Abstract

Artificial Neural Networks have received renewed interest over the last few years. Applications have been described in a huge variety of engineering problems: machine vision, speech recognition, and control problems being typical areas of interest.

This short paper reviews the use of ANNs in medicine, and includes references to expert systems, medical imaging, signal processing and classification of diseases.

The Use of Artificial Neural Networks in Medicine

Denis Anthony
Department of Engineering
Warwick University

September 11, 1991

KEYWORDS

Neural Networks, Expert Systems, Back Propagation, Medicine

1 Introduction

For some classes of problems which require a computed solution, a conventional programming approach has not given satisfactory performance. Machine vision, speech recognition and odour detection are examples of tasks where artificial sensors in conjunction with conventional software have failed to approach the capability of biological systems. The sensory organs of animals are far better able to detect subtle changes, and analyse the inputs than current man-made systems. In part this may be considered a hardware problem, but the algorithms applied to the hardware are also critical.

Take vision as an example. The visual system uses concentric rings of retinal cells. These are massively connected to intermediate levels of neurons before termination in the visual cortex. It has been pointed out, for example in [1], that the number of connections in the human visual system is enormous, given roughly 10^8 neurons in the retina, and each cell is connected to between hundreds to tens of thousands of others. Such a system is able, even with the slow neuronal transfer speeds, which are in the order of milliseconds, to allow real-time response to complicated scenes at around 20-30 frames per second. This is far more impressive than the performance of very powerful computers using sequential algorithms. But the number of steps performed by any one neuron can only be in the order of tens to hundreds. It is inconceivable that the complex visual performance of humans can be handled by so few serial steps, it must involve massive parallelism, whereby many different but small tasks are computed simultaneously.

It may be that the parallelism offered by the central nervous system is more appropriate than serial programming in very fast processors for biological tasks. A better paradigm for image analysis and other "hard" problems may be to mimic the mammalian visual system by using massively parallel computation, using simple individual processes. A large variety of such artificial neural networks exists, which may be employed dependent upon the application.

Medicine, concerned with biological function, may be expected to be an arena in which ANNs show benefit. There are several large areas of application of ANNs to medicine. Because some ANNs can solve nonlinear problems (e.g. back propagation

using at least one hidden layer) control problems with nonlinearities may be dealt with. Prostheses could benefit from ANN control, e.g. manipulation using an artificial hand. Synthetic sensory organs may be built, or impaired organs assisted by networks.

Classification problems can be solved using ANNs. These include pattern recognition problems, such as found in imaging techniques, and the diagnosis of diseases.

The following sections will explore a cross section of applications of ANNs in the medical area.

2 Prostheses

As [2] has noted, prostheses have been used for centuries, however the best of them are not as good as the original healthy part. Artificial limbs for example lack the range and variety of movement. The use of "smart" structures with embedded actuators and sensors will become more common, and it is likely that some of these will benefit from the use of neural circuitry. Such a "neural prosthesis" could react to signals in a closed loop [2].

Sensory prostheses have particular difficulties to overcome. The ear and eye (e.g.) manage to cover enormous dynamic ranges. They are active and adaptive, the eye reacts to differing light intensities by altering the pupil size (e.g.) and exhibit nonlinear behaviour. Work has been done on the design of adaptive sensors, e.g. [3] who built an electronic cochlea. ANNs have been used to help analyse data from an electronic nose [4], where the ANN could distinguish with some success the aromas of different brands of beers and alcohols. Work is currently being undertaken to determine whether the ANN can classify coffee aromas.

The somatosensory system has been simulated using neural nets. It is known [5] that while peripheral fibre response are isomorphic to the stimulus form, cortical neurons are not, implying some transformation has occurred. A Walsh-Haramard transform, applied in multiple stages in a linear neural net gave outputs which were compatible with anatomical and physiological data [5].

Colour vision has been simulated, essentially mapping the 3 colour-sensitive cones of the retina to the 7 colours of the rainbow [6]. The net was trained on monochromatic light, and learned to associate the light with the relevant colour. The network gave results which were physiologically plausible; colour confusion lines of the macaque were similar to those generated by the hidden units of the network.

Electromyographic (EMG) patterns may be used to control prosthetic devices. Motor prostheses using EMG signals currently require the patient to learn how to actuate the control device using movements that are not necessarily natural. It would be better for the device to learn how the human intends to control it, by being trained on the EMG signals the patient will use to request a movement. Preliminary work has been reported in this area [7], where FFT analysed EMG signals were used in a back propagation ANN, and target outputs were given which were control movements. An alternative approach to this control problem is to use an auto-associative network (Hopfield net) to extract the relevant features from the data, which are then used as input to a classification (back propagation) network [8]. The effect is to map the features onto four control signals which move a prosthetic device.

3 Expert Systems

An expert system may be built by imputting numerical values representing knowledge items into the input layer of an ANN. Whereas a conventional expert system might ask the user to input the answer to a question, which could be a logical answer (yes/no), a level of confidence (a number in range [0.0,1.0] e.g.) or some text, the ANN can only accept a numerical input. Thus the input may be coded as 0.0 for no, 1.0 for yes, or a value given which represents a confidence level, or a value given for some essentially numerical data item (e.g. haemoglobin level) which would be scaled to lie in an appropriate range.

Expert systems have been built in this manner, indeed one of the most commonly documented medical applications of ANNs has been the manufacture of expert systems. Chest pain (for example) has been classified into high, medium and low risk cardiac patients [9]. In a comparison with a symbolic expert system which was designed to diagnose headaches, an MLP was built which took as inputs numeric coded answers to 216 questions into 216 input units [10]. These were fed via a hidden layer of 72 units to an output layer which represented 23 diseases as 23 output units. The outputs were either turned on (threshold 0.75), off (threshold 0.25) or in a "don't know" state. More than one disease was allowed to be present, and in fact from a training group of 300 patients 6% had more than one disease known by the system.

The symbolic expert system performed slightly better than the MLP, with an accuracy of 70% compared with 67% for the MLP for a test set of 94 patients.

However the symbolic expert system, though slightly more accurate, required approximately 5,000 hand crafted (and hence labour intensive) relations between symptom and disease. These factors were automatically produced by the ANN.

Where a patient makes an incorrect input to a system by (e.g.) misunderstanding a question, errors may creep into the final diagnosis. By determining which inputs are critical the MLP may ask for confirmation of certain questions.

In addition to answers to medical history questions, laboratory test results and electrocardiogram (ECG, also called EKG in the USA), findings were input to a MLP to diagnose acute myocardial infarction (heart attack) [11].

From an initial set of 41 inputs, one for each sign or symptom or laboratory result, various input sets were chosen, and by trial and error a 20 unit input set was found to be satisfactory. A detection rate of 92% and false alarm rate of 4% was found on testing, where 356 patients presenting were put into either the training or test group. The best previous reported performance had been 88% and 26% respectively for detection and false alarm rates [11].

A particularly challenging diagnostic test is to differentiate patients who have sustained infarction, but have no clear cut ECG evidence. Using only these cases as inputs, a detection rate of 80% was discovered. The authors concluded the MLP was not totally reliant upon ECG findings to give a diagnosis. A similar expert system using MLPs have been constructed to diagnose carcinoma of the lung [12].

The design of nets is often arbitrary, or a result of trial and error. Designing the architecture of the network with consideration of the particular problem to be solved is likely to result in more compact nets and faster learning. Typically all units in a layer are connected to all units in adjacent layers. If certain relations among the data are known these may be used to reduce the connectivity.

Reduction in connections will of itself speed the network convergence, and is less

likely to result in local minima. Furthermore if the network structure is predicated upon known relations, these relations will be learned more quickly. ANNs may also be used as part of an expert system, an approach used in [13] where a net classified ultrasonic defects as part of a larger system.

Hypertension has been diagnosed and treated using a multi-net approach [14]. Three separate nets were used:-

- A reference generating module (RGM) to learn normal variation of blood pressure related to age and sex during a 24 hour period.
- A drug compatibility module (DCM) to learn which drugs are best suited to a particular clinical picture.
- A therapy selecting module (TSP) to prescribe the treatment.

The nets were trained separately, and then linked together. The inputs to the RGM were age and sex, 24 outputs being the hourly blood pressures. The inputs did not include clinical data, so the net learned about typical systolic and diastolic time variations due to age and sex, unrelated to disease. The DCM took clinical data and gave outputs of appropriate drugs.

The TSM, which was the largest and most complicated network, was input blood pressure reading to the input layer. The first hidden layer was connected to a second hidden layer which was connected also to input units containing clinical data, and a further hidden layer combined both previous layers as inputs and drug compatibility data. This net took considerable time to converge, and thus the RGM and DCM trained nets were subsequently used to input into this module for training, as the relevant raw data was input to these nets. Thus the RGM and DCM acted as data compression nets for extracting relevant details for the TSM to train on.

The resultant combined net gave a sensitivity of 92% and a specificity of 96%. The appropriateness of treatment was high also, 82% of patients being deemed by a expert to have been given reasonable treatment.

Low back disorders [15] have been classified using ANNs, and have been compared with other techniques, which included a fuzzy logic system and clinical judgment. ANNs outperformed humans in most cases, and interestingly when they failed, clinicians tended to misdiagnose the same cases. Thus the authors concluded the ANNs were learning to make the same mistakes, as well as to diagnose correctly, as the clinicians. Four diagnostic classes were used in this study, and nets trained to perform the total classification performed less well than four separate nets trained to diagnose each type solely. Thus splitting the task up into subnets seems to be better than trying to solve a complete problem with one net. [16]

ANNs have also been used to diagnose dyspepsia [17], and headache [18], where in both cases standard back propagation was used.

4 Image Processing

By the use of thermal imaging tumour staging may be accomplished. Seven stages of tumour were identified in a study using white mice [19]. Taking the pixels from a thermal image as input to an ANN, it was possible to classify the tumours. However the convergence was slow, as the number of units was substantial. Reducing the input space

was accomplished by using the relevant frequencies after fourier analysis. The resultant net was smaller, ran faster, and gave acceptable classifications.

5 ECG Pattern Recognition

Applications for ECG include :-

- The detection of abnormal ECGs. A bi-directional associative memory model was used to detect abnormal ECG patterns [20]. This is an easier problem than classifying ECGs. Abnormal ECGs were detected with less than 1% error. The utility of this method is that patterns detected as abnormal may be stored for later analysis. Also normal and abnormal patterns may be time averaged.

- Back propagation has been used to classify ECG signals based on five parameters extracted from the ECG reading [21].

- Classification of arrythmias has been achieved using an auto-associative model [22]. The net distinguished between normal QRS and arrythmias with high accuracy. Classification of ECGs using 5 features extracted from the readings and used as input to an MLP has been used to place the ECGs into 6 diagnostic classes [23].

- Rather than parameterise ECG readings, [24] used digitised patterns. There is a problem inherent in this approach, it will not be known a priori where the QRS starts. One could pre-process the data to determine where the QRS is located, but in [24] the network was made to be translation independent. A trained multilayer network was able to distinguish with high accuracy between normal sinus rhythm, ventricular tachycardia and ventricular fibrillation.

- Data compression using an MLP has been used to reduce the data of a digital holter monitor. The network construction was similar to earlier work [25] where an MLP was used to learn the identity mapping, using inputs as targets, and where the single hidden layer gave a reduced dimensional representation of the data. In [26] an addition made to the method of [25] was to monitor the error of the trained network in real time, and to be training a second network on current data. If the error of reconstructing the data frequently went above a given threshold, the second network was copied to the original network, and thereafter used as the compressor. All networks are saved along with the outputs of the hidden layer, which may be later used to reconstruct the data.

6 EEG Pattern Recognition

Electroencephalographic (EEG) data has been analysed using back propagation ANNs to identify spikes [27]. In preliminary experiments the ANN was able to identify the spikes with a high precision on test data. EEG patterns are difficult to classify due to their noise and complexity [28]. Typical EEG studies use external stimuli to provoke an EEG reaction. In [28] readiness potential, the EEG pattern produced just prior to voluntary muscle movement, was used to train an ANN. The muscle movements were those produced when the subject pronounced syllables. The ANN was able to distinguish correctly 16 of 30 new patterns as belonging to one of 5 syllable groups.

Evoked EEG potentials were used to classify neurological diseases [29]. It was suggested that the ANN needed to be supplemented with additional knowledge to diagnose accurately, and that this could be implemented by pre-processing the data, or by building a co-operative system. The former approach would recognise "well-behaved" patterns and process them in an ANN. The latter approach would take the results from an ANN and combine it with other data to effect a diagnosis.

7 Ultrasound

Ultrasound is used as an imaging modality. There have been attempts to produce tomographic images, in a similar fashion as in X-Ray, NMR and nuclear studies. However there are problems associated with the relatively large wavelength (diffraction, refraction effects) and due to anisotropy. Tomogram reconstruction techniques often use simplifying assumptions, to avoid considering these effects.

ANNs are being utilised to solve the inverse problem, ie. given the times of flight (or amplitude, or some other attribute) of the ultrasonic beam through a specimen, deduce the structure that would produce this data [30]. While this is geared to non destructive testing of materials, the relevance to medical imaging is obvious. Previous work has been undertaken in which an MLP was trained with the transducer pressures obtained with a range of cylindrical objects as test specimens. The ANN was able to learn the size of the scanned object from the projection data [31].

Holographic images have been created using ultrasonic rays, where a 3 layer MLP learned to output images of letters of the alphabet [32].

The wave velocity within a tissue is usually considered to be constant, but may vary by up to 8% [33]. An MLP with 2 hidden layers was used to correct the delay noise that occurred as a result of the assumption of constant velocity. A substantial improvement was obtained in simulations.

8 Impedance Imaging

By applying a small current to two electrodes, and using the output voltages from an array of electrodes surrounding the organ/limb etc., it is possible to reconstruct an image. The reconstruction techniques have difficulties however, and may simplifying assumptions are used.

MLPs have been used [34] to solve the inverse problem, ie. given the voltages, what conductivity distribution would have caused such outputs.

9 Nuclear Medicine

MLPs have been used to differentiate between normal and Thallium-201 scintigrams which show coronary artery disease. For each of 5 segments of 3 views, 3 parameters were computed, and these 45 values fed into an MLP [35, 36]. Test data gave an accuracy of 82% for diagnosing cardiac disease.

Pulmonary scintigrams have been classified wuth very limited success using back propagation using principal components of images [37, 38, 39].

10 Nuclear Magnetic Resonance (NMR)

Back propagation has been employed [40] to segment NMR into 7 distinct tissue types (background, fat and skin, stroke, ventricle (CSF), grey matter, and white matter). In this implementation three images were used based on spin-lattice relaxation time (T1) spin-spin relaxation time (T2) and proton-density (PD). In practice a less computationally intensive set of parameters was used (T1-weighted, T2-weighted and PD-weighted). An ANN took the pixels as input, one at a time, and learnt to which class the pixels belonged, ie. to which tissue type. A separate network was constructed with three inputs, one for each parameter.

Using confusion matrices as an objective test the ANNs were compared with the Maximum Likelihood Classification (MLC) technique. MLC gave a better accuracy, with an overall classification of 94% compared with 91% for the ANN. However subjectively the images produced by the ANN were considered to be superior. The images given by the 3 input ANN were better than the single input ANN.

Neural models of vision systems have been developed by [41] which have been used to process NMR images [42]. It was concluded that the ANN extracted invariant features needed for pattern recognition.

11 X-Ray

Very little work has been undertaken in the area of radiological applications of ANNs. A simulation study has shown some potential efficacy in this area however [43]. A synthetic nodule was created by placing a 3 by 3 pixel square centrally in a 5 by 5 pixel grid. Pseudo-gaussian noise was added, and a network trained with a set of images, half of which contained a nodule. Receiver operating curves were plotted for test (unseen) images with varieties of signal to noise ratios (SNR). The training set had an SNR of 15.0, and test sets of SNR 1.5 and 3.0 were used. In comparison to human observers, who had been trained to distinguish the nodules, the ANN outperformed the human observers for both SNRs.

It was concluded that ANN could detect signals with improved performance compared with human observers. However it is noted that the signals were very simple, and not representative of clinical radiographs [44].

12 Cytology

ANNs have been compared with conventional techniques to automate identification of pre-cancerous cervical cell [45]. A frequency domain approach was used to compute a discriminant function which was used to classify the cells as normal or abnormal. This was compared with a conventional morphology approach which used various parameters including ratio area-of-nucleus to area-of-cytoplasm. Both methods were compared with the back-propagation network using 80 features derived from the frequency domain.

Of the three methods the morphology approach was the best with 99% accuracy and 3.8% maximum false negative, while using identical data the frequency domain approach gave 98% accuracy, and the ANN 95%. Thus ANNs are not necessarily the best method to use, and conventional techniques should be considered, and where appropriate used as a comparison to ANNs. However if accuracy may be compromised the ANN may be

less labour intensive than a conventional approach.

13 Other Medical Uses of ANNs

A variety of less common applications of ANNs exist, a few of which are discussed below.

13.1 Drug Therapy

An MLP has been used to estimate the pharmacokinetic properties of drugs. A comparison was made between an MLP, maximum Likelihood (ML) estimation and maximum a posteriori probability (MAP). The MLP performed well in estimating distribution volume and elimination clearance of theophylline [46, 47].

13.2 Mammograms

Work is currently under way [48] to remove irrelevant detail from mammograms using ANNs. The mammogram image has a background and textural information etc. which is easily ignored by a human operator, to allow attention to be focussed on the area of the image containing the clinical detail. An ANN is being explored to determine its effectiveness in emulating the human factor.

A similar technique was described in [49] and [37] where segments extracted from a nuclear image, some of which contained lung fields, and some which were artefacts, were classified by a Multilayer Perceptron (MLP). The MLP was able to correctly classify "unseen" (ie. not presented to the network in training) segments so that a subsequent net could be used to accept only segments of interest.

13.3 Arterial Pressure Waveforms

Arterial pressure waveforms have been analysed by two different ANN paradigms, back propagation and competitive learning [50]. In the latter there is unsupervised learning, ie. there is no "teacher" involved. The back propagation net successfully classified all test patterns as artefact or artefact-free. The competitive model had a high success rate of 93%, but consistently missed artefacts, classifying them as artefact-free.

14 Problems and Difficulties with the ANN Methods

One of the limitations of ANNs is the difficulty in determining what is happening within the network. The ANN is typically treated as a "black box", with known inputs and hopefully outputs which are appropriate to the problem under study. Using a "black box", where this is defined to be a system whose inputs and outputs are examinable, but whose internal workings are obscure, is appropriate where the relationships between inputs and outputs are not well known, provided one is aware of certain problems, some of which are discussed below.

A standard technique of evaluating nets is to present test data which the network has not been trained on, and report the accuracy of the net on this data. However as [51] state, this may not be sufficient. For the test set to be of any use certain criteria should be met :-

- Repeatability. The data collected should have low inter and between observer error.
- Representativity. Adequate amount of data for each class should be present.
- Accuracy. The data should be as complete as possible, and each item as accurate as possible.

In a medical system, where the outputs may be used to make decisions which could be critical to a patient, the effects of errors are pertinent. In an expert system (e.g.) one may analyse the effects of incorrect data entry by flipping the state of the input from a correct to an incorrect value, and determining the effect of this change on the resultant output, which may be a diagnosis (say) [52].

If the classification of patterns is taken as the sole criterion of success of an ANN, one may have undesirable performance for specific problems. For example it is more important to avoid misdiagnosing serious head injury as psychosomatic than vice versa, and therefore a network may need to show a worse absolute classification rate to avoid disasterous mis-diagnoses. The prognosis of head injury has been explored using an ANN system [53]. A cost matrix was used to penalise misclassifications that are most important to get right.

While the internal functioning of the network is difficult to analyse, some work has been done to gain some idea of what is going on. Rules may be generated from an expert system network. This may be achieved [18, 10]. by setting all inputs off, and monitoring the effect of the output states of diseases on the trained network. Where an additional input is turned on that switches the output off again, a negative rule is obtained. Combinatorial explosion may be avoided by limiting the number of affirmative and negative components to a rule. A list of rules drawn up in this fashion was judged by medically trained personnel to be reasonable.

There were some discrepancies between the rules generated by doctors and those generated by the MLP. In these cases it may be worth re-considering the medical judgments, rather than assuming the MLP is wrong. Thus the MLP may be seen as a useful indicator of possible clinical rule adjustments.

In most of the studies described above, ANNs were trained in "supervised" mode. This means that training data sets with target outputs were given to the ANN. It is difficult and expensive to obtain large enough data sets in many cases. Furthermore the targets are often assigned by subjective assessment, e.g. a physician's diagnosis, and may not be correct. Unsupervised learning may be achieved via ANNs. Kohonen nets [54] take input data and arrange it into classes, with no supervision. More conventional unsupervised classification techniques include the K-Means method. Ultrasonic liver image data has been classified into normal versus chronic active hepatitis (e.g.) using non ANN clustering techniques [55].

The use of unsupervised ANN learning may be of use where large training sets are difficult to obtain, or whose accuracy is doubtful. A particular application that has been suggested is in EMG analysis. Given an EMG pattern, which is composed of several motor unit action potentials (MUAP), each of which is composed of many single fibre action potentials, a problem to be solved is how many MUAPs there are, and to describe the SUAPs, and the firing pattern. This problem has been simulated using an ANN with Gaussian nodes [56], where a partition into clusters is attempted such that the clusters form SFAPs from the same MUAP.

If an ANN is trained beyond the point of convergence, the total error may still be falling slightly, but the classification of unseen data may become worse [57]. The training

data is subject to overfitting after convergence, and generally this should be avoided. One technique to spot this is to use a test set to report an error, but not use it in training. Ie. a training set is used to update the weights that is separate from the test set. Overfitting will be evident from an increased test set error.

Networks are notoriously difficult to design. Some work has been done using genetic algorithms to find optimal structures, e.g. [58]. Networks may also be pruned of nodes where the least utility is located, as in [59, 60]. Generally networks are designed by trial and error however. ANNs are sensitive to initial conditions, and the parameters used (momentum and learning rate e.g.). The parameters used are often reached by trial and error, though some work has been done to automatically set the parameters, and to dynamically tune them to the current state of the network [61, 62, 63].

It is generally stated that ANNs are resistant to noise and erroneous inputs, though rarely is this aspect covered in actual implementations. As [64] has shown, ANNs trained on noiseless data are not as resistant to noise as those trained on data with noise. Thus in simulations, exact derived data may not be adequate for training a system that will use real data. If the actual network is likely to be subject to corruption, as opposed to the data input to it, the network may need to be overtrained. In simulations nets which had been trained past convergence were more hardware resistant than those whose training stopped at convergence [64, 65]. The information becomes more evenly spread through the network after convergence is reached. Thus one may wish to effect a compromise between fault tolerance and overtraining if faults in hardware are likely.

15 Conclusion

Neural networks have been found to be effective in a wide variety of medical applications. Their use should however be the subject of some caution, as they may be used inappropriately. The choice of network paradigm, the design of the network architecture, the choosing of parameters are areas which need to be considered when implementing a network solution.

References

[1] Chan L. *Adaptive and Invariant Connectionist Models For Pattern Recognition.* PhD thesis, Cambridge University Engineering Department, 1989.

[2] Ham F. Neurocontrol: Past, present and future. In *Southcon 90 Conference Record*, pages 206–211, Orlando Florida, 1990.

[3] Lyon R.F and Mead C. An analog electronic cochlea. *IEEE Transactions*, ASSP-36 7:1119–1134, 1988.

[4] Gardner J.W, Hines E.L, and Wilkinson M. Application of artificial neural networks to an electronic olfactory system. *Measurement in Science Technology*, 1:446–451, 1990.

[5] Bankman I.N, Johnson K.O, and Hsiao S.S. A neural network model of transformations in the somatosensory system. In *Images of the 21st Century: IEEE Eng. in Medicine and Biology 11th Annual Conference* [66], pages 2060–2061.

[6] Usui S, Nakuchi S, and Miyake S. Neural network model of colour vision. In *Images of the 21st Century: IEEE Eng. in Medicine and Biology 11th Annual Conference* [66], pages 2044–2045.

[7] Hiraiwa A, Shimohara K, and Tokunaga Y. EMG pattern analysis and classification by neural network. In *IEEE International Conference on Systems, Man, and Cybernetics* [67], pages 1113–1115.

[8] Kelly M.F, Parker P.A, and Scott R.N. The application of neural networks to myoelectric signal analysis - a preliminary-study. *IEEE Transactions On Biomedical Engineering*, 37(3):221–230, 1990.

[9] Hart A. and Yyatt J. Connectionist models in medicine: An investigation of their potential. In Hunter J. Cookson J. and Wyatt J., editors, *Lecture Notes in Medical Informatics*, volume 38, pages 115–124. Springer-Verlag, 1989.

[10] Saito K and Nakano R. Medical diagnostic expert system based on PDP model. In *Proc 3rd Int. Joint Conference on Neural Nets*, volume I, pages 255–262, San Diego, California, 1989. IEEE.

[11] Baxt W.G. Use of artificial neural network for data analysis in clinical decision-making: The diagnosis of acute coronary occlusion. *Neural Computation*, 2(2):480–489, 1990.

[12] Cohen M.E, Hudson D.L, and Anderson M.F. Combination of a neural model and a rule-based expert system to determine efficacy of medical testing procedures. In *Images of the 21st Century: IEEE Eng. in Medicine and Biology 11th Annual Conference* [66], pages 1991–1992.

[13] Hallam N.J, Hopgood A.A, and Woodcock N. Defect classification in welds using a feedforward network within a blackboard system. In *Inns Int. Neural Networks Conference* [68], pages 353–356.

[14] Poli R, Cagnoni S, Livi R, Coppini G, and Valli G. A neural network expert system for dianosing and treating hypertension. *Computer*, March:64–71, 1991.

[15] Bounds D.G and Lloyd P.J. A multi layer perceptron network for the diagnosis of low back pain. In *Proc. 2nd International Joint Conference on Neural Networks* [69], pages 481–489.

[16] Bounds D.G and Lloyd P.J. A comparison of neural network and other pattern recognition approaches to the diagnosis of low back disorders. Technical report, Royal Signals and Radar Establishment, Malvern, Great Britain, 1990.

[17] Scalia F, Marconi L, Ridella S, Arrigo P, Mansi C, and Mela G.S. An example of back propagation: Diagnosis of dyspepsia. In *IEE Int. Conference on Neural Networks*, pages 332–336, London, 1989.

[18] Saito K and Nakano R. Rule extraction from facts and neural networks. In *Inns Int. Neural Networks Conference* [68], pages 379–382.

[19] Egbert D.D and Rhodes E.E. Preprocessing of biomedical images for neurocomputer analysis. In *Proc. 2nd International Joint Conference on Neural Networks* [69], pages I–561 –568.

[20] Cheung J.Y and Hull,Jr., S.S. detection of abnormal electrocardiograms using a neural network approach. In *Images of the 21st Century: IEEE Eng. in Medicine and Biology 11th Annual Conference* [66], pages 2015–2016.

[21] Pietka E. Neural nets for ECG classification. In *Images of the 21st Century: IEEE Eng. in Medicine and Biology 11th Annual Conference* [66], pages 2021–2022.

[22] Lin K and Chang W.H. Classification of QRS pattern by an associative memory model. In *Images of the 21st Century: IEEE Eng. in Medicine and Biology 11th Annual Conference* [66], pages 2017–2018.

[23] Xue Q, Hu Y, and Tompkins W.J. A neural network weight pattern study with ECG pattern recognition. In *Images of the 21st Century: IEEE Eng. in Medicine and Biology 11th Annual Conference* [66], pages 2023–2024.

[24] Lee S.C. Using a translation-invariant neural network to diagnose heart arrhythmia. In *Images of the 21st Century: IEEE Eng. in Medicine and Biology 11th Annual Conference* [66], pages 2025–2026.

[25] Cottrell G.W, Munro P, and Zipser D. Learning internal representations from gray-scale images: An example of extensional programming. *Cognitive Science Soc. Annual Conference*, 9th Procs.:461–473, 1987.

[26] Nagasaka Y, Iwata A, and Suzumura N. Data compression using neural network for digital holter monitor. In *Images of the 21st Century: IEEE Eng. in Medicine and Biology 11th Annual Conference* [66], pages 2019–2020.

[27] Eberhart R.C, Dobbins R.W, and Webber W.R.S. EEG waveform analysis using casenet. In *Images of the 21st Century: IEEE Eng. in Medicine and Biology 11th Annual Conference* [66], pages 2046–2057.

[28] Hiraiwa A, Shimohara K, and Tokunaga Y. EEG topography recognition by neural network. In *IEEE International Conference on Systems, Man, and Cybernetics* [67], pages 1116–1117.

[29] Bruha I and Madhaven G.P. Need for a knowledge-based subsystem in evoked potential neural-net recognition system. In *Images of the 21st Century: IEEE Eng. in Medicine and Biology 11th Annual Conference* [66], pages 2042–2043.

[30] Anthony D.M, Hines E.L, Hutchins D, and Mottram T. Ultrasound tomography imaging of composite defects using neural networks. Technical report, Department of Engineering Warwick University, Great Britain, 1991.

[31] Conrath B.C, Daft C.M.W, and O'Brien W.D. Applications of neural networks to ultrasound tomography. *Ultrasonics Symposium*, pages 1007–1010, 1989.

[32] Watanabe S and Yoneyama M. Ultrasonic robot eyes using neural networks. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 37(3 May):141–147, 1990.

[33] Nikoonahad M and Liu D.C. Medical ultrasound imaging using neural networks. *Electronic Letters*, 26(8 Apr 14):545–546, 1990.

[34] Neaves P. Design and development of a complex impedance measurement system for impedance tomograhpy. Master's thesis, Department of Engineering Warwick University, 1991.

[35] Sochor H, Dorffner G, and Porenta G. Classification of thallium-201 scintigrams using a neural network trained by back propagation. *Journal of Nuclear Medicine*, JULY:1314, 1988.

[36] Porenta G, Dorffner G, Schedlmayer J, and Sochor H. Parallel distributed processing as a decision support approach in the analysis of thallium-201 scintigrams. In *Proc. Computers in Cardiology*, pages 259–262, Washington, 1988.

[37] Anthony D.M. *The Use of Artificial Neural Networks in Classifying Lung Scintigrams*. PhD thesis, University of Warwick, 1991.

[38] Anthony D.M, Hines E.L, Taylor D, and Barham J. The use of neural networks in classifying lung scintigrams. In *Inns Int. Neural Networks Conference* [68], pages 71–74.

[39] Anthony D.M. Barham J. Hines E.L. and Taylor D. A comparison of image compression by neural networks and principal compoents analysis. In *Proc. 4th International Joint Conference on Neural Networks* [70], pages 339–344.

[40] Ozkan H, Hendrick M.S, Sprenkels G, Benoit M.S, and Dawant M. Multi-specral magnetic resonance image segmentation using neural networks. In *Proc. 4th International Joint Conference on Neural Networks* [70], pages 429–434.

[41] Grossberg S and Mingolla E. Neural dynamics of perceptual grouping: Textures, boundaries and emergent segmentation. *Perception and Psychophysics*, 38:144–210, 1985.

[42] Lehar S.M, Worth A.J, and Kennedy D.N. Application of the boundary contour/feature contour system to magnetic resonance brain scan imagery. In *Proc. 4th International Joint Conference on Neural Networks* [70], pages 435–440.

[43] Boone J.M, Sigillito V.G, and Shaber G.S. Neural networks in radiology: An introduction and evaluation in a signal detection task. *Medical Physics*, 17:2:234–241, 1990.

[44] Baker B., Curry S., and Baumrind S. Neural network method for solving pattern recognition problems in craniofacial X-ray image analysis. *Proc. 11th Annual Int Conf on Engineering in Medicine and Biology*, 11(5):1646, 1989.

[45] Ricketts I.W. Neural networks and conventional approaches to the detection of abnormal cells. Dept. of Mathematics and Computer Science, University of Dundee.

[46] Shadmehr R and D'Argenio D.Z. Connectionist modeling vs. bayesian procedures for sparse data pharmokinetic parameter estimation. In *Images of the 21st Century: IEEE Eng. in Medicine and Biology 11th Annual Conference* [66], pages 2058–2059.

[47] Shadmehr R. A neural network for nonlinear bayesian estimation in drug therapy. *Neural Computation*, 2:216–225, 1990.

[48] Research Bulletin University of Dundee Microcomputer Centre. Microcompter Centre, University of Dundee.

[49] Anthony D.M, Hines E.L, Taylor D, and Barham J. An investigation into the use of neural networks for an expert system in nuclear medicine image analysis. In *IEE conference on image processing*, pages 338–342, Warwick University, Coventry, 1989.

[50] Sebald A.V. Use of neural networks for detection of artifacts in arterial pressure waveforms. In *Images of the 21st Century: IEEE Eng. in Medicine and Biology 11th Annual Conference* [66], pages 2034–2035.

[51] Hart A. Evaluating black boxes as medical decision aids. *Medical Informatics*, 15(3):229–236, 1990.

[52] Hyman W.A and Mitta D.A. Statistical analysis of the effect of input errors on expert system advice. In *Images of the 21st Century: IEEE Eng. in Medicine and Biology 11th Annual Conference* [66], pages 1771–1772.

[53] Lowe D and Webb A. Exploiting prior knowledge in network optimization: an illustraton from medical prognosis. *Network*, 1:199–323, 1990.

[54] Kohonen T. Self organized formation of topologically correct feature maps. *Biological Cybernetics*, 43:59–69, 1982.

[55] Momenan R, Wagner R, Loew M, Insana M, and Garra B. Unsupervised learning for characterization of tissue from acoustical speckle in ultrasound images. *unknown*, 1987.

[56] Huang Q, Graupe D, Huang Y, and Liu R. Identification of firing patterns of neuronal signals. In *The 28th IEEE Conference on Decision and Control*. IEEE, 1989.

[57] Chauvin Y. Generalization performance of overtrained back-propagation networks. In Almeida L.B and Wellekens C.J, editors, *In Neural Networks*, pages 46–55. Springer-Verlag, 1987.

[58] Harp S.A, Samad T, and Guha A. Towards the genetic synthesis of neural networks. *Int. Conference on Genetic Algorithms*, pages 360–369, 1989.

[59] Sankar A and Mammone R.J. Optimal pruning of neural tree networks for improved generalization. In *Proc. 4th International Joint Conference on Neural Networks* [71].

[60] Sankar A and Mammone R.J. Speaker independent vowel recognition using neural tree networks. In *Proc. 4th International Joint Conference on Neural Networks* [71].

[61] Chan L. and Fallside F. An adaptive training algorithm for back propagation networks. *Computer Speech and Language*, 2:205–218, 1987.

[62] Vogl T.P, Mangis J.K, Rigler A.K, Zink W.T, and Alkon D.L. Accelerating the convergence of the back-propagation method. *Biological Cybernetics*, 59:257–263, 1988.

[63] Lister J.B, Schnurrenberger H, and Mamillod P. Implementation of a multi-layer perceptron for a non-linear control problem. Technical report, Ecole Polytechnique Federale de Lausanne, Dec 1990.

[64] Nijhuis J, Hofflinger B, van Schaik A, and Spaanenburg L. Limits to the fault-tolerance of a feedforward neural network with learning. *Proceedings IEE Fault-Tolerant Computing Symposium 20*, pages 228–235, 1990.

[65] Nijhuis J and Spaanenburg L. Fault tolerance of neural associative memories. *IEE Proceedings*, 136(E):389–394, 1989.

[66] IEEE. *Images of the 21st Century: IEEE Eng. in Medicine and Biology 11th Annual Conference*, 1989.

[67] *IEEE International Conference on Systems, Man, and Cybernetics*, volume 3, 1989.

[68] *Inns Int. Neural Networks Conference*, Paris, 1990.

[69] IEEE. *Proc. 2nd International Joint Conference on Neural Networks*, San Diego, California, July 1988.

[70] IEEE. *IJCNN International Joint Conference on Neural Networks*, volume 1, San Diego, 1990.

[71] IEEE. *IJCNN International Joint Conference on Neural Networks*, volume 1, Seattle, 1990.