 53.
5. Braams BJ, Gilje W, Lackner K. Fast determination of plasma parameters through function parameterization. *Nuclear Fusion* 1986; 26: 699–708.
6. Gruber O, Gernhardt J, McCarthy P, Lackner K, Seidel U, Schneider W, Woyke W, Zehrfeld HP. Proceedings of the 17th Symposium on Fusion Technology, Rome, Italy, 1992; 2: 1042–1046.
7. Hofmann F, Tonetti G. Fast identification of plasma boundary and X-points in elongated tokamaks. *Nuclear Fusion* 1988; 28: 519.
8. Lister JB, Marmillod P, Moret JM. Control systems for the TCV tokamak. Report LRP 332/87, Centre de Recherche en Physique des Plasmas, Lausanne, 1987.
9. Lister JB, Schnurrenberger H. Fast non-linear extraction of plasma parameters using a neural network mapping. *Nuclear Fusion* 1991; 31: 1291–1300.
10. Bishop CM, Cox P, Haynes PS, Roach CM, Smith MEU, Todd TN, Trotman DL. A neural network approach to tokamak equilibrium control. In *Neural Network Applications* (JG Taylor, ed.). Springer-Verlag, 1992; 114–128.
11. Lagin L, Bell R, Davis S, Eck T, Jardin S, Kessel C, Mcenerney J, Okabayashi M, Popyack J, Sauthoff N. Application of neural networks for real-time calculations of plasma equilibrium parameters for PBX-M. Proceedings of the 17th. Symposium on Fusion Technology, Rome, Italy, 1993; 2: 1057–1061.
12. Hornik K, Stinchcombe M, White H. Multilayer feedforward networks are universal approximators. *Neural Networks* 1989; 2: 359–366.
13. Press WH, Flannery BP, Teukolsky SA, Vetterling WT. *Numerical Recipes in C: The Art of Scientific Computing* (2nd Ed.), Cambridge University Press, 1992.
14. Fukunaga K. *Statistical Pattern Recognition* (2nd Ed.). Academic Press, 1990.
15. Bishop CM, Haynes PS, Roach CM, Smith MEU, Todd TN, Trotman DL. Hardware implementation of a neural network for plasma position control in COMPASS-D. Proceedings of the 17th. Symposium on Fusion Technology, Rome, Italy, 1993; 2: 997–1001.
16. Bishop CM, Roach CM, von Hellerman M. Automatic analysis of JET charge exchange spectra using neural networks. *Plasma Physics and Controlled Fusion* 1993; 35: 765–773.
17. Bishop CM, Strachan IGD, O'Rourke J, Maddison G, Thomas PR. Reconstruction of tokamak density profiles using feedforward networks. *Neural Computing and Applications* 1993; 1: 4–16.
18. Bartlett DV, Bishop CM. Development of neural network techniques for the analysis of JET ECE data. Proceedings of the 8th International Workshop on ECE and ECRH, 1992.