Review

Review of neural network applications in medical imaging and signal processing

A. S. Miller[1] B. H. Blott[1] T. K. Hames[2]

[1] Department of Physics, University of Southampton, Highfield, Southampton, Hants. SO9 5NH, UK
[2] Department of Medical Physics & Bioengineering, Southend General Hospital, Prittlewell Chase, Westcliff-on-Sea, Essex S50 0RY, UK

Abstract—*The current applications of neural networks to in vivo medical imaging and signal processing are reviewed. As is evident from the literature neural networks have already been used for a wide variety of tasks within medicine. As this trend is expected to continue this review contains a description of recent studies to provide an appreciation of the problems associated with implementing neural networks for medical imaging and signal processing.*

Keywords—*Medical imaging, Medical signal processing, Neural networks*

Med. & Biol. Eng. & Comput., 1992, 30, 449–464

1 Introduction

NEURAL NETWORKS are computational techniques for recognising patterns, which have undergone a renaissance of interest over the past six years. Networks employing various architectures and learning algorithms are finding uses in both academic research and industrial applications. Research into neural networks has been active since the 1940s but underwent a decline in the late 1960s following the work of Minsky (see his latest book (MINSKY and PAPERT, 1988)), which showed that the network designs available at that time were unable to solve relatively trivial problems. Recently neural network research has received fresh impetus as new network learning paradigms have been developed. The most widely used of these is the back-propagation algorithm, introduced by RUMELHART *et al.* (1986). The field continues to attract new thinking, and a wide variety of network architectures and learning rules are now being investigated (WIDROW and LEHR, 1990). A comprehensive description of neural networks relevant to biological and biomedical research is given by CLARK (1991), but with the main emphasis given to neural modelling. Beale's and Wasserman's introductory texts on neural networks also provide sound insight into the theoretical aspects of neural networks (BEALE and JACKSON, 1990; WASSERMAN, 1989). In this review we limit the discussion to the analysis of medical data acquired *in vivo*.

A neural network consists of simple processing units, called 'neurons' or 'nodes', which bear only a passing resemblance to actual biological neurons. Each neuron is connected to other neurons in the network by unidirectional connections of different strength or weight. Early designs were based on perceptrons, processing elements which summed their weighted inputs and applied a linear transfer function to produce an output. Currently networks exist which use both linear and nonlinear functions, and either binary or continuous activations.

First received 26th June 1991 and in final form 30th January 1992

The neurons are usually arranged in a series of layers, bounded by input and output layers encompassing a variable number of hidden layers, connected in a structure which depends on the complexity of the problem to be solved. Connections usually feed from the input to the output layers, a feedforward network, although feedback connections from hidden layer to input layer, for example, are possible. Information from the input layer is thus propagated down through the network to the output layer. A layer is fully connected if each node in the layer is connected to all the nodes in an adjacent layer (Fig. 1). However, it can be advantageous to use partially connected layers. It should be noted that the input layer serves only as a means of acquiring input and distributing it to the other network neurons.

The most important features of networks are their ability to learn from examples, and because the learnt information is stored across the network weights, to generalise. This means that appropriate classifications will be made even for input patterns not actually included in the training set, provided that the training set covered a representative group of patterns. This ability to learn and generalise means that neural networks have the potential for solving image- and signal-processing problems which are

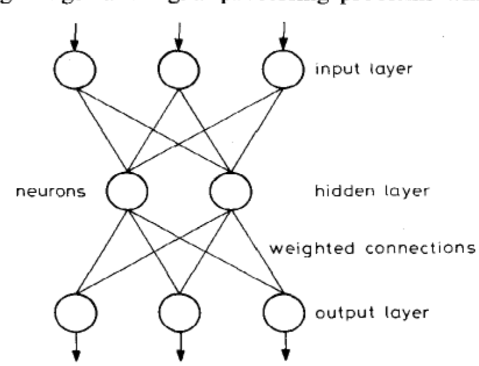

Fig. 1 *Three-layer feedforward neural network*

not readily tractable using rule-based conventional classifiers. Although neural networks offer a new approach to existing imaging problems, no theories exist as to whether a network can solve any given task. Some identification tasks are easily solved by simple networks employing short training times; higher orders of image complexity require larger networks and longer training times. BROWN (1989) gives an example of this in a simple application of neural networks to medical imaging. Choosing the correct presentation strategy and network architecture still relies on a certain amount of trial and error on the part of the researcher. This problem is addressed in the discussion section.

Neural networks can be divided into two main classes, those employing supervised learning and those with unsupervised learning. Hybrid systems using both strategies have also been developed.

1.1 *Supervised learning*

For a network to be trained using a supervised learning system a 'training' data set of example inputs and their corresponding desired outputs is required. The network weights, which will eventually store the learnt patterns, are initially set to small random values. During learning the example inputs are presented to the network and the resultant and desired outputs are compared. An error term is then calculated which is used to alter network weights. One example of a supervised learning algorithm is back-propagation (RUMELHART *et al.*, 1986); this is dealt with in some depth here as many of the applications reviewed use this learning scheme. Back-propagation networks use a sigmoid transfer function rather than the simpler threshold type function, which introduces two important properties. First, the sigmoid is nonlinear, allowing the network to perform complex mappings of input to output vector spaces, and secondly it is continuous and differentiable, which allows the gradient of the error to be used in updating the weights. Using the back-propagation algorithm it is possible to train networks containing hidden layers and so perform complex nonlinear mappings between the input and output domains, as opposed to the simpler linear mappings possible with perceptron networks. Networks of this type are sometimes referred to as 'multilayer perceptrons employing back-propagation learning'.

In training, the input example is presented to the network and the output activations are calculated using the existing network weights. An error term, based on the difference between the calculated and desired output, is then propagated back through the network to calculate changes of the interconnection weights between each layer. By repeatedly presenting the network with examples and changing the weights to minimise the error term, the network is trained to give the desired output for a given class of input. Weight changes can be carried out after each example or after each presentation of the complete training set, sometimes called cumulative back-propagation. By plotting the error term against the number of presentations a 'learning curve' can be produced, which ideally converges to zero as the number of presentations increases.

Adaptive linear combiners (ALCs) are much less complex than back-propagation networks, but employ much the same strategies (WIDROW and WINTER, 1988). The output of an adaptive linear element ('adaline') is a linear combination of the weighted sum of the inputs. No sigmoid transfer function is used. An error term is then calculated and used to modify the network weights. The weighted sum is fed through a thresholding unit to give a binary output. Adaptive filters use an ALC with inputs from time-sequenced signal samples. The ALC output is a linear combination of both current and past input signal where the weights can be trained to provide the desired output signal.

Associative memories are another form of supervised learning network. The network consists of two sets of neurons, representing input and output fields. All neurons in the input field are connected to all neurons in the output field. By adjusting the weights between the neurons it is possible to store exemplar input patterns and their output classes in the network. Subsequent input patterns will produce outputs corresponding to the closest exemplar class. Autoassociative memories are trained to produce an output identical to the input, so when they are fed with noisy or incomplete data they make an informed 'guess' of the missing original data, hence effectively yielding a reduction in the noise level.

1.2 *Unsupervised learning*

Unsupervised learning, as employed by KOHONEN (1990) and Hopfield (HOPFIELD and TANK, 1986) networks for example, requires no exemplar outputs. The unsupervised network simply consists of input and hidden neurons. The network learns by associating different input pattern types with different clusters of hidden nodes. When trained, different groups of hidden neurons respond to different classes of input patterns. Adaptive resonance theory (ART) (CARPENTER and GROSSBERG, 1987) provides another form of unsupervised learning network. Input patterns are added to existing feature clusters or used to form a new cluster if the 'Hamming' distance from existing clusters is too great; the Hamming distance is defined as the number of bits in the first pattern that do not match corresponding bits in the second pattern. ART-based systems are also able to learn new examples after the initial training has been completed. This is not possible on most systems as network weights become contaminated by further training only on the subset of new examples rather than the whole example set together.

1.3 *Hybrid systems*

Counter-propagation networks are examples of network architecture which use both supervised and unsupervised learning (HECHT-NEILSEN, 1987). The network consists of three layers and is trained in two stages. First the Kohonen unsupervised learning algorithm is employed to adjust the weights between the input and hidden layers. During the second phase of training the Grossberg Outstar supervised learning method is used to set the weights between the hidden layer and the output. Different input patterns are presented to the network and activations of the hidden-layer neurons are calculated, the highest activated neuron is then set to unity and the rest to zero. The Outstar algorithm then adjusts the weights by an amount proportional to the difference between the actual and desired output. This whole process can be thought of as employing unsupervised learning to produce an internal set of clusters in the hidden layer and then employing supervised learning to map those cluster patterns onto the desired output. Counter-propagation networks learn rapidly but are not always able to learn complex mappings because the hidden layer can only have one non-zero output at any one time.

1.4 *Hardware/software implementation*

Neural networks can be implemented in software or hardware. Software implementations are generally slower, requiring large computer facilities to train networks for real problems, but usually offer more flexibility in choice of network architecture and learning algorithm. There are now several software packages available for PCs and workstations; some of these are reviewed by BEARD (1990). In addition public domain software is also available for some systems. Existing hardware-based networks are faster than readily available software packages by several orders of magnitude (HUANG and KUNG, 1989). Very-large-scale integrated circuits are now being used for network implementation, and research into optical neurocomputers is also a growing research field. Both these approaches are expected to produce systems that are much more powerful than existing hardware implementations. However, a halfway house is provided by the availability of suitable PC cards and vector processor accelerators for software-based network simulations, which allow network training to be carried out on small low-cost systems in a realistic timescale.

2 Applications

Neural networks are rapidly finding new application areas within medical science. In addition to widespread application in diagnosis, much development work is being undertaken in image processing and interpretation as well as signal processing and analysis, and these are the central topics reviewed here. Computer-aided diagnosis, the other main field of neural network applications in medical science, is comprehensively reviewed by STUBBS (1990). The applications discussed here fall into three broad categories as follows:

Interpretation: data are collected and presented to the network; the network produces an output which can be directly interpreted in terms of variables which describe the state of the system.

Enhancement: data are collected and presented to the network; the network produces an output which can be used to enhance the image or signal, thus aiding interpretation. Enhancement covers both the creation of new information to be displayed in conjunction with the original data, and reconstruction from noisy data.

Compression: using neural networks to compress data for transmission or storage; once taught the network can represent all the information in an image as the activation levels of its hidden layer neurons. If there are fewer hidden layer neurons than inputs, the image is compressed.

Development of neural networks is already taking place in a wide range of modalities found in the analysis of *in vivo* medical data. These can be grouped under the following general areas. First, there are the long-established 'projected image' techniques such as X-ray radiographs and lung scintigrams. Secondly, there are echo techniques, mainly represented by ultrasound scans. Thirdly, tomographic methods, e.g. computer tomography (CT) and magnetic resonance imaging (MRI), which now include electrical impedance tomography (EIT), offer low resolution, low cost and continuous monitoring. Lastly there are diagnostic techniques such as ECG, EEG and EMG which rely on bioelectric (or biomagnetic) changes with time. These diagnostic techniques are quite distinct from those reviewed by STUBBS (1990) as the neural network is directly involved in the data acquisition, rather than relying on data collected by the physician.

2.1 *X-ray imaging*

The Craniofacial Research Instrumentation Laboratory, University of California, USA, has undertaken a preliminary investigation into the use of a neural network for craniofacial image analysis (BAKER *et al.*, 1989). Existing systems used techniques such as template matching or feature observation comparisons to pick out anatomical landmarks or implanted markers. Both have serious disadvantages. Template matching is sensitive to image distortions due to rotations, translations and scale change. Feature comparison requires careful selection of important characteristics. Neural networks are thought to be a promising avenue of research. No actual results are given, as the work seems to be in its early stages. Clearly there is room for such work although the images are less complex than those provided by MR imaging.

In addition to this application radiological diagnosis of neonatal chest radiographs has been achieved using neural networks (GROSS *et al.*, 1990). At this stage the network inputs, in the form of abnormalities in the radiograph, are provided by a radiologist. This application is diagnostic, as the network relies on inputs provided by a physician. A fully automated system might be possible at a later date. Researchers from the CRIN/INRIA, France, have developed a connectionist system which may be useful for tasks such as X-ray image interpretation (ALEXANDRE *et al.*, 1990). Their system, which they call a 'cortical column', has limited connectivity but includes feedforward, feedback and intra-areal connections (GUYOT *et al.*, 1989). The main part of their publication describes their model, followed by a brief mention of an application to the interpretation of X-ray images of vertebra. The work supports the notion that complex systems based on neurobiological theory can help in the construction of applied neural network systems. Work reported by CONRATH *et al.* (1989) on ultrasound tomography image reconstruction is also relevant to X-ray tomography reconstruction. This work has been included in the ultrasound section.

2.2 *Lung scintigrams*

The University of Warwick and the Walsgrave Hospital, Coventry, UK, have used neural networks to classify lung scintigrams (ANTHONY *et al.*, 1990). The network was used to determine the angle at which the image was taken and whether the image was of perfusion or ventilation. The system was also used to identify pulmonary embolus (PE) and chronic obstructed airways disease (COAD) with partial success. Previous work had shown that providing the network with raw scintigram data produced unsatisfactory results, the system not even being able to distinguish between different image viewpoints. A multilayer perceptron network, consisting of 128 input, eight output and a varied number of hidden layer neurons was used. It was trained using the back-propagation algorithm, and was used in conjunction with 'principal component analysis' (PCA) (JOLLIFFE, 1986) to preprocess the raw image data (BARBER, 1980). Results for unseen data (i.e. images not included in the training set) were not entirely satisfactory. The viewpoint and perfusion/ventilation images could be identified correctly in most cases. However, this information is already available to the physician. Correct classification of PE and COAD cases, the main aim of the research, was poor. A greater number of training images and a reduction in image noise level were thought likely to improve results. In addition the number of inputs was thought to be too few to provide sufficient information for correct classification.

Further work has been carried out using neural networks to clean up segments of images. Sections of images were standardised using a medial filter technique and used to train the network. The network performed better than a method based on the means of segments when used to provide values for missing pixels. It was shown that the number of training images should be at least of the order of the number of weighted connections in the network, to force the network to generalise about its stored information. The Coventry-based group also looked at data compression on scintigram images using neural networks and PCA (ANTHONY *et al.*, 1989). They concluded that PCA was superior, producing a better quality image in a shorter time.

2.3 *Ultrasound imaging*

A neural network has been used to estimate the ejection fraction of the heart from ultrasound Doppler signals by researchers at the University of Oklahoma and the Case Western Reserve University, USA (KARKHANIS *et al.*, 1990). Conventional methods for estimating the ejection fraction from Doppler data use a simplified model of the heart and relate peak flow acceleration to the ejection fraction. The networks all had 25 input neurons representing normalised points from a 60-point power spectrum of the Doppler pulse, as well as the area under the curve and the DC component. Two to six neurons were used in the hidden layer, and one output neuron represented the scaled ejection fraction.

Several different variations on the back-propagation learning algorithm were investigated: standard and cumulative back-propagation, cumulative back-propagation with a linear output layer transfer function and simulated annealing. Four hidden layer neurons, in conjunction with the back-propagation algorithm, were found to give best results, achieving a correlation of 97 per cent on training data. Testing with unseen data proved to be far less satisfactory, with no significant correlation between the network output and the measured ejection fraction. This was thought to be due to the high level of noise in the rather small training and testing data sets. However, it was found that the network placed more emphasis on the area and DC component inputs than on the spectral components.

Researchers at Cornell Medical College and the Polytechnic University, USA, have implemented a neural network system for recognition of tumours in ultrasound images of the eye (SILVERMAN and NOETZEL, 1990). Back-propagation networks were taught to recognise tumours and classify them by type, the eventual aim of the work being an automatic location and classification system. The network for tumour location used a 9×9 grid of input pixels, two hidden layers of ten neurons each and a 9×9 grid of output pixels. The network was trained using subregions from the original 128^2 image and a corresponding image containing tags for tumour pixels. Drawing the subregions from various parts of the main image ensured positional invariance. To ensure scale invariance the network training input was varied in magnification, comprising 9×9 pixel blocks extracted from 128^2, 64^2 and 32^2 size images each centred on the same central point. No quantitative results are given for this part of the application, but after output thresholding the network did select tumour regions successfully on unseen images.

Automatic tumour classification was investigated using five parameters: tumour height, base width, and from the backscatter spectrum of the tumour interior the spectral slope, intercept and standard error. Several network architectures were investigated with none, one or two hidden layers of various sizes. Four output nodes were used to represent the diagnostic classes: spindle cell melanoma, mixed/epithelioid melanoma and metastatic carcinoma. The network with one hidden layer of four neurons achieved a successful classification on unseen data of 78 per cent and 91 per cent without the melanoma subtype classification. This classification was repeated using discriminant analysis, which was found to give slightly poorer results, 72 per cent and 85 per cent, respectively.

By combining the two networks described above an automated tumour classification system was produced. The tumour was first located and then the inner half of the tagged pixels was used for spectral analysis. Eight out of ten of the tumours tested were given the same classification by the network as by a human expert. It was not possible to obtain an absolute accuracy because these cases were not histologically confirmed. The system was highly CPU intensive requiring several hours on a Cray-XMP supercomputer. Once trained the network could be implemented in hardware with image subregions processed in parallel.

DAPONTE and SHERMAN (1991) from the Southern Connecticut State University and the University of Bridgeport, USA, have used networks for texture classification of ultrasound images of the liver. 64×64 pixel regions of interest were chosen to be as close to the organ centre as possible while avoiding major blood vessels. The system was designed to distinguish between images exhibiting signs of liver paremchyma and normal images. Two commercially available software packages were used for network implementation, the Nestor Network Development Software running a restricted coulomb energy network and NeuralWare Professional II running a counter-propagation network: both systems were implemented on a PC. The two networks, respectively, achieved classification accuracies of 90·6 per cent and 81·2 per cent on unseen test cases. The best statistical discriminant methods give classifications to an 80 per cent level; the coulomb energy network was therefore deemed to be significantly better.

Neural network applications in medical ultrasound imaging have also been investigated by Siemens Ultrasound & Bio-Imaging Research Inc. (NIKOONAHAD and LIU, 1990). Ultrasound imaging systems usually use beamformers which are designed on the assumption of a constant sound velocity in the tissue. However, variations in the velocity of sound in the body wall tissue may be as large as 8 per cent, which generates noise from the variation in echo delay across the array, producing a degradation of the image quality. The signal-enhancement system described estimates the arrival time for each channel and uses these as inputs to the neural network.

The network output is a delay function for a uniform medium. The difference between the original delay and the output delay function is used as a control signal for the delay preprocessor. The preprocessor shifts the echoes by the correct amount for all channels. The corrected data are then fed into the beamformer, which creates the image to be displayed. A fully connected network employing back-propagation learning was used. The network comprised 64 input and 64 output neurons, one per channel, and two hidden layers with eight neurons in each. The network was trained using a software simulation of the proposed system, covering depths varying from 10 to 57 mm. The delay curvature changes more rapidly close to the array, requiring closer spacing of the training depths. The network produced satisfactory results at all depths up to 240 mm, and a substantial reduction in the sidelobe level of the image point-spread-function was observed. A hardware implementation and further optimisation of this system are also under investigation.

Researchers at the University of Illinois, USA, have used neural networks for the reconstruction of an ultrasound tomograph (CONRATH *et al.*, 1989). A three-layer back-propagation network was used for image reconstruction. The input layer was fed the real and imaginary pressures from 16 transducers around a test cylinder of 75 mm radius. The output was specified to be the radius of a nonattenuating test cylinder inside the tank. Various numbers of hidden neurons were tried, with 13 being found to give best results. The network was trained for test objects of 0·454–0·754 mm radius at a fixed position within the tank. No tests on unseen data were carried out, but the network learnt to recognise the radius of test cylinders, to an RMS accuracy of 0·008 mm, over the training set.

If neural networks could be shown to offer advantages for the solution of inverse problems of this type there would be great scope for their use in the medical physics field, not only in ultrasound and X-ray tomography but also in the nonlinear processing required for electrical impedance tomography. The case remains to be made in the face of already well established and understood back-projection techniques which can be implemented in real time on existing computer systems. The neural network approach would require very large networks to produce full image reconstruction, Conrath *et al.* suggest 10^8 elements. Although creating a network of this size might be possible on optical or VLSI network implementations, training such a network using currently available learning algorithms would be extremely difficult.

Neural networks show much promise in the area of adaptive imaging systems. An ambitious nonmedical example, using a supervised network, is the adaptive optical system for the Multiple Mirror Telescope in Arizona (ANGEL *et al.*, 1990) which has provided encouraging preliminary results (SANDER *et al.*, 1991). Parallels can be drawn between this work and the adaptive focusing through the body wall discussed by NIKOONAHAD and LIU (1990).

2.4 *MR, CT and PET imaging*

Interpreting the complex images produced by magnetic resonance is an area in which neural networks are finding several applications. Currently images are interpreted visually by skilled physicians with little availability of quantitative tools; for example it would be helpful to be able to readily calculate volumes and surface areas of features. In addition automatic tissue typing would allow faster image interpretation and reduce the amount of manual work required to use images.

Applications of neural networks to these imaging modalities fall into four broad categories: pixel classification, feature extraction/enhancement, automated diagnosis from image parameters and image compression. These are dealt with in order and discussed at the end of each section.

2.4.1 *Pixel classification.* Groups from the Vanderbilt University and Medical Center, Nashville, USA, have used neural networks to enhance MR images in two dimensions using all the readily available parameters: T_1, T_2 and proton density. The authors term this multispectral image analysis, in analogy to LANDSAT image processing (ÖZKAN *et al.*, 1990). An interactive environment provided by a Sun workstation was used to identify tissue types by considering the relative intensities of three parameters: the longitudinal relaxation time T_1, the transverse relaxation time T_2 and the proton density. The contribution of these three parameters to the overall pixel intensity is normalised with respect to the overall image parameters: the pulse echo time and the pulse repetition time. Using just T_2-weighted and proton density-weighted images it has proved possible to identify several types of tissue: background, fat and skin, cerebrospinal fluid (CSF), normal white and grey matter, white matter lesions and other lesions (an infarct for example). White and grey matter were not identified separately as no T_1-weighted images were available.

The selection of the learning data for training the neural network was made well within areas of a given tissue type to avoid incorrect classification, and to avoid biasing equal sized regions were used for each tissue type. The network comprised two input neurons, each assigned to the pixel intensity on one of the images, six output neurons, one for each tissue type, and up to a maximum of 15 neurons in the hidden layer. More than one hidden layer was also tried, but with little improvement in the results. The network was trained using the back-propagation algorithm over 50 cycles. Increasing the number of presentations of the data and increasing the number of hidden neurons led to an increase in classification accuracy, which levelled off at ten hidden neurons and 100 presentations.

The authors constructed a 'confusion matrix' with matrix elements C_{ij} representing the probability of classifying tissue type i as type j. A comparison was made between the neural network and a maximum likelihood classifier (Bayesian classifier): an overall classification accuracy for both techniques of over 90 per cent was obtained, the maximum likelihood classifier achieving slightly better results by a few per cent on some tissue types. Inspection of images showed that the neural network produced better delineation of tissue regions with fewer isolated points than the Bayesian classifier. Location and extent of the white matter lesions also corresponded well to expert opinion. Further corroboration of these classification techniques is already being undertaken, in the form of comparisons of classifier results and expert classification of tissue regions, without prior knowledge of case histories.

Three-dimensional image segmentation was also investigated using consecutive MR slices with regions of interest drawn on each slice (DAWANT *et al.*, 1990). As already observed in the two-dimensional case neural networks produced visually superior images and the confusion matrices were not significantly different from the Bayesian classifier. The three-dimensional segmentations were an improvement on those obtained in two dimensions for both types of classifier.

It should be emphasised that the network was taught using manually selected tissue regions and the assignment of a specific tissue type to such a region is a matter of judgment. Any erroneous region definition will lead to errors in network results at a later stage. This is minimised by using small regions well away from tissue type boundaries. To a certain extent even using training examples provided by experts is unsatisfactory, as the network will simply produce what the physicians want to see, rather than what might actually be present. This dilemma can only be resolved by histological studies. Neural networks do however seem to offer an alternative to existing image classification techniques, which are reviewed by SCHERTZ *et al.* (1985) and also by ALAUX and RINCK (1989). The work of DAWANT *et al.* (1990) represents the main work in the area of pixel classification, with other authors reviewed here adopting a similar approach.

Further work on tissue classification on MR images has been undertaken at the University of South Florida, USA (Schellenberg *et al.*, 1990). Multispectral MR images of gel and biological phantoms, and a human volunteer, were obtained for T_1, T_2 and proton density modalities. A comparison was made between a four-layer neural network, trained using back-propagation, and a maximum likelihood classifier. For training, regions of interest were again chosen to be well away from boundaries. Confusion matrices showed little difference between the two methods, both giving nearly perfect classifications with only a few erroneous pixels out of over 1000. For the head images the neural network seemed to produce better definition of edge pixels. This work is similar to the two-dimensional image segmentation of Dawant.

A different approach to the problem of MR image segmentation is being undertaken at King's College, London, UK (Toulson and Boyce, 1991), using only one imaging modality. The inputs to the network are generated by a function which is sensitive to the contents of a group of neighbouring pixels. This system allows pixels to be classified in context, rather than just using the intensity value. The network is trained using manually segmented images and so is prone to the same potential errors as the previous system. The network is, however, said to reduce the image noise especially in the skin/scalp region of the image.

All three of these approaches produce classification levels comparable with those achieved by conventional classifiers. There seems to be general agreement between the different research groups that image segmentation is improved, especially in boundary regions. However, no quantitative evidence of this improvement has yet been produced. A combination of the two basic strategies: the multispectral approach of Dawant *et al.* and Schellenberg *et al.*, and the neighbouring-pixel approach of Toulson and Boyce, may lead to further improved results.

2.4.2 *Feature extraction/enhancement.* The Department of Electrical Engineering at Linköping University, Sweden, has used neural networks to identify blood vessels in three-dimensional MR images by employing a special example of group-theoretic filters called rotation invariant operators (Grönqvist and Lenz, 1989). The three-dimensional image is first preprocessed, using the rotationally invariant operators for, say, line or edge attributes, to generate feature vectors for each pixel which contain information about the contents of the surrounding pixels. The length of each feature vector then specifies the attribute, edge or line feature, say, which dominates in that region of the image.

A fully connected three-layer neural network was used, trained using the back-propagation algorithm. The training used sample feature vectors and their corresponding classes as input and output data. Two MR images, each consisting of 12 layers of 256×256 12-bit pixels, were pointwise subtracted and thresholded to extract all points belonging to blood vessels. This resulting image was used as a target pattern for the network. One MR image was then filtered with three edge and five line filters, to provide feature vectors. Before training the network the user has to mark a few blood vessels and uninteresting areas in a connected region on the image. Thirteen feature vectors were used to make up the training set.

The trained neural network can be used, after the feature preprocessing, to detect complex structures of arbitrary dimensions. However, in common with conventional classifiers, the system was sensitive to the training set used. Among other pattern-recognition problems of medical importance the work is to be extended to detect cancerous regions within MR images.

The University of Colorado, USA, is using a neural network for lesion detection in MR images of the head (Raff and Newman, 1990). T_1 and T_2 images were used in conjunction with a 'novelty' filter network, a type of autoassociative memory. The novelty filter was trained to recognise the characteristics of normal tissue and filter these from other image regions, thus highlighting any abnormal tissue regions such as lesions. T_1 and T_2 images from 16 patients were used to form a test data set. Regions of normal brain tissue, 32×32 pixels in size, were extracted from each image to form a set of 1024 element training vectors (no other details of the form of network input or preprocessing were given). When the novelty filter was applied to an image of the brain with MS lesions two lesions were clearly identifiable. Such a system shows promise for lesion size quantification and time evolution of diseases. This work contains only preliminary results and includes no quantitative data on tissue classification. Further work to generate statistical results on the accuracy of image classification, as provided by other workers in the field, is still required.

The Boston University Centre for Adaptive Systems, USA, are using a neural network for highlighting features within MR images of the brain (Lehar *et al.*, 1990). They have used a boundary contour system (BCS) and feature contour system (FCS) developed by Grossberg and Mingolla and Grossberg and Todorovic (Grossberg and Mingolla, 1987) which employs adaptive resonance theory (Carpenter and Grossberg, 1987). This is a highly specialised system modelled far more closely on actual brain structure than is found in most other network applications. Example images are given for both BCS and FCS image enhancement. Both techniques seem to improve the image quality. No quantitative results are given, and would be difficult to derive, so assessment of the system is difficult. Accurate edge and feature detection of this type could prove useful for three-dimensional image reconstruction from multiple MR slices where accurate boundary definition is important for reliable reconstruction.

The Biomedical Engineering Department at the University of Virginia, USA, has used neural networks for translation invariant aorta segmentation (Katz and Merickel, 1989). Eight different slice positions for T_1, T_2 and proton density images were taken from a patient with advanced atherosclerosis which was confirmed by prior angiography. The 128×128 pixel images were reduced to 32×32 averaged pixels to limit the computational load. These were then manually tagged to produce training data. A multilayer back-propagation network with either one or two hidden layers was used for classification. To ensure translational invariant recognition the network input was restricted to a receptive field which could be moved around the image. The network therefore had $3n^2$ inputs, where n is the number of pixels along the side of the field of interest. n was varied from 1 to 7 pixels. The network was trained on random regions selected from the image, biased to account for the relatively small size of the aorta region. By plotting the components of the three image modalities it was shown that there was intermixing of aorta and non-aorta pixels in the input space, and that intensities vary even in images of the same subject. Various different receptive field sizes were tried.

A 5×5 receptive field was shown to be able to compensate for intermixing by using contextual image information. An overall segmentation accuracy of over 95 per cent was achieved by a network with a 5×5 pixel receptive field (75 input neurons), 150 hidden neurons and one output neuron. The main problem with this technique is the intrapatient, and to a greater degree interpatient, image variability, as could be seen in the variation of intensities between images of the same patient. The averaging used in the image reduction tends to make segmentation at tissue boundaries difficult. Further work using median-based methods to reduce images is being investigated.

The techniques of feature extraction employ a wide range of network architectures. Although back-propagation networks are used by Grönqvist and Lenz and Katz and Merickel, other networks are achieving successful image segmentations with novel architectures. Like conventional feature extraction techniques, neural networks trained to perform feature-related tasks are restricted to the particular features on which the network was trained. This limitation is not a problem provided that the network can generalise to other unseen images from different patients. If recognition of more than one feature is required then multiple networks could be trained, each to recognise a separate feature.

2.4.3 *Automated diagnosis.* Automated diagnosis of Alzheimer's Disease (AD) from PET scans has been undertaken by researchers at the University of Miami, USA (KIPPENHAN and NAGEL, 1990). A back-propagation network was used with input values representing variations from the mean of the cerebral glucose metabolism in the eight lobes of the brain: frontal, parietal, temporal and occipital (left and right). The network had eight inputs, two hidden neurons and two outputs representing normal and AD. Data sets from 15 normal and 11 AD subjects were used for training, and the same number of cases was reserved for testing the network. The network classified cases with a sensitivity and specificity in excess of 90 per cent, which is better than existing techniques, especially in terms of specificity.

Diagnostic parameter identification by neural networks from PET scans has been investigated by researchers at the University of Tennessee, USA (MILLER *et al.*, 1990). In practice PET scan data and physiological models can be used to determine diagnostic indicators. However, the non-linear minimisation method, normally used to identify key parameters, is too slow for use in a clinical environment. This application uses a back-propagation network instead of the conventional algorithm. The eventual aim of the work is to perform parameter identification at many points within image regions of interest. Initial results suggest that the neural network can produce similar results to the minimisation algorithm, with an average difference between the two techniques of less than 6 per cent, but in a shorter time. In a few cases the network produced erratic results. It is thought that this is due to the network attempting to extrapolate results outside the knowledge domain of the training set. Different forms of data preprocessing and training data selection are now being considered, to improve on these results.

2.4.4 *Image compression.* Researchers at the Universidade de Aveiro and the IST Lisbon, Portugal, have used neural networks to compress CT images (PINHO and ALMEIDA, 1991). The image of 256 × 256 pixels was split into 4 × 4 or 8 × 8 pixel blocks which overlapped by 25 per cent. These blocks constituted the input and output layers of the network, which contained one hidden layer. Compression is possible because, once taught, the weights of the neuron interconnections are fixed, and all the image data can be represented as the output values of the hidden layer neurons. A compression factor of 2 is achieved by teaching the network with half the number of neurons in the hidden layer as in the input or output layer. Compression factors of 2, 4, 5·33 and 8 were obtained by training networks with varying numbers of hidden nodes. Three images were used for training and a further nine for testing performance.

A compression factor of four was achieved without significant picture degradation. Even with a compression factor of eight, image quality was still deemed to be acceptable. Better image quality was obtained using the 8 × 8 pixel segmentation. Unfortunately no quantitative results for restored image quality were provided. Training the network with a back-propagation learning algorithm took between 0·5 and 80 CPU hours (using an HP650 9000 series computer). After the network was trained, compression and decompression times were short.

The system provides a useful means of reducing image data for transfer, so allowing images to be viewed away from the CT facility. A similar data compression system could be used to reduce MR images. With the increasing complexity of medical imaging systems, effective data reduction and storage will become more important.

2.5 *Retinal analysis*

The Departments of Ophthalmology and Electrical & Computer Engineering at the University of California, USA, have incorporated neural networks into a '*st*ructured *a*nalysis of the *re*tina' system, STARE (GOLDBAUM *et al.*, 1989). The STARE project aims to develop an image-understanding system to diagnose retinal diseases. Analysis of large numbers of colour photographs (or other stored images) and fluorescein angiograms of the ocular fundus will enable the STARE system to perform a variety of functions: to alert the ophthalmologist to images which require further attention, measure features within the image and calculate changes within images taken at different times. Neural networks were used in the latter stages of the system, after image processing, to produce a diagnosis based on the locations of all lesions present in the original image. The network consisted of 158 inputs, five neurons in the hidden layer and nine outputs, one for each category of diagnosis. The inputs corresponded to the number of lesions in 11 segmented regions of the image, for up to 24 types of lesion. Nine diagnostic cases were used to test the system: background diabetic retinopathy, proliferative diabetic retinopathy, macroaneurism, Coat's disease, central retinal artery occlusion, central retinal vein occlusion, branch retina vein occlusion, macular degeneration and normal retina.

Eleven examples of each diagnosis were collected, ten for learning and one of each case for testing. To optimise the results from this relatively small data set the test case was cycled giving 11 separate learning and testing sets. The combined statistical result for all these sets was an 83 per cent successful diagnosis from the images. A larger learning set is expected to improve this figure. Lesion position and classification for training is currently carried out manually, making production of a larger training set time-consuming.

2.6 *Electrical impedance tomography*

The Physics Department at Southampton University, UK, is investigating neural network applications within electrical impedance tomography (MILLER *et al.*, 1992). Neural networks are to be used to define lung and heart regions within the image of a thorax. By marking organ regions on 152 images taken over a 15 s sample time it is then possible to choose the appropriate pixel, or group of pixels, within each organ for producing an organ activity plot over time. This will facilitate the automated production of real-time organ activity plots, which will be essential for patient monitoring.

Neural networks have several possible applications within the EIT field. The low cost of existing EIT systems make them suitable for continuous patient monitoring. Neural networks may offer a means of automating some aspects of data analysis, as well as providing a signal as to the patient's overall wellbeing. Such networks, once trained, can be implemented on small, cheap, PC-based systems used for EIT image reconstruction. Inverse problems, although potentially solvable using neural networks, present a much greater challenge to existing network technology. The large networks and training sets needed to produce solutions to the inverse problem will require extensive computer facilities. Several conventional algorithms are available for producing images from EIT voltage data, and can be run in real time on a PC-based transputer system or similar hardware. Neural networks, being essentially nonlinear, appear to offer advantages in treating the EIT problem, which is itself highly nonlinear. However, the problem is so ill conditioned that ensuring correct network learning is a nontrivial problem.

2.7 *ECG signal processing*

There is much interest in the processing of ECG signals with neural networks. The work to date can be placed in three broad categories: signal classification, reconstruction and compression. The main thrust of existing research appears to be in the signal classification area. Various different groups have used neural networks for automated signal classification. Long-period ECG traces taken over 24 h contain of the order of 100 000 beats: automated beat classification is therefore desirable. Such classification can then be used to diagnose potential heart conditions.

2.7.1 *Classification.* A comparison of neural network and conventional statistical classifiers has been undertaken by researchers at the University of Leuven, Belgium (Bortolan *et al.*, 1990). A large ECG database containing signals from 3266 patients with normal QRS duration was used for the study. The network input included the ECG signal, parameterised into 37 standard variables (Blada *et al.*, 1977), and the sex and age of the patient. Various different network architectures were investigated, with one or two hidden layers of various sizes. A set of nonlinear transforms was used on the input parameters to map the input patterns onto a larger pattern space. Several different transforms were used. The most successful network employed the transform $\{x_i, \sin(\pi x_i), \cos(\pi x_i), \sin(2\pi x_i), \cos(2\pi x_i)\}$ for each of the 39 variables, giving a total of 195 input nodes. The network had no hidden layers and seven output nodes, representing normal, left and right ventricular hypertrophy, biventricular hypertrophy and anterior, inferior and combined myocardial infarction. The network compares well with the conventional linear discriminant analysis and multigroup logistic discriminant analysis methods, which give a total accuracy on a test set of 67·2 per cent and 66·3 per cent, respectively, compared with the network's 63·0 per cent. A partial accuracy figure (i.e. cases misclassified as similar waveform types) of 83·9 per cent for the logistic analysis is comparable to the 82·9 per cent achieved by the neural network.

A Group from the Institute for Technology Policy, New York, and Nestor Inc., USA, have collaborated to use a network to classify abnormal beats from long-period ECG traces, used to diagnose patients with suspected heart conditions (Carroll *et al.*, 1989). ECG beat data are coded as character strings. These are then classified using a conventional algorithm to compare input strings with known data. A similarity value is calculated according to character content and position. 230 beats equally weighted over 11 types was used to teach the network. A further 230 beats were then used for testing. A 1 per cent 'false-negative' rate (i.e. abnormal beats classified as normal) was obtained, and 74 per cent of abnormal beats were successfully identified. These figures are expected to improve for a larger training set. The system was also used to try to identify nine different abnormal beat types. This identification proved harder, with some beat types being classified correctly in only 10 per cent of cases.

Lee at Johns Hopkins University, USA, (Lee, 1989) has attempted to classify ECG signals according to beat type using a network in which the weights are constrained to give invariance under translation of the inputs. This is equivalent to a network with an additional layer connecting each time-series input to a set of consecutive time slices. By modifying the network in this way, the input can be made insensitive to rate and shape changes within different signal types. A back-propagation network was used with 51 inputs for the ECG signal, seven hidden layer neurons and one to three outputs. Networks with three outputs were trained to distinguish both normal/abnormal beats and ventricular tachycardia (VT) and ventricular fibrillation (VF). Data sets from 54 patients were used to train and test the network, which was found to be able to distinguish normal and VT beats with a degree of accuracy in excess of 95 per cent, for all but one case. The results for classifying VF beats were inconclusive, the network misclassifying them as normal or VT. Comparison of this work with that of Carroll *et al.* (1989) is difficult due to differences in the classification task attempted.

Additional work, reported by Xue *et al.* (1990) at the University of Wisconsin, USA, shows that the translationally invariant network is equivalent to training a network with unconstrained weights but with modified input data. The input time series is overplayed with time-shifted signals so that each network input sees a set of consecutive time slices. By analysing the principal components of the trained network weights (Xue *et al.*, 1989) it was possible to show that both methods reached a similar solution. The weight pattern of the shifted signal network showed weights distributed more randomly, which implies better generalisation properties for the network. The shifted signal network was also shown to converge faster than the translationally invariant network.

The network of Lee (Lee, 1989) employs little preprocessing of the ECG data, other than centring the waveforms. A similar classification scheme has been investigated at the University of Ottawa, Canada (Yeap *et al.*, 1990), but employing a preprocessor to extract key features from the ECG signal. Five features were used: QRS signal width, QRS amplitude, QRS offset, T-wave slope and prematurity; these were all normalised to lie between 0 and 1 for input to the network. A back-propagation network with five inputs, one for each feature, two hidden layers of 20 neurons each, and one output neuron for normal/abnormal beat indication was used. Training used four different example tapes from the American Heart Association database with a greater than 14 per cent occurrence of abnormal beats. The network was then tested using all the remaining database tapes, including those of ventricular tachycardia. The network showed a poor 44 per cent false positive rate and a 76 per cent false negative rate for unseen data. Much better false positive and false negative rates were obtained for the test data, 13 per cent and 27 per cent, respectively.

Researchers at the Technical University of Silesia, Poland (PIETKA, 1989), have also used five key ECG signal features as a form of network preprocessing: QRS duration, P-Q interval length, R-R rate and the delay in negative slope in V_1-V_2 and V_5-V_6. The network attempts to classify left and right bundle branch block (LBBB and RBBB), supraventricular premature beats (SVPBs) and ventricular beats (VBs), in addition to normal/abnormal beats. A back-propagation network with five inputs, two hidden layers with a variable number of nodes, and six outputs was trained on two types of training data. The first training set contained the features from 300 real ECG signals, and the second was generated by restricting the feature parameters for each signal type according to a cardiologist's opinion. The network learnt to classify unseen data to an accuracy of over 60 per cent for all six beat types, and to over 80 per cent in three of the beat types. An increase in accuracy was seen when the number of hidden nodes was increased from 12 and 4 to 18 and 6. This result seems to suggest that using signal features allows the network to produce better classifications. Which key features are chosen may also be important. It is difficult to make any such comparison as Pietka and Yeap use vastly differing data sets, both in size and range of signal types.

A back-propagation network has also been used to distinguish between supraventricular (SVT) and ventricular tachycardia (VT) exhibiting a wide QRS complex. The research group at the University of Limburg, The Netherlands (DASSEN *et al.*, 1990), used only four variables for beat classification: longest RS interval, presence/absence of AV dissociation and morphology criteria of leads V_1 and V_6. The resulting system recognised over 80 per cent of beats from an unseen data set.

TSAI *et al.* (1990) from the Christian University, Taiwan, employ even more extensive preprocessing of ECG signals before network classification. The system uses both low- and high-pass filters followed by a 128-point FFT to produce a 64-point power density spectrum which is then fed into the neural network. A back-propagation network with 64 inputs, 32 hidden layer neurons and five outputs was used. The outputs represent normal, LBBB, RBBB, premature ventrivular contraction and atrial premature beats. The results of this work are however unclear. The training set contained only 80 beats and classification figures given are for a test set which also included the training data. For this combined 'test' data an overall classification of over 80 per cent was obtained.

Abnormal ECG signal recognition has also been attempted using other network architectures. Bidirectional associative memory (BAM) networks have been used for signal classification by groups at the University of Oklahoma, USA (CHEUNG *et al.*, 1990). The BAM network consisted of 20 × 18 binary inputs arranged as voltage values normalised over 20 inputs for 18 sampling times. Fifteen binary value outputs were used. Six normal and six abnormal QRS patterns were manually selected and stored in the network along with their corresponding output patterns, 111100000000000 and 000000000001111, respectively. Individual beats were separated and aligned in time by detecting the large R-wave in the QRS pattern using a signal velocity trigger. This preprocessing, in addition to removing the DC component of the signal, greatly reduced the complexity of the recognition task. The system was able to classify real recorded ECG data to an accuracy of better than 99 per cent.

Using neural networks to track the prevailing heart rate from an ECG signal has been investigated by a group at the University of Alabama, USA (STRAND and JONES, 1990). Existing systems are slow to respond to rapid changes in heart rate. The network used consisted of a seven-neuron input layer, a 15-neuron hidden layer and one output neuron. The seven inputs, representing the seven previous R-wave intervals, and the output, representing the prevailing interval, both mapped onto the range (0, 1). Unlike most other neural networks reviewed here this net has feedback connections between the output and hidden neurons. The back-propagation algorithm was used for network training. The training set comprised 270 R-wave sequences with a minimum of 14 intervals measured from 19 ECGs. Heart rates in the range 55–140 beats $\min^{-1}$ were included plus arrhythmic conditions such as frequent ventricular premature beats and anomalous conditions such as abrupt changes of rate. For unseen test data the network tracked the prevailing heart rate to within 1 per cent for 84 per cent of cases and to within 4 per cent for all cases. These results were for a training and test set which did not include all types of R-wave interval pattern, nor did it include noise-induced anomalous R-waves. Training on a more extensive data set is required before the network can be incorporated into a real system.

The classification of multiple lead ECG images has been investigated by a group at the University of Amsterdam, The Netherlands (LINNENBANK *et al.*, 1990). 12 × 16 pixel maps were constructed from 62 lead ventricular tachycardia (VT) and QRS integral (QRSI) data. These were classified with respect to 18 reference patterns. Two networks were trained for classification, one using 192 inputs (12 × 16 pixels) containing the raw pixel image, and the other using 384 inputs representing two input images, one containing zero potential lines and the other a map of enhanced extremes. Both networks had an 18-neuron hidden layer and 18 output neurons, one for each reference pattern. Both networks were trained with and without noise in the training set. An Amiga 500 microcomputer was used for all network simulations. After training image recognition could be carried out in less than 1 s. Correct classifications of 65 per cent were obtained with the network employing feature inputs and noisy training data. A further 26 per cent of the data set were classified in adjacent classes.

Comparisons of the various different preprocessing strategies and network architectures used for ECG classification are hampered by the size and generality of the data used to train and test the network, the difficulty of the classification task attempted and to a lesser degree nonuniformity in the statistical tests used by researchers to check their results. Without direct comparisons using similar data sets, selecting the best system is problematic. Deciding which information is relevant for signal classification and how best to present this to the network is also difficult to decide, primarily due to lack of good comparisons. Various levels of preprocessing have been employed, in addition to the selection of key features for beat classification.

A different approach to heart monitoring has been taken by researchers at the University of Paderborn, Germany. A neural network has been used to estimate the duration of systolic and diastolic heart phases and to diagnose congenital and acquired heart disease from a phonographic input (BARSCHDORFF *et al.*, 1990).

2.7.2 *Reconstruction and compression.* Work has been reported by the Istituto per i Circuiti Elettronici, Genova, Italy, investigating biomedical neural network applications (ARRIGO *et al.*, 1989). A general implementation of a back-propagation network, on an FPS M64 Series Minisupercomputer, was used to investigate various biomedical applications such as diagnosis of dyspepsia, pH-metric analysis for ulceration diagnosis and ECG signal reconstruction. The two diagnostic applications both use separately collected symptom and pH-metric data to train and test the network. Both cases gave promising initial results. Larger training sets are to be used to improve these still further. The main goal is to extract rules for the relationship of symptoms to diagnosis.

Further work by the same group on the diagnosis of dyspepsia is reviewed by STUBBS (1990). The ECG signal reconstruction application used a completely connected four-layer network with 11 inputs, 11 neurons in the first hidden layer and six in the second, and a single output for the reconstructed signal. The 11 inputs were set to values of the ECG signal at increasing, equally spaced, time intervals T (where T is a multiple of the sample time interval T'). The output was the signal value at time T' after the end of the input sequence. The network was successfully taught to reconstruct the morphological characteristics of an ECG signal. No quantitative results of this signal-enhancement application were shown.

Another suggested application in ECG signal processing is the use of neural networks as adaptive filters for cancelling the maternal heartbeat when recording fetal ECGs (WIDROW and WINTER, 1988). Some preliminary results are given which show that such a system does clean the fetal signal. Ordinary filtering techniques are inadequate for this problem as the maternal and fetal ECG signals are spectrally overlapped.

The Nagayo Institute of Technology, Japan, has used neural networks to compress ECG signals on a Holter monitoring system (IWATA *et al.*, 1989). Existing Holter monitoring systems use a magnetic cassette to record two channels of patient ECG data over 24 h periods. The memory required to store such data is approximately 20 Mbyte. Existing memory cards only have a 256-512 kbyte capacity, and so a data compression of at least 1 : 30 is required. Eigenvector analysis can provide a compression of 1 : 50 but the computation of linear equations in 50 dimensions cannot be achieved fast enough with the low-level computational resources available on a portable system. The neural network compresses data because, once taught, the weights of the neuron interconnections are fixed, and all the image data can be represented as the output values of the hidden layer neurons.

Two neural networks are used because ECG signals change with body movements. One network compresses data by storing its hidden layer activations in memory and simultaneously tests itself by calculating the error between the network output and the original signal. When the network output error exceeds a specified threshold the original signal is stored in memory, rather than the hidden layer activations. The second network simply uses the input signal as training data input and output pairs, and modifies its weights accordingly using back-propagation. If the compression network error frequently exceeds the threshold then a new set of network weights from the training network is recorded in memory and copied onto the compression network. These new weights are used for signal compression. Signal data can be reconstructed at a later date by feeding the hidden layer activations into a network with the correct stored weights to produce an output signal. It was found that ECG signals with more beat types required more hidden units to store the signal correctly.

Using this system 24 h of ECG data can be compressed into 260 kbyte of memory. A neural network accelerator based on four digital signal processors was developed to allow network learning to be performed in real time on the ECG data. Data compression using neural networks therefore seems feasible for a system with less than 512 kbyte of memory.

2.8 *EMG and EEG signal processing*

A joint collaboration between the MDRTC Neurodiagnostic Unit, Cyprus, and Newcastle General Hospital, UK, has investigated the use of neural networks for classifying electromyograph signals (SCHIZAS *et al.*, 1989). Macroelectromyography (macro EMG) uses an electrode with a relatively large surface area to record the overall activity of a muscle motor unit. This overall activity during a weak muscle contraction is known as the macro motor unit potential (macro MUP). The reorganisation of the motor unit, which occurs as a result of disorders of the peripheral nerve and muscle, can be classified using the macro EMG technique. Classification is based on the amplitude of the macro MUP.

Data from 41 patients were analysed. Cases of Becker's muscular dystrophy, spinal muscular atrophy and motor neurone disease, in addition to normal cases, were included in the sample. Twenty macro MUPs from each patient were taken. Each macro MUP was preprocessed and values for maximum amplitude, area, average power and duration were calculated. These values for 20 cases were used to train two different three-layer perceptron networks, an 8 : 16 : 32 : 4 and an 80 : 120 : 240 : 4 network (figures denote neurons per layer, input layer first). The eight-input network used the mean and standard deviation of the four macro MUP signal parameters from all 20 signals. For the 80 input net each MUP signal was input separately.

The eight-input network was found to learn very slowly, the learning curve being noisy and not completely converging, even after several thousand propagations. The 80-input net converged after about 100 propagations. This result suggests that the reduced data did not contain information required by the network to converge. Even with the 80-input network, testing with previously unseen cases produced successful classification in at best 60 per cent of cases. This result suggests that, although the network had been taught to recognise the teaching data, it had not learnt to generalise to unseen signals. A larger data set would allow the network to generalise further and produce better results. However, the existing data took eight years to collect, so simply using a larger data set is not a quick solution.

The marked difference in performance of the two network architectures and data presentation strategies highlights the need for definite predictive theories for network performance. Arbitrary choices must be made about the network which can significantly affect performance. A reduced number of hidden-layer neurons might allow the network to generalise more effectively. Other types of network, principally Kohonen unsupervised learning networks, may be investigated by the group.

Classification of electromyogram (EMG) signals for prosthesis control is another area in which neural networks are finding applications. EMG signals from the amputee can be classified into different prosthesis movements. It has been shown that linear separation of the EMG signals for the different movement classes is not possible. Neural networks have therefore been tried to provide better signal classification, and for faster feature extraction.

KELLY *et al.* (1990) used a Hopfield network (HOPFIELD and TANK, 1986) for signal feature extraction and compared its performance with two existing techniques. Existing algorithms such as sequential least squares (SLS) in conjunction with the moving-average model (GRAUP, 1975) achieve a high degree of class separation, but require much user training and are too slow for clinical applications. The moving-average EMG signal model predicts the next signal value in a series as a linear sum of some of the previous values. The SLS algorithm is used to modify the summation coefficients to minimise the difference between the actual and predicted time series values.

A Hopfield network was trained to calculate the parameters for the moving-average EMG signal model. A male subject with a lower arm amputation just below the elbow was used for the investigation. Two Beckman electrodes positioned on the triceps brachii and biceps brachii were used to acquire training data. The Hopfield network was found to be faster than the SLS algorithm and could produce the same time series coefficients. The accuracy of the network solution could be increased by altering the convergence criterion ε, but only at the expense of computation time. The network was found to be five to ten times faster than the SLS algorithm, especially for large numbers of points in the time series. A comparison with a recursive least squares (RLS) algorithm was also carried out as the computational speed of the SLS method is known to be suboptimal. The network produced results in a comparable time to the RLS method for a time series of length 100 points, but was significantly faster for a 1000-point series.

A multilayer perceptron employing back-propagation learning was employed to classify EMG signals for four distinct arm functions. The first moving average coefficient and the signal power from a single EMG channel were used as inputs to a three-layer network with two inputs, four hidden-layer neurons and four output nodes. The signals showed good clustering within the feature space but not sufficient to be linearly separable. The network was shown to have produced decision boundaries which successfully classified the training data. A similar investigation of this application has been carried out by a group at the Amirkabir University of Technology, Iran (KHOSHABA *et al.*, 1990). Integral absolute value EMG signals from the biceps and triceps were combined with information given by a human supervisor to form a training set. The network was successfully trained to recognise different EMG signal types from the patient. Neither of these applications used mention the use of unseen data for testing the network classification.

Groups from the Johns Hopkins University and Hospital, USA, have used neural networks for electroencephalogram (EEG) waveform classification (EBERHART *et al.*, 1989). Currently system development is being carried out using CASENET, a flexible neural network simulation package. The long-term goal of this project is to produce a portable device for spike/seizure detection using low-cost hardware. Trained neurologists have some difficulty in recognising spikes. The six neurologists who annotated the recordings used to train the network were found to differ in their classification of some spike cases. Only spikes marked by four or more of the neurologists were considered to be 'true' spikes.

The initial aim was to produce a system that could identify 80 per cent of 'true' spikes. Various strategies for presenting single channel data to the network were considered. Scanning of the raw EEG data using a sliding window was rejected on the grounds of increased computational load during the learning phase: a larger training set would be required to account for the appearance of narrow spikes anywhere within the window. The training set could be reduced by centring the candidate spikes in a 240 ms window in the input field. The network, which used either 48 or 60 inputs depending on the sampling rate, employed a hidden layer of eight nodes and two output nodes. The outputs were used to indicate 'true' spikes and spikes not identified by any of the neurologists but flagged by the Johns Hopkins' Epilepsy Monitoring Unit software. A third output may be added at a later date, to indicate spikes marked by fewer than four of the neurologists.

A second method involved additional preprocessing to calculate nine spike parameters, which were then presented to the network for spike recognition. The increased computational overhead for preprocessing was offset by the reduced network size and training time.

Both methods produced systems of adequate performance for practical use, if extended to a multiple-channel system, and both were judged to be 'superior to any algorithm known to the authors for EEG spike detection'. An initial system has been developed to evaluate a minimum of four channels, three bipolar and one reference using 192 raw data inputs, 16 hidden nodes and two outputs. For a spike to be considered positive, it must appear on at least one of the bipolar channels and the reference channel. Accuracies of over 90 per cent for an unseen test data set have been achieved, a result well above the 80 per cent goal. Further work is being undertaken using the different network configurations available through CASENET. The system requires more refinement before being clinically useful.

A comparison of different methods of EEG signal classification for automatic sleep scoring has also been investigated by researchers from the Universities of Florida and Tennessee, USA, and the Universidade de Aveiro, Portugal (PRINCIPE *et al.*, 1989*b*). A rule-based expert system, a belief automaton, a system employing a probabilistic theory of evidence (PRINCIPE *et al.*, 1989*a*), and a neural network were all used to classify EEG data from patients ranging from 13 to 70 years of age. The neural network consisted of 24 input units containing EEG, EOG and EMG tokens for the current minute, four output units representing sleep stages: 0 (awake), 1/5, 2, and 3/4. A training set of six examples for each of the four outputs was used in conjunction with the back-propagation algorithm. The network produced comparable results to that of the two other systems, giving an average agreement with human experts of 85·2 per cent, as compared with 86·3 per cent for the expert system and 84·9 per cent for the belief automaton.

It is suggested that the similarity in performance, despite the wide range of classification strategies, is due to all the systems being of a reactive type. Reactive systems do not contain internal feedback loops, so that outputs can always be traced back to the initial input conditions. It is suggested that creative systems, containing feedback, may have higher information-processing qualities. Until such systems are developed more effort should be put into the development of better input descriptor signals.

A neural network employing spectral components of the EEG signal as inputs has been used for differentiating the effects of three benzodiazepines on sleep EEGs (GEVINS *et al.*, 1988). The main subject of the work is clinical, rather than neural network based. The authors conclude that neural networks have two main advantages over analysis of variance: they do not require the assumption of a multivariate normal distribution, and they can distinguish conditions with overlapping distributions and variances.

It is difficult to assess the performance of different systems for EEG signal classification as they are not attempting identical classification tasks. The use of neural networks for EEG signal classification obviously shows promise, although work in some areas may be hampered by lack of training and testing data. In addition, PRINCIPE *et al.* (1989*a*) suggest that classification performance has an upper limit due to the reactive nature of all the currently available systems. Initial investigations of neural networks of EMG signal classification for prosthesis control do however appear promising, as parallel aspects of the network are utilised to produce faster signal classification.

2.9 *Thermography and thermoregulation diagnostics*

A comparison of neural networks and conventional algorithmic techniques for extracting the characteristics of whiplash injury from thermographic images has been undertaken by groups at the University of Nevada, USA (EGBERT *et al.*, 1990). This work followed a previous study of neural networks for the evaluation of tumour growth in laboratory mice (EGBERT *et al.*, 1988). Thermography is used in the study of neurological and musculoskeletal disorders, providing skin temperature data and so giving information on heat loss and vascular and metabolic activity. Currently image interpretation is carried out by an expert physician.

Data were collected from two groups of subjects, with and without whiplash injuries to the neck. Twenty images from normal and 20 images from abnormal patients were obtained for training and testing. Thermographic images of whiplash injuries usually exhibit asymmetry and a number of hot spots. These features were used as the basis for diagnosis.

Substantial preprocessing of the data was carried out before attempting any classification. This reduced the problem domain from a 320×256 pixel image to a one-dimensional vector of fewer than 100 elements. The image was first filtered to remove any high-frequency thermal noise. Then all local maxima were detected and the pixels belonging to each maximum extracted as the principal image features. The image was considered as having polar co-ordinates and divided into $5°$ sections to $\pm 35°$ of the central axis. The total number of maxima pixels per $5°$ section and the average temperature difference of all pixels in paired sectors left and right of the $0°$ line were selected. The polar nature of the feature extraction and mask positioning make them translation and rotation invariant and scale invariant. The only part of this processing task not performed automatically is the alignment of the polar co-ordinate system.

Two different commercially available neural networks were used for image classification, and compared with conventional algorithmic classification techniques. Both neural networks employed back-propagation learning with various combinations of features used as inputs. Two output classes were used, flagging normal and abnormal images. The best results were obtained using both population of maxima and temperature difference from all 16 polar sectors, making an input vector of 32 elements. In each case half the available data were set aside for testing.

The Brainmaker software (California Scientific Software) achieved an overall correct classification of 90 per cent compared with 80 per cent for the ANZA software and co-processor (Hecht-Nielson Neurocomputers). This was thought to be due to the modifiable sigmoid function provided by Brainmaker. The results both compared well with a 75 per cent successful classification of a conventional nearest-neighbour classifier. In all cases, however, the classification of normal cases was significantly better than that for abnormal cases, usually by more than 20 per cent. Further work is in progress to expand the size of the training and testing database. Investigation into the use of an adaptive resonance network for examining data clustering is also being undertaken. Additional image features are also being developed to allow successful classification of more subtle image patterns.

Neural networks have also been applied to thermoregulation diagnostics by a group at the Medizinsche Universität zu Lübeck, Neumünster, Germany (PAUL and VON GOLDAMMER, 1990). The technique relies on detecting the shift in temperature, at certain points on the skin surface, induced by cold stress. These temperature changes can indicate changes in the underlying organs. The temperature of 60 points on the patient's body are measured before and after a 10 min period of unclothed exposure to a climate of 18°C and 60 per cent humidity. In a healthy patient face temperatures are higher and body temperatures are lower after exposure. A transputer implementation of a 10×10 Kohonen network was used for the study. The network used 120 inputs, one for each temperature measurement, and took 3·5 h to train on 10^5 examples. Initial results are described as interesting, but no specific details are given.

2.10 *Spirogram and CO_2 waveform interpretation*

Researchers at the North Carolina State University have used neural networks for spirogram interpretation (GAGE and MILLER, 1990). The standard parameters used for pulmonary function evaluation are chosen largely because of the ease of calculating them from the strip chart produced by most spirometers. Information contained in some of these parameters is thought to be largely redundant. A multilayer perceptron employing back-propagation learning was used to classify forced expired volume signals. Eighty normalised, logarithmically spaced points from the volume/time curves were used instead of the usual parameters, in addition to sex, age and height data. Two experiments were carried out to use the network to distinguish between normal, restricted and obstructed curves and normal and asthmatic curves. Both experiments achieved a successful classification of over 90 per cent. The classification accuracy was found to be a function of training set size, with larger training sets giving better results, as would be expected. In the second experiment the network performed comparably to a conventional classifier using the normal parameters.

The University of Arizona College of Medicine, USA, has used neural networks for classifying CO_2 waveforms (NAVABI *et al.*, 1990). This is a relatively simple task, and although it provides no new information to the anaesthesiologist, such automation may be useful when combined with other information in providing detection of successful intubation. Two three-layer back-propagation networks were used to classify CO_2 waveforms into spontaneous, mechanical and mechanical with attempted spontaneous breathing. The CO_2 waveform was first edge-filtered, as classification can be achieved by examining the leading and trailing edges of the waveform. Both networks had 50 inputs, six hidden layer neurons and one output. The first network was trained to detect mechanical and spontaneous ventilation from the 50 samples preceding the detection of the rising edge. The second network detected spontaneous breathing events from the 50 samples preceding the falling edge. Data from 17 surgical cases involving intubation were used to train and test the system. The system classified 93 per cent of breaths as spontaneous/mechanical correctly and detected 85 per cent of mechanical with attempted spontaneous breathing. The system is limited by noise, which affects the performance of the edge filter.

2.11 *Echocardiographic imaging*

Evaluation on two-dimensional echocardiographic images by a neural network has been investigated by researchers from the University of Toledo, USA, and the Academy of Medical Sciences, Beijing, China (CIOS *et al.*, 1990). Data from two separate regions of the myocardial tissue were obtained: the ventricular septum and the left ventricular posterior wall. For the first region six normal subjects, five with myocardial infarctions of the anterior ventricular septum (AVS) and seven with hypertrophic cardiomyopathy (HCM) were studied. For the second region five normal subjects and six subjects with myocardial infarctions of the posterior wall (PW) were considered. Neural networks were used to distinguish between these different cases.

One-hundred input neurons were used to define a 10×10 pixel region of interest within the image. The number of hidden neurons was found by trial and error to yield best results if 36 were used. The number of output neurons depended on the imaging region: for the ventricular septum region three were used and two for the left ventricular posterior wall. Two types of input layer were used, a fully connected layer and a patterned layer. The patterned layer used overlapping regions of the two-dimensional image as inputs to hidden layer neurons, providing information about the image structure and reducing the number of network weights.

Owing to the small amount of patient data different networks were trained using different portions of the data set for training and testing. The results were then averaged to produce an overall result. The fully connected network provided slightly better results than the patterned network, specificities of over 70 per cent being obtained for all diagnostic classes. Sensitivities of 71 per cent for HCM and 83 per cent for PW were also obtained. The sensitivity for AVS was 0 per cent. This was thought to be due to the fact that AVS is difficult to discern even by a trained cardiologist from a single two-dimensional section.

2.12 *Arterial pressure waveform analysis*

Researchers at the University of California, USA, have used neural networks for the classification of arterial pressure waveforms (SEBALD, 1989). Two neural network architectures were applied to waveform classification, back-propagation and competitive learning. The back-propagation network had 21 input neurons, 21 hidden neurons and one output. The competitive learning network had 21 input and output nodes. The back-propagation network was found to give superior classification performance, correctly classifying 100 per cent of the test data waveforms, compared with the competitive learning network's 93 per cent correct classification.

2.13 *Cineangiogram analysis*

An adaptive linear learning network has been applied to the processing of angiograms for the identification of coronary arteries, by researchers at the University of Rhode Island, USA (NEKOVEI and SUN, 1990). A two-part system was used, a variable threshold portion to calculate the threshold of each pixel from its neighbours and a single-layer adaptive network. The system was implemented on a SUN workstation in C and used a set of eight cineangiographic images digitised into 256×256 pixel frames. The system converged to a satisfactory segmentation after 10–30 iterations. Example images show good artery segmentation. No quantitative results are given and no images segmented by other techniques are shown, so assessment of such a system is difficult.

3 Discussion

One obvious hurdle for those wishing to use neural networks to solve image and signal processing problems is the choice of network. The multilayer perceptron, employing the supervised back-propagation learning algorithm, has already gained much ground in engineering and has been used in many of the applications reviewed here. A wide variety of other network architectures and learning paradigms exist (HINTON, 1989; ALEKSANDER and MORTON, 1990), some of which have been applied to medical imaging. Unsupervised networks, for example, do not require a set of ideal solutions for training, and have the advantage of not requiring a solution set, which must usually be created by an expert in the field.

In all cases the form in which data is presented to the network, and the network size, must still be specified. Currently there are no theories governing the best way of constructing a network for a given problem. Networks must be designed by trial and error: this empirical approach to network design is difficult to surmount. Network training times may be long, which precludes trying large numbers of networks and differing presentation strategies. Careful preprocessing of the input data can dramatically reduce the number of input nodes, and so reduce the network size and training times. Conversely, poor preprocessing can remove information required by the network to converge and so increase the training overhead. Increasing the number of input neurons may make the network learning task easier, but will also increase the number of network weights which must be set.

An optimum training set is thought to be one containing many more examples than the number of network weights, and so coding strategies must not increase the number of weights significantly if they are to be effective. Once the input and output layers have been defined, the number and size of the hidden layers must also be specified. Too few hidden nodes will inhibit learning, too many will allow the network to learn each training pattern individually and no generalisation will take place.

Again, no real theories exist as to the number of nodes most suited to a problem of any given complexity. However, new algorithms which examine the significance of each hidden neuron and connection weight, and add or delete neurons and connections accordingly, appear frequently in the literature. The network pruning method of KARNIN (1990) is a simple example of such an algorithm. This lack of theoretical background for neural networks, and the subsequent ad hoc approach to neural network application design, is mainly due to the emergent nature of the field. The success of networks at solving otherwise intractable problems in the field of artificial intelligence has led to their rapid deployment in a wide variety of application areas.

In some of the cases reviewed here the amount of data available for training and testing of the network was limited. This significantly degraded the results obtained. At least one-third of the data must be set aside for testing the taught network, further reducing the data available for teaching. This can be maximised by creating different training and testing sets from the available data and combining the results to give an overall assessment of the network performance. More efficient network paradigms can be used to produce improved results without the need to resort to further examination of patients to obtain more data. New training algorithms and paradigms frequently appear in the literature. There is much scope to use these within the medical field where training data is often limited.

It should also be noted that the training data must still form a representative sample of the set of all possible inputs if the network is to perform correctly. Supplying the network with reliable data for training also presents problems. Even if raw data are available, generating example answers with which to train the network presents problems. Often manually selected image regions or signal samples are used for training. This manual selection places a bias on any results obtained — the user has effectively taught the network what he or she wishes to see. Producing sets of desired outputs for training and testing supervised networks can also be unduly time consuming, often requiring the help of qualified experts. Unsupervised learning networks overcome some of these problems by not requiring an exemplar output during training.

To produce reliable results accurate data must be used to train the network. One obvious solution to the problem of restricted training and testing data is to create simulated data using either a computer based or physical model. Great care must be taken when modelling such data as any incorrect assumptions in the model will be reflected in the network's effectiveness on real data. Careful testing of the network performance when given real data will still be required. Such testing will necessarily involve a comparison of network classification with that of experts in the field.

Software-based networks allow more flexibility during the system design phase, but usually require large computer resources. Many of the applications reviewed are eventually to be used for patient monitoring, which precludes the use of centralised computer facilities. As smaller and more portable hardware-based neural networks become readily available such applications can be realised. PC-based neural network software packages utilising the Intel 1860 and other high-performance processors, in the form of PC cards, are already available. The processor card is used to accelerate the calculations involved in network simulation, and represents an intermediate stage in the progression from software to hardware-based systems. Such systems have the advantage of offering the flexibility of a software simulation, combined with the improved processing power usually associated with the hardware-based systems.

MR images show the greatest potential for neural network applications. Their complex nature, with several subimages for each slice, makes manual interpretation difficult. Basic tissue typing and feature identification have already been investigated. An image-reducing code, similar to that implemented for CT images, could also prove useful.

The other main application of networks is for signal calculation. Several applications involving the flagging of large numbers of signals are reviewed. Such work could allow networks to be used in online monitoring of patients in addition to fast collation of previously obtained results.

Neural networks have yet to be applied to sensor fusion within medical physics. Sensor fusion involves using the network to analyse data from several different monitoring units, and combining them to give an overall signal denoting the state of the patient. This technique represents an advance in signal classification as it provides information as to the overall wellbeing of the patient, as opposed to just one particular aspect of their health.

4 Conclusion

A comprehensive review of the increasing range of applications of neural networks to medical imaging and signal processing has been presented. The indications are that neural networks represent a major advance in the imaging and signal-processing fields, but with only the first stage of the necessary development effort having been completed to date. The principal problems which must be addressed when producing a complete network application are: collecting and classifying sufficient training and testing data, choosing a valid data presentation strategy and an appropriate network architecture, and implementing the network on a sufficiently powerful computer system to achieve net convergence in a reasonable time. With current advances in software and hardware none of these present insuperable problems.

5 Postscript

Subsequent to the completion of the main body of this review several other conferences of relevance have taken place: The International Joint Conference on Neural Networks (Seattle, Washington, USA, 8th-12th July 1991, and Singapore, 18th–21st Nov. 1991), the IEE 2nd International Conference on Artificial Neural Networks (18th–20th Nov. 1991, Bournemouth, UK), the 13th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (Orlando, Florida, USA, 31st Oct.–1st Nov. 1991) and the International Joint Conference on Neural Networks (Baltimore, Maryland, USA, 7th–11th June 1992).

Acknowledgments—This work was partially funded by a SERC studentship (ASM). The authors would like to acknowledge the *Neuron Digest* electronic mailboard's organisers and contributors, who helped to track down some of the current medical applications of neural networks. *Neuron Digest* is available from 'NEURON-REQUEST@CATTELL.PSYCH.UPENN.EDU'.

References

ALAUX, A. and RINCK, P. A. (1989) Multispectral analysis of magnetic resonance images: a comparison between supervised and unsupervised classification techniques. In Proc. Int. Symp. on Tissue Characterisation in Magnetic Resonance Imaging, European Soc. for Mag. Res. in Med. & Biol, ISBN 354051532, 19th–21st April, Wiesbaden, Germany, 165–169.

ALEKSANDER, I. and MORTON, H. (1990) *An introduction to neural computing*. Chapman & Hall, New York.

ALEXANDRE, F., GUYOT, F. and HATON, J.-P. (1990) A connectionist network with two complementary visual processing systems for X-ray image interpretation. In Proc Int. Neural Networks Conf., July, Paris, France, Part 1, 496–499.

ANGEL, J. R. P., WIZINOWICH, P., LLOYD-HART, M. and SANDLER, D. (1990) Adaptive optics for array telescopes using neural network techniques. *Nature*, **348**, 221–224.

ANTHONY, D. M., HINES, E. L., TAYLOR, D. and BARHAM, J. (1989) A study of data compression using neural networks and principal component analysis. Proc. IEE Colloq. Biomed. Appl. of Digital Sig. Proc., IEE Coll. Dig. 1989/144, Nov., London, UK, Paper 2.

ANTHONY, D. M., HINES, E. L., TAYLOR, D. and BARHAM, J. (1990) The use of neural networks in classifying lung scintigrams. Proc of Int. Neural Network Soc. Int. Neural Networks Conf., 17th–21st June, San Diego, California, USA, Part 1, 71–74.

ARRIGO, P., CORANA, A., GUILIANO, F., MARCONI, L., MORANDO, M., RIDELLA, S., ROLANDO, C. and SCALIA, F. (1989) Neural networks: computer simulation and biomedical applications. Proc. 2nd Italian Workshop on Parallel Architectures and Neural Networks. CAIANIELLO E. R. (Ed.), World Scientific, Societa Italiana Reti Neuroniche, ISBN 981020146X, 26th April, Vietri Sul Mare, Salerno, Italy, 205–212.

BAKER, B., CURRY, S. and BAUMRIND, S. (1989) A neural network method for solving pattern recognition problems in craniofacial X-ray image analysis. In Proc. 11th Ann. Int. Conf. IEEE Eng. in Med. & Biol. Soc., 9th–12th Nov., Seattle, Washington, USA, Section 5, 1646.

BARBER, D. C. (1980) The use of principal components in the quantitative analysis of gamma camera dynamic studies. *Phys. in Med. & Biol.*, **25,** 283–292.

BARSCHDORFF, D., ESTER, S., DORSEL, T. and MOST, E. (1990) A new phonographic technique for congenital and acquired heart disease using neural networks (in German, abstract in English). *Biomed. Technik*, **35,** 271–279.

BEALE, R. and JACKSON, T. (1990) *Neural computing: an introduction.* Adam Hilger, UK.

BEARD, N. (1990) Having a brainwave. *Personal Comput. World*, Feb., 186–190.

BLADA, R. A., DILLER, G., DEARDORFF, E., DOUE, J. and HSIEH, P. (1977) The HP ECG analysis program. In *Trends in computer processed electrocardiograms.* VAN BEMMEL, J. H. and WILLEMS, J. L. (Eds.), North-Holland, Amsterdam, 197–204.

BORTOLAN, G., DEGANI, R. and WILLEMS, J. L. (1990) Design of neural networks for classification of electrocardiographic signals. Proc. 12th Ann. Int. Conf. IEEE Eng. in Med. & Biol. Soc., 1st–4th Nov., Philadelphia, Pennsylvania, USA, Section 3, 1467–1468.

BROWN, D. G. (1989) Neural networks for higher order medical imaging tasks. Proc. Medinfo '89 and 6th Ann. Conf. on Medical Informatics, ISBN 0444881387, 16th–20th Oct. and 11th–15th Dec., Beijing and Singapore, 77–81.

CARPENTER, G. and GROSSBERG, S. (1987) A massively parallel architecture for a self-organising neural pattern recognition machine. *Computer Vision Graph. & Sig. Proc.*, **37,** 54–115.

CARROLL, T. O., VED, H. and REILLY, D. (1989) Neural network ECG analysis. Proc. Int. Joint Conf. on Neural Networks, (abstracts only). 18th–22nd June, Washington DC, USA, Part 2, 575.

CHEUNG, J. Y., HULL, S. S. Jr, YEO, C. S. Y. and KOHLI, P. (1990) Recognition of abnormal EKG signals by a neural network approach, *Microcomputer Appl.*, **10,** (2), 48–53.

CIOS, K. J., CHEN, K. and LANGENDERFER, R. A. (1990) Use of neural networks in detecting cardiac diseases from echocardiographic images. *IEEE Eng. in Med. & Biol. Mag.*, Sept., 58–60.

CLARK, J. W. (1991) Neural network modelling. *Phys. in Med. & Biol.*, **36,** 1259–1317.

CONRATH, B. C., DAFT, C. M. W. and O'BRIEN, W. D. Jr (1989) Applications of neural networks to ultrasound tomography. Proc. IEEE Ultrasonics Symp., Oct., Chicago, Illinois, USA, 1007–1010.

DAPONTE, J. S. and SHERMAN, P. (1991) Classification of ultrasonic image texture by statistical discrimination analysis of neural networks. *Computer Med. Imaging Graph.*, **15,** (1), 3–9.

DASSEN, W. R., MULLEREERS, R., SMEETS, J., DEN DULK, K., CRUZ, F., BRUGADA, P. and WELLENS, H. J. J. (1990) Self-learning neural networks in electrocardiography. *J. Electrocardiol.*, (US), **23,** Suppl., 200–202.

DAWANT, B. M., ÖZKAN, M., SPRENKELS, H. G., ARAMATA, H., KAWAMURA, K. and MARGOLIN, R. A. (1990) A neural network approach to magnetic resonance imaging tissue characterisation. In *Communication control and signal processing.* Proc. Bilkent Int. Conf. on New Trends in Comm. Control & Sig. Proc., ARIKAN, E. (Ed.), 2nd–5th July, Bilkent University, Ankara, Turkey, Elsevier, Amsterdam, Part 2, 1803–1809.

EBERHART, R. C., DOBBINS, R. E. and WEBBER, W. R. S. (1989) CASENET: a neural network tool for EEG waveform classification. Proc. IEEE Symp. Computer Based Medical Systems, 25th–27th June, Minneapolis, Minnesota, USA, 60–68.

EGBERT, D. D., RHODES, E. E. and GOODMAN, P. H. (1988) Preprocessing biomedical images for neurocomputer analysis. IEEE Int. Conf. on Neural Networks, July, San Diego, California, USA, Part 1, 561–568.

EGBERT, D. D., KABURLASOS, V. G. and GOODMAN, P. H. (1990) Neural network discrimination for subtle image patterns. IEEE Int. Joint Conf. on Neural Networks, 17th–21th June, San Diego, California, USA, Part 1, 517–524.

GAGE, H. D. and MILLER, T. K. (1990) Mapping neural networks for the analysis of the forced expired volume signal. Proc. 3rd Ann. Symp. on Computer-Based Medical Systems, June, Chapel Hill, North Carolina, USA, 367–373.

GEVINS, A. S., STONE, R. K. and RAGSDALE, S. D. (1988) Differentiating the effects of three benzodiazapines on non-REM sleep EEG spectra: a neural network pattern classification analysis. *Neuropsychobiol.*, **19,** (2), 108–115.

GOLDBAUM, M. H., KATZ, N. P., CHAUDHURI, S. and NELSON, M. (1989) Image understanding for automated retinal diagnosis. Proc. 13th Ann. Symp. Computer Appl. in Medical Care, Nov., Washington DC, USA, Chap. 5–8, 756–760.

GRAUP, D. (1975) Functional separation of EMG signals via ARMA identification methods for prosthesis control purposes. *IEEE Trans.*, **SMC-5,** 252–259.

GRÖNQVIST, A. and LENZ, R. (1989) Detection of blood vessels in 3-D MR-images. Proc. Int. Joint Conf. on Neural Networks, 18th–22nd June, Washington DC, USA, Part 1, 145–149.

GROSS, G. W., BOONE, J. M., GRECO-HUNT, V. and GREENBERG, B. (1990) Neural networks in radiologic diagnosis: II. Interpretation of neonatal chest radiographs. *Invest. Radiol.*, **25,** 1017–1023.

GROSSBERG, S. and MINGOLLA, E. (1987) Neural dynamics of surface perception: boundary webs, illuminants, and shape-from-shading. *Computer Vision Graph. & Image Proc.*, **37,** 116–165.

GUYOT, F., ALEXANDRE, F., HATON, J. P. and BURNOD, Y. (1989) A potentially powerful connectionist unit: the cortical column. NATO Advanced Workshop on Neuro Computing, NATO, Bonas, France, 22nd–27th Aug. 1988, 197–206.

HECHT-NIELSEN, R. (1987) Counter-propagation networks. Proc. 1st Int. Conf. on Neural Networks, June, San Diego, California, USA, Part 2, 19–32.

HINTON, G. E. (1989) Connectionist learning procedures. *Artif. Intell.*, **40,** 185–234.

HOPFIELD, J. J. and TANK, D. W. (1986) Computing with neural circuits: a model. *Science*, **223,** 625–633.

HUANG, J. N. and KUNG, S. Y. (1989) Parallel algorithms/architectures for neural networks. *J. VLSI Sig. Proc.*, **1,** 221–251.

IWATA, A., NAGASAKA, Y. and SUZUMURA, N. (1989) A digital Holter monitoring system with dual 3 layer neural networks. Proc. Int. Joint. Conf. on Neural Networks, 18th–22nd June, Washington DC, USA, Part 2, 69–74.

JOLLIFFE, I. T. (1986) *Principal component analysis.* Springer-Verlag, Berlin.

KARKHANIS, P. A., CHEUNG, J. Y. and TEAGUE, S. M. (1990) Using a PC based neural network to estimate the ejection fracion of a human heart. *Microcomputer Appl.*, **9,** (3), 99–107.

KARNIN, E. D. (1990) A simple procedure for pruning back-propagation trained neural networks. *IEEE Trans.*, **NN-1,** 239–242.

KATZ, W. T. and MERICKEL, M. B. (1989) Translation invariant aorta segmentation from magnetic resonance images. Proc. Int. Joint Conf. on Neural Networks, 18th–22nd June, Washington DC, USA, Part 1, 327–333.

KELLY, M. F., PARKER, P. A. and SCOTT, R. N. (1990) The application of neural networks to myoelectic signal analysis: a preliminary study. *IEEE Trans.*, **BME-37,** 221–230.

KHOSHABA, T., BADIE, K. and HASHEMI, R. M. (1990) EMG pattern classification based on back propagation neural network for prosthesis control. Proc. 12th Ann. Int. Conf. IEEE Eng. in Med. & Biol. Soc., 1st–4th Nov., Philadelphia, Pennsylvania, USA, Section 3, 1474–1475.

KIPPENHAN, J. S. and NAGEL, J. H. (1990) Diagnosis and modelling of Alzheimer's disease through neural network analysis of PET studies. Proc. 12th Ann. Conf. IEEE Eng. in Med. & Biol. Soc., 1st–4th Nov., Philadelphia, Pennsylvania, USA, Section 3, 1449–1450.

KOHONEN, T. (1990) The self-organizing map. *IEEE Proc.*, **78,** 1464–1480.

LEE, S. C. (1989) Using translation-invariant neural network to diagnose heart arrhythmia. Proc. 11th Ann. Int. Conf. IEEE Eng. in Med. & Biol. Soc., 9th–12th Nov., Seattle, Washington, USA, Section 6, 2025–2026.

LEHAR, S. M., WORTH, A. J. and KENNEDY, D. N. (1990) Application of the boundary contour/feature system to magnetic resonance brain scan imagery. Proc. Int. Joint Conf. on Neural Networks, 17th–21st June, San Diego, California, USA, Part 1, 435–440.

LINNENBANK, A. C., SIPPENSGROENEWEGEN, A. and GRIMBERGEN, C. A. (1990) Artificial neural networks applied in multiple lead electrocardiography rapid quantitative classification of ventricular tachycardia QRS internal patterns. Proc. 12th Ann. Int. Conf. IEEE Eng. in Med. & Biol., 1st–4th Nov., Philadelphia, Pennsylvania, USA, Section 3, 1461–1462.

MILLER, A. S., BLOTT, B. H. and HAMES, T. K. (1992) Neural networks for electrical impedance tomography image characterisation. *Clin. Phys. Physiol. Meas.*, **13,** Suppl. A, 119–123.

MILLER, L. F., SMITH, G. T., WU, Y. and UHRIG, R. E. (1990) Evaluation of neural networks for parameter identification from positron emission tomography scans. *Trans. Am. Nuclear Soc.*, **62,** 5–6.

MINSKY, M. L. and PAPERT, S. A. (1988) *Perceptrons: an introduction to computational geometry*, 2nd edn. MIT Press, Cambridge, Massachusetts, USA.

NAVABI, M. J., WATT, R. C., MYLREA, K. C. and HAMEROFF, S. R. (1990) Classification of CO_2 waveforms using artificial neural networks. Proc. 12th Ann. Conf. IEEE Eng. in Med. & Biol. Soc., 1st–4th Nov., Philadelphia, Pennsylvania, USA, Section 3, 1455–1456.

NEKOVEI, R. and SUN, Y. (1990) An adaptive algorithm for coronary artery identification in cineangiograms. Proc. 12th Ann. Int. Conf. of the IEEE Eng. in Med. & Biol. Soc., 1st–4th Nov., Philadelphia, Pennsylvania, USA, Section 3, 1459–1460.

NIKOONAHAD, M. and LIU, D. C. (1990) Medical ultrasound imaging using neural networks. *Electron. Lett.*, **26,** 545–546.

ÖZKAN, M., SPRENKLES, H. G. and DAWANT, B. M. (1990) Multispectral magnetic resonance image segmentation using neural networks. Proc. Int. Joint Conf. on Neural Networks, 17th–21st June, San Diego, California, USA, Part 1, 429–434.

PAUL, J. and VON GOLDAMMER, E. (1990) Neural net applications in medicine. In *Symbols versus neurons*. STENDER, J. and ADDIS, T. (Eds.), IOS Press, Amsterdam, 215–231.

PIETKA, E. (1989) Neural nets for ECG classification. Proc. 11th Ann. Int. Conf. IEEE Eng. in Med. & Biol., Soc., 9th–12th Nov., Seattle, Washington, USA, Section 6, 2021–2022.

PINHO, A. J. and ALMEIDA, L. B. (1991) Biomedical image compression based on artificial neural networks. Proc. 1st European Conf. on Biomed. Eng., 17th–20th Feb., Nice, France, 90–91.

PRINCIPE, J. C., CHANG, T. G., GALA, S. K. and TOME, A. P. (1989*a*) Information processing models for automatic sleep scoring. Proc. 11th Ann. Int. Conf. IEEE Eng. in Med. & Biol. Soc., 9th–12th Nov., Seattle, Washington, USA, Section 6, 1804–1805.

PRINCIPE, J. C., GALA, S. and CHANG, T. (1989*b*) Sleep staging based on the theory of evidence. *IEEE Trans.*, **BME-36,** 503–509.

RAFF, U. and NEWMAN, F. D. (1990) Lesion detection in radiologic images using an autoassociative paradigm: preliminary results. *Med. Phys.*, (US), **17,** 926–928.

RUMELHART, D. E., HINTON, G. E. and MCCLELLAND, J. L. (1986) A general framework for parallel distributed processing. In *Parallel distributed processing: explorations in the microstructure of cognition Volume 1 Foundations*. RUMELHART, D. E., MCCLELLAND, J. L. and The PDP Research Group (Eds.), MIT Press, Cambridge, Massachusetts, USA, 45–76.

SANDER, D. G., BARRETT, T. K. and PALMER, D. A. (1991) Use of a neural network to control an adaptive optics system for an astronomical telescope. *Nature*, **351,** 300–302.

SCHELLENBERG, J. D., NAYLOR, W. C. and CLARKE, L. P. (1990) Application of artificial neural networks for tissue classification from multispectral magnetic resonance images of the head. Proc. 3rd Ann. IEEE Symp. on Computer-Based Medical Systems, June, Chapel Hill, North Carolina, USA, 350–357.

SCHERTZ, L. D., VANNIER, M. W., GADO, M. H. and BUTTERFIELD, R. L. (1985) Statistical assessment of MR image classification techniques. In Proc. 7th Ann. Conf. IEEE Eng. in Med. & Biol. Soc., Sept., Chicago, Illinois, USA, 585–592.

SCHIZAS, C. N., PATTICHIS, C. S., SCHOFIELD, I. S., FAWCETT, P. R. and MIDDLETON, L. T. (1989) Artificial neural net algorithms in classifying electromyographic signals. Proc. 1st IEE Conf. on Artificial Neural Networks, IEE Conf. Publ. 1989/313, 16th–18th Oct., London, UK, 134–138.

SEBALD, A. V. (1989) Use of neural networks for detection of artifacts in arterial pressure waveforms. Proc. 11th Ann. Int. Conf. IEEE Eng. in Med. & Biol. Soc., 9th–12th Nov., Seattle, Washington, USA, Section 6, 2034–2035.

SILVERMAN, R. H. and NOETZEL, A. S. (1990) Image processing and pattern recognition in ultrasonograms by backpropagation. *Neural Networks*, **3,** 593–603.

STRAND, E. M. and JONES, W. T. (1990) A neural network for tracking the prevailing heart rate of the electrocardiogram. Proc. 3rd Ann. Symp. on Computer-Based Medical Systems, June, Chapel Hill, North Carolina USA, 358–365.

STUBBS, D. F. (1990) Multiple neural network approaches to clinical expert systems. *SPIE Appl. of Artif. Neural Networks*, **1294,** 433–441.

TOULSON, D. L. and BOYCE, J. F. (1991) Segmentation of MR images using neural nets. IEE Colloq. on Image Proc. in Med., IEE Coll. Dig. 1991/84, 19th April, London, UK, Paper 5.

TSAI, Y. S., HUNG, B. N and TUNG, S. F. (1990) An experiment on ECG classification using back-propagation neural network. Proc. 12th Ann. Int. Conf. IEEE Eng. in Med. of Biol. Soc., 1st–4th Nov., Philadelphia, Pennsylvania, USA, Section 3, 1463–1464.

WASSERMAN, P. D. (1989) *Neural computing theory and practice*. Van Nostrand Reinhold, New York.

WIDROW, B. and WINTER, R. (1988) Neural nests for adaptive filtering and adaptive pattern recognition. *IEEE Computer*, March, 25–39.

WIDROW, S. and LEHR, M. A. (1990) 30 years of adaptive neural networks: perceptron, madaline, and backpropagation. *Proc. IEEE*, **78,** 1415–1441.

XUE, Q., HU, Y. H. and TOMPKINS, W. J. (1989) A neural network weight pattern study for ECG pattern recognition. Proc. 11th Ann. Int. Conf. IEEE Eng. in Med. & Biol. Soc., 9th–12th Nov., Seattle, Washington, USA, Section 6, 2023–2024.

XUE, Q., HU, Y. H. and TOMPKINS, W. J. (1990) Training of ECG signals in neural network pattern recognition. Proc. 12th Ann. Int. Conf. IEEE Eng. Med. & Biol. Soc., 1st–4th Nov., Philadelphia, Pennsylvania, USA, Section 3, 1465–1466.

YEAP, T. H., JOHNSON, F. and RACHNIOWSKI, M. (1990) ECG beat classification by a neural network. Proc. 12th Ann. Int. Conf. IEEE Eng. in Med. & Biol. Soc., 1st–4th Nov., Philadelphia, Pennsylvania, USA, Section 3, 1457–1458.

Authors' biographies

Ade Miller gained his B.Sc. in Physics from the University of Southampton, UK, in 1990 and is currently researching in astrophysics for a Ph.D. His research interests include neural network applications to both medical and astronomical imaging systems; specific projects involve electrical impedance tomography and gamma ray telescope imaging.

Dr B. H. Blott is on the Physics staff of the University of Southampton, UK. His research interests are in the biophysics of cell membranes, particularly the effects of electric and magnetic fields thereon, and in medical imaging.

Dr T. K. Hames is currently Director of Technical Services at Southend General Hospital, UK. His research interests are in the application of signal and image processing to the development of noninvasive assessment techniques including ultrasound and electrical impedance plethysmography.