

University of East London Institutional Repository: <http://roar.uel.ac.uk>

This paper is made available online in accordance with publisher policies. Please scroll down to view the document itself. Please refer to the repository record for this item and our policy information available from the repository home page for further information.

To see the final version of this paper please visit the publisher's website. Access to the published version may require a subscription.

Author(s): Dybowski, Richard.

Article Title: Neural computation in medicine: Perspectives and prospects

Year of publication: 2000

Citation: Dybowski R. (2000) "Neural computation in medicine: Perspectives and prospects". In Malmgren, H., Borga, M., Niklasson, L. (eds.) Artificial Neural Networks in Medicine and Biology (Proceedings of the ANNIMAB-1 Conference, Goteborg, 13-16 May 2000) Springer, 2000, pp 26-36.

Link to published version:

<http://www.springer.com/computer/communications/book/978-1-85233-289-1>

ISBN: 978-1852332891

Publisher statement: <http://www.springer.com/open+access?SGWID=0-169302-12-467999-0>

Information on how to cite items within roar@uel:

<http://www.uel.ac.uk/roar/openaccess.htm#Citing>

Neural Computation in Medicine: Perspectives and Prospects*

Richard Dybowski

Medical Informatics Laboratory, Division of Medicine,
King's College London (St Thomas' Campus), London, UK

Abstract

In 1998, over 400 papers on artificial neural networks (ANNs) were published in the context of medicine, but why is there this interest in ANNs? And how do ANNs compare with traditional statistical methods? We propose some answers to these questions, and go on to consider the 'black box' issue. Finally, we briefly look at two directions in which ANNs are likely to develop, namely the use of Bayesian statistics and knowledge-data fusion.

1 Cognitive and data modelling

In the wake of publication of the back-propagation algorithm in 1986, the number of ANN-oriented articles featured in the Medline database has grown from 2 in 1990 to 473 in 1998. But why is there this interest in ANNs within medicine?

The design of ANNs was originally motivated by the phenomena of learning and recognition, and the desire to model these cognitive processes. The *cognitive-modelling* branch of ANN research is still active, and it is relevant to medicine in providing models of psychological and cerebral dysfunction [1, 2, 3]. However, in the late 1980s, a more pragmatic stance emerged, and ANNs became to be seen also as tools for *data modelling*, primarily for classification.

Medicine involves decision making, and classification is an integral part of that process, but medical classification tasks, such as diagnosis, can be far from straightforward. At least two sources of difficulty can be identified.

- Clinical laboratories are being subjected to an ever-increasing *workload*. Much of the data received by these laboratories consists of complex figures, such as cytological specimens – objects traditionally interpreted by experts – but experts are a limited resource.
- The *complexity* of patient-related data can be such that even an expert can overlook important details. With the production of high volumes of possibly high-dimensional data from instruments such as flow-cytographs, and the integration of disparate data sources through *data fusion*, this is becoming an increasing problem.

*In Malmgren H, Borga M, Niklasson L. (eds.) *Proceedings of the ANNIMAB-1 Conference (Artificial Neural Networks in Medicine and Biology)*, Göteborg, 13-16 May 2000 (Springer, 2000) pp. 26-36.

Presented with such problems, it is natural that the medical fraternity should turn to ANNs in the hope that these adaptable models could alleviate at least some of the problems; however, the use of ANNs in medicine has raised a number of issues. One of these is the relationship between ANNs and statistics; another is the use of ‘black box’ systems in medicine. We look at both these issues and then consider two directions in which ANNs are likely to develop.

2 A statistical perspective

2.1 Multilayer perceptrons

Multilayer perceptrons (MLPs) with sigmoidal hidden-node functions are the ANNs most commonly used in medicine. Medical data are diverse and, at times, highly complex [4], but MLPs are universal approximators. The flexibility of MLPs has enabled them to be applied to a wide variety of medical fields [5, 6, 7, 8, 9, 10, 11], including oncology [12], dentistry [13], bacteriology [14], and sleep research [15]. For example, Tarrasenko [16] describes how an MLP was used to associate an electroencephalograph (EEG) with deep sleep, REM sleep or wakefulness. EEGs from nine patients were digitized and, in order to reduce the dimensionality of feature space, the EEGs were characterized by the ten coefficients of an autoregressive time-series model. These coefficients were used as the input vectors for an MLP with three outputs. The number of hidden units was varied from 4 to 20, and five initially random weight vectors were used. The best-performing MLP had a misclassification rate of 10.6%. In comparison, nearest-neighbour classification gave a misclassification rate of 18.3%.

The genesis and renaissance of ANNs took place within various communities, and papers published during these periods reflected the disciplines involved: biology and cognition; statistical physics; and computer science. But it was not until the early 1990s that a probability-theoretic perspective emerged, which regarded ANNs as being within the framework of statistics [17, 18, 19, 20].

A recurring theme of this literature is that many ANNs are analogous, or identical to, existing statistical techniques. For example, a popular statistical method for modelling the relationship between a binary response variable and a vector of covariates is *logistic regression*, but a single-layer perceptron with a logistic output function and trained by a cross-entropy error function will be functionally identical to a main-effects logistic regression model. Furthermore, an MLP can be regarded as both a non-linear extension of logistic regression and as a particular type of projection pursuit regression model. However, Ripley & Ripley [21] point out that the statistical algorithms for fitting projection pursuit regression are not as effective as those for fitting MLPs.

One may ask whether the apparent similarity between ANNs and existing statistical methods means that ANNs are redundant. One answer to this is given by Ripley [22]:

The traditional methods of statistics and pattern recognition are ei-

ther *parametric* based on a family of models with a small number of parameters, or *non-parametric* in which the models used are totally flexible. One of the impacts of neural network methods on pattern recognition has been to emphasize the need in large-scale practical problems for something in between, families of models with large but not unlimited flexibility given by a large number of parameters.

Another response is to point out that the widespread fascination for ANNs has attracted many talented researchers and potential users into the realm of data modelling. It is true that the neural-computing community re-discovered some statistical concepts already in existence, but this influx of participants has created new ideas and refined existing ones. These benefits include the *learning of sequences* by time delay and partial recurrence [23], and the creation of powerful visualization techniques, such as *generative topographic mapping* [24]. Thus the ANN movement has resulted in statisticians having available to them a collection of techniques to add to their repertoire. Furthermore, the placement of ANNs within a statistical framework has provided a firmer theoretical foundation for neural computation. This has led to new developments such as the Bayesian approach to ANNs [25], and to improvements to existing neural methods [26].

ANNs can be used jointly with conventional statistical methods. In the scheme suggested by Goodman & Harrell [27], the performance metrics of a generalized linear model are compared with those from an MLP. If the two approaches are not found to differ with statistical significance, the linear model is chosen since it is simpler with respect to computation and interpretation. The authors demonstrate this approach in the context of a coronary artery bypass dataset.

MLPs are trained using exemplars from each class of interest, but many pathologies are uncommon. In such situations a balanced training set should be used for the under-represented classes [28]. However, this may not be feasible when abnormalities are very rare, in which case *novelty detection* should be considered [29].

It should be emphasized that, even with correct training, an ANN will not necessarily be the best choice for a classification task in terms of accuracy. This has been highlighted by Wyatt [30], who wrote:

Neural net advocates claim accuracy as the major advantage. However, when a large European research project, StatLog, examined the accuracy of five ANN and 19 traditional statistical or decision-tree methods for classifying 22 sets of data, including three medical datasets [31], a neural technique was the most accurate in only one dataset, on DNA sequences. For 15 (68%) of the 22 sets, traditional statistical methods were the most accurate, and those 15 included all three medical datasets.

But one should add the comment made by Michie et al [31] on the results of the StatLog project:

With care, neural networks perform very well as measured by error rate. They seem to provide either the best or near best predictive performance in nearly all cases ...

Nevertheless, in order to justify the implementation of an ANN, its ease of implementation and performance should be compared with that obtained from one or more appropriate standard statistical techniques. Furthermore, one should also be aware of emerging non-neural alternatives, such as *support vector machines* for classification and *independent component analysis* for visualization, and also make comparisons with these, if appropriate.

2.2 Self-organizing feature maps

The most common neural system for unsupervised training is Kohonen's *self-organizing feature maps* (SOFMs) [32]. The aim of SOFMs is to map an input vector to one of a set of neurons arranged in a lattice, and to do so in such a way that positions in input space are topologically ordered with locations on the lattice. Hertz et al [33] liken this process to an elastic net, existing in input space, which wants to come as close as possible to the input vectors of a training set.

Amongst the interesting medical applications of SOFMs are their use for classification of craniofacial growth patterns [34], the extraction of information from electromyographic signals with regard to motor unit action potentials [35], and magnetic-resonance image segmentation [36]. With regard to the first example, an SOFM was used to extract the most relevant information from mandibular growth data. The position of a patient on the resulting map could be used to aid orthodontic diagnosis and treatment.

The example by Glass & Reddick [36] demonstrates how an SOFM can be used for image segmentation. An SOFM was trained with pixels from viable tumors and necrotic tissue, as visualized by magnetic resonance images. An MLP was then trained to distinguish between these two types of tissue on the basis of the SOFMs final input-node weights. Consequently, the SOFM-MLP combination characterized the type of tissue represented by a new pixel. This provided a non-invasive technique to assess an osteosarcoma patient's response to chemotherapy.

The so-called Growing Cell Structure [37] is related to SOFMs, and Walker et al [38] used this technique for the cytodiagnosis of breast carcinoma. In addition to performing classification, the method also enabled them to visualize how the input variables were involved in the classification.

Although the SOFM algorithm provides a means of visualizing the distribution of data points in input space, Bishop [39] points out that this can be weak if the data do not lie within a two-dimensional subspace of the higher-dimensional space containing the data. Another observation, made by Balakrishnan et al [40] is that Kohonen feature maps are similar to the statistical technique of *k*-means clustering, yet it has been our observation that many papers describing an SOFM do not compare the efficacy of it with another visualization technique applied to the same dataset.

2.3 Radial-basis function networks

The second most commonly used ANNs in medicine are the *radial-basis function networks* (RBFNs). An RBFN can be regarded as a type of generalized linear discriminant function, a linear function of functions that permits the construction of non-linear decision surfaces. The basis functions of an RBFN define *local responses* (*receptive fields*). Typically, only some of the hidden units (basis functions) produce significant values for the final layers. This is why RBFNs are sometimes referred to as *localized receptive field networks*. In contrast, all the hidden units of an MLP are involved in determining the output from the network (they are said to form a *distributed representation*). The receptive-field approach can be advantageous when the distribution of the data in the space of input values is multimodal. Furthermore, RBFNs can be trained more quickly than MLPs [41], but the number of basis functions required grows exponentially with the number of input nodes, and an increase in the number of basis functions increases the time taken, and amount of data required, to train an RBFN adequately.

Construction of an RBFN is a two-step process. The first step is typically performed using an unsupervised method, such as k -means clustering or an SOFM, to define the basis functions. This exploits the distribution of the classes in feature space. The second step uses supervised learning for the output-layer weights via linear optimization, which is faster than the typical training algorithms used for MLPs. The potential advantages of these two steps have been investigated for a number of medical applications. For example, an RBFN approximated well the nonlinearity of heart dynamics [42]. Each basis function provided a local reconstruction of the dynamics in the space spanned by the basis functions. Other examples include the classification of cervical-tissue fluorescent spectra [43], the identification of bacteria (clustering with respect to mass spectra) [44] and spinal disorders (clustering with respect to motion data) [45]. However, even when an SOFM suggests good discrimination between the classes of interest, an RBFN may not perform as well as an MLP. For example, when the EEG classification mentioned in Section 2.1 was repeated using an RBFN in place of the MLP, the misclassification rate increased slightly to 11.6% [16].

When conditions are such that an RBFN can act as a classifier [46, 47], an advantage of the local nature of RBFNs compared with MLP classifiers is that a new set of input values that falls outside all the localized receptor fields could be flagged as not belonging to any of the classes represented. In other words, the set of input values is novel. This is a more cautious approach than the resolute classification that can occur with MLPs, in which a set of input values is always assigned to a class, irrespective of the values.

2.4 Adaptive Resonance Theory Networks

Although not a common ANN, an *adaptive resonance theory* (ART) *network* is noteworthy. The ART process [48] can be regarded as a type of hypothesis

test [49]. A pattern presented at an input layer is passed to a second layer, which is interconnected to the first. The second layer makes a guess about the category to which the original pattern belongs, and this hypothetical identity is passed back to the first layer. The hypothesis is compared with the original pattern and, if found to be a close match, the hypothesis and original pattern reinforce each other (*resonance* is said to take place). But if the hypothesis is incorrect, the second layer produces another guess. If the second layer cannot eventually provide a good match with the pattern, the original pattern is learned as the first example of a new category. Spencer et al [50] incorporated ART into a system capable of discovering temporal patterns in ICU data and predicting the onset of haemodynamic disorders.

Although ART provides unsupervised learning, an extension called ARTMAP [51] combines two ART modules to enable supervised learning to take place. Harrison et al [52] describe how ARTMAP can be used to update a knowledge base. They do so in the context of ECG diagnosis of myocardial infarction and the cytopathological diagnosis of breast lesions.

In spite of resolving the stability/plasticity dilemma [53], the ART algorithms can be sensitive to noise [54]. Furthermore, Ripley [22] questions the virtue of the ART algorithms over adaptive k -means clustering [55].

3 The “black-box” issue

A criticism leveled against neural networks is that they are ‘black-box’ systems [30, 56] (although this can be less of a problem for the mathematically minded [57]). By this it is meant that the manner in which a neural network derives an output value from a given feature vector is not comprehensible to the non-specialist, and that this lack of comprehension makes the output from neural networks unacceptable.

There are a number of properties that we desire in a model, two of which are accuracy (the ‘closeness’ of a model’s estimated value to the true value) and interpretability. By *interpretability*, we mean the type of input-output relationships that can be extracted from a model and which are comprehensible to the intended users of the model.

An interpretable model is advantageous for several reasons. Firstly, It could be educational by supplying a previously unknown but useful input-output summary. This, in turn, can lead to new areas of research. Secondly, it could disclose an error in the model when an input-output summary or explanation contradicts known facts.

Does the lack of interpretability, as defined above, make a model unacceptable? That depends on the purpose of the model. Suppose that the choice of a statistical model for a given problem is reasonable (on theoretical or heuristical grounds), and an extensive empirical assessment of the model (for example, by cross-validation and prospective evaluation) shows that its parameters provide an acceptable degree of accuracy over a wide range of input vectors. The use of such a model for prediction would generally be approved, subject to a

performance-monitoring policy. Why not apply the same reasoning to neural networks, which are, after all, non-standard statistical models?

But suppose that we are interested in *knowledge discovery*; by this we mean the extraction of previously unknown but useful information from data. With a trained MLP, it is difficult to interpret the mass of weights and connections within the network, and the interactions implied by these. However, the goal of *rule extraction* [58] is to map the (possibly complex) associations encoded by the functions and parameters of a trained ANN to a set of comprehensible if-then rules. If successful, such a mapping would lead to an interpretable collection of statements describing the associations discovered by the ANN, which, in turn, may lead to clinical insight.

Another approach to providing interpretability is to use a *hybrid neuro-fuzzy system* [59]. These are feed-forward networks built from if-then rules containing linguistic terms based on domain knowledge, and the membership functions associated with the fuzzy rules can be tuned to data by means of a training algorithm. Both rule extraction and hybrid neuro-fuzzy systems are responses to those clinicians unwilling to use a predictive model lacking interpretability, even when the model is highly accurate.

4 Future directions

4.1 Bayesian neural computation

Whereas classical statistics attempts to draw inferences from data alone, *Bayesian statistics* goes further by allowing data to update prior beliefs via Bayes' theorem. Advantages to neural computation provided by the Bayesian framework include

- a principled approach to fitting an ANN to data via regularization,
- allowance for more than one possible MLP architecture by *model averaging*, and for multiple solutions to the training of a single MLP architecture by using a *committee of networks*,
- automatic selection of features to be used as input to an MLP from a set of candidate features (*automatic relevance determination*).

However, there are disadvantages to the Bayesian approach, namely that it tends to require more computation than the maximum likelihood approach, and it requires explicit descriptions of the prior probability distributions, which are often subjective.

Bayesian ANNs have not yet found their way into general use, but, given their capabilities, we expect them to take a prominent role in medical neural computation.

4.2 Knowledge-data fusion

If both data and domain knowledge are available, how can these two resources be used together? This is the focus of *knowledge-data fusion*. One approach is use networks in which nodes and their connections have an interpretation in the domain of interest. An example of this is seen with hybrid fuzzy-neuro systems. Another example is the concept of *probabilistic networks*; however, unlike an MLP, the nodes of a probabilistic network are related stochastically, and inference can take place in different directions along the network.

The stochastic relationships of a probabilistic network are defined by conditional probabilities at the non-terminal nodes, and for a conditional probability with a complex likelihood, an MLP may be an effective model. This idea of using MLPs at the nodes of a probabilistic networks is one approach to implementing *data fusion*, such as the integration of electrophysiology signals with both radiological images and patient records for decision support. By using knowledge-data fusion to integrate data from disparate sources, with ANNs acting as modules for the framework, virtual patient records within a hospital-wide intranet could be updated in a principled manner.

References

- [1] L. Aakerlund and R. Hemmingsen. Neural networks as models of psychopathology. *Biological Psychiatry*, 43(7):471–82, 1998.
- [2] W.W. Tryon. A neural network explanation of posttraumatic stress. *Journal of Anxiety Disorders*, 12(4):373–85, 1998.
- [3] C.A. Galletly, C.R. Clark, and A.C. McFarlane. Artificial neural networks: a prospective tool for the analysis of psychiatric disorders. *Journal of Psychiatry & Neuroscience*, 21(4):239–47, 1996.
- [4] R. Dybowski and V. Gant. Artificial neural networks in pathology and medical laboratories. *Lancet*, 346:1203–1207, 1995.
- [5] J.T. Wei, Z. Zhang, S.D. Barnhill, K.R. Madyastha, H. Zhang, and J.E. Oesterling. Understanding artificial neural networks and exploring their potential applications for the practicing urologist. *Urology*, 52(2):161–172, 1998.
- [6] W. el Deredy. Pattern recognition approaches in biomedical and clinical magnetic resonance spectroscopy: A review. *NMR in Biomedicine*, 10(3):99–124, 1997.
- [7] A. Tewari. Artificial intelligence and neural networks: Concept, applications and future in urology. *British Journal of Urology*, 80 Supplement 3:53–8, 1997.
- [8] J.P. Hogge, D.S. Artz, and M.T. Freedman. Update in digital mammography. *Critical Reviews in Diagnostic Imaging*, 38(1):89–113, 1997.
- [9] P.Y. Ktonas. Computer-based recognition of EEG patterns. *Electroencephalography & Clinical Neurophysiology – Supplement*, 45:23–35, 1996.
- [10] W. Penny and D. Frost. Neural networks in clinical medicine. *Medical Decision Making*, 16(4):386–389, 1996.
- [11] D. Itchhaporia, P.B. Snow, R.J. Almassy, and W.J. Oetgen. Artificial neural networks: Current status in cardiovascular medicine. *Journal of the American College of Cardiology*, 28(2):515–521, 1996.

- [12] R.N. Naguib and G.V. Sherbet. Artificial neural networks in cancer research. *Pathobiology*, 65(3):129–139, 1997.
- [13] M.R. Brickley, J.P. Shepherd, and R.A. Armstrong. Neural networks: A new technique for development of decision support systems in dentistry. *Journal of Dentistry*, 26(4):305–309, 1998.
- [14] R. Goodacre, M.J. Neal, and D.B. Kell. Quantitative analysis of multivariate data using artificial neural networks: A tutorial review and applications to the deconvolution of pyrolysis mass spectra. *Zentralblatt fur Bakteriologie*, 284(4):516–539, 1996.
- [15] C. Robert, C. Guilpin, and A. Limoge. Review of neural network applications in sleep research. *Journal of Neuroscience Methods*, 79(2):187–193, 1998.
- [16] L. Tarassenko. *A Guide to Neural Computing Applications*. Arnold, London, 1998.
- [17] J.S. Bridle. Probabilistic interpretation of feedforward classification network outputs, with relationships to statistical pattern recognition. In F. Fogelman-Soulie and J. Herault, editors, *Neurocomputing: Algorithms, Architectures, and Applications*, pages 227–236. Springer-Verlag, Berlin, 1991.
- [18] B.D. Ripley. Statistical aspects of neural networks. In O.E. Barndorff-Nielsen, J.L.Jensen, and W.S. Kendell, editors, *Networks and Chaos — Statistical and Probabilistic Aspects*, pages 40–123. Chapman & Hall, London, 1993.
- [19] S.-I. Amari. Mathematical methods of neurocomputing. In O.E. Barndorff-Nielsen, J.L.Jensen, and W.S. Kendell, editors, *Networks and Chaos — Statistical and Probabilistic Aspects*, pages 1–39. Chapman & Hall, London, 1993.
- [20] B. Cheng and D.M. Titterton. Neural networks: A review from a statistical perspective. *Statistical Science*, 9:2–54, 1994.
- [21] B.D. Ripley and R.M. Ripley. Neural networks as statistical methods in survival analysis. In R. Dybowski and V. Gant, editors, *Clinical Applications of Artificial Neural Networks*, chapter 9. Cambridge University Press, Cambridge. In press.
- [22] B.D. Ripley. *Pattern Recognition and Neural Networks*. Cambridge University Press, Cambridge, 1996.
- [23] S. Haykin. *Neural Networks: A Comprehensive Foundation*. Macmillan, New York, 1994.
- [24] C.M. Bishop, M. Svensen, and C.K.I. Williams. GTM: the generative topographic mapping. *Neural Computation*, 10(1):215–234, 1997.
- [25] D.J.C. MacKay. A practical Bayesian framework for back-propagation networks. *Neural Computation*, 4(3):448–472, 1992.
- [26] A. Flexer. Connectionists and statisticians, friends or foes? In J. Mira and F. Sandoval, editors, *Proceedings of the International Workshop on Artificial Neural Networks*, LNCS 930, pages 454–461. Springer, 1995.
- [27] P.H. Goodman and F.E. Harrell. Neural networks: Advantages and limitations for biostatistical modelling. Technical report, Washoe Medical Center, University of Nevada, October 1998.
- [28] L. Tarassenko. Neural networks. *Lancet*, 346:1712, 1995.

- [29] L. Tarassenko, P. Hayton, N. Cerneaz, and M. Brady. Novelty detection for the identification of masses in mammograms. In *Proceedings of the 4th IEE International Conference on Artificial Neural Networks*, pages 442–447, Cambridge, 1995. Cambridge University Press.
- [30] J. Wyatt. Nervous about artificial neural networks? *The Lancet*, 346:1175–1177, 1995.
- [31] D. Michie, D.J. Spiegelhalter, and C.C. Taylor, editors. *Machine Learning, Neural and Statistical Classification*. Ellis Horwood, New York, 1994.
- [32] T. Kohonen. Self-organizing formation of topologically correct feature maps. *Biological Cybernetics*, 43:59–69, 1982.
- [33] J. Hertz, A. Krogh, and R.G. Palmer. *Introduction to the Theory of Neural Computation*. Addison-Wesley, Reading, MA, 1991.
- [34] C.J. Lux, A. Stellzig, D. Volz, W. Jager, A. Richardson, and G. Komposch. A neural network approach to the analysis and classification of human craniofacial growth. *Growth, Development, & Aging*, 62(3):95–106, 1998.
- [35] C.I. Christodoulou and C.S. Pattichis. Unsupervised pattern recognition for the classification of EMG signals. *IEEE Transactions on Biomedical Engineering*, 46(2):169–178, 1999.
- [36] J.O. Glass and W.E. Reddick. Hybrid artificial neural network segmentation and classification of dynamic contrast-enhanced MR imaging (DEMRI) of osteosarcoma. *Magnetic Resonance Imaging*, 16(9):1075–1083, 1998.
- [37] B. Fritzke. Growing cell structures – a self-organising network for unsupervised and supervised learning. *Neural Networks*, 7:1441–1460, 1994.
- [38] A.J. Walker, S.S. Cross, and R.F. Harrison. Visualisation of biomedical datasets by use of growing cell structure networks: A novel diagnostic classification technique. *Lancet*, 354:1518–1521, 1999.
- [39] C.M. Bishop. *Neural Networks for Pattern Recognition*. Clarendon Press, Oxford, 1995.
- [40] P.V. Balakrishnan, M.C. Cooper, V.S. Jacob, and P.A. Lewis. A study of the classification capabilities of neural networks using unsupervised learning: A comparison with k-means clustering. *Psychometrika*, 59:509–525, 1994.
- [41] J. Moody and C. Darken. Fast learning in networks of locally-tuned processing units. *Neural Computation*, 1:281–294, 1989.
- [42] A. Bezerianos, S. Papadimitriou, and D. Alexopoulos. Radial basis function neural networks for the characterization of heart rate variability dynamics. *Artificial Intelligence in Medicine*, 15(3):215–234, 1999.
- [43] K. Tumer, N. Ramanujam, J. Ghosh, and R. Richards-Kortum. Ensembles of radial basis function networks for spectroscopic detection of cervical precancer. *IEEE Transactions on Biomedical Engineering*, 45(8):953–961, 1998.
- [44] R. Goodacre, E.M. Timmins, R. Burton, N. Kaderbhai, A.M. Woodward D.B. Kell, and P.J. Rooney. Rapid identification of urinary tract infection bacteria using hyperspectral whole-organism fingerprinting and artificial neural networks. *Microbiology*, 144(5):1157–1170, 1998.
- [45] J.B. Bishop, M. Szpalski, and S.K. Ananthraman D.R. McIntyre M.H. Pope. Classification of low back pain from dynamic motion characteristics using an artificial neural network. *Spine*, 22(24):2991–2998, 1997.

- [46] R. Dybowski. Classification of incomplete feature vectors by radial basis function networks. *Pattern Recognition Letters*, 19:1257–1264, 1998.
- [47] I.T. Nabney. Efficient training of RBF networks for classification. Technical report NCRG/99/002, Neural Computing Research Group, Aston University, 1999.
- [48] G.A. Carpenter and S. Grossberg. A massively parallel architecture for a self-organizing neural pattern recognition machine. *Computer Vision, Graphics, and Image Processing*, 37:54–115, 1987.
- [49] M. Caudill and C. Butler. *Naturally Intelligent Systems*. MIT Press, Cambridge, MA, 1990.
- [50] R.G. Spencer, C.S. Lessard, F. Davila, and B. Etter. Self-organising discovery, recognition and prediction of haemodynamic patterns in the intensive care unit. *Medical & Biological Engineering & Computing*, 35(2):117–123, 1997.
- [51] G. Carpenter, S. Grossberg, and J. Reynolds. ARTMAP: Supervised real-time learning and classification of nonstationary data by a self-organizing neural network. *Neural Networks*, 4:565–588, 1991.
- [52] R.F. Harrison, S.S. Cross, R.L. Kennedy, C.P. Lim, and J. Downs. Adaptive resonance theory: A foundation for ‘apprentice’ systems in clinical decision support? In R. Dybowski and V. Gant, editors, *Clinical Applications of Artificial Neural Networks*, chapter 13. Cambridge University Press, Cambridge. In press.
- [53] S. Grossberg. Adaptive pattern classification and universal recording: I. parallel development and coding of neural feature detectors. *Cybernetics*, 23:121–134, 1976.
- [54] B. Moore. ART1 and pattern clustering. In D. Touretzky, G. Hinton, and T. Sejnowski, editors, *Proceedings of the 1988 Connectionist Models Summer School*, pages 174–185, Sa Mateo, CA, 1989. Morgan Kaufmann.
- [55] D.J. Hall and D. Khanna. The ISODATA method of computation for relative perception of similarities and differences in complex and real computers. In K. Enslein, A. Ralston, and H.S. Wilf, editors, *Statistical Methods for Digital Computers*, pages 340–373. Wiley, New York, 1977.
- [56] D. Sharp. From “black-box” to bedside, one day? *The Lancet*, 346:1050, 1995.
- [57] R.A. Kemp, C. MacAulay, and B. Palcic. Opening the black box: the relationship between neural networks and linear discriminant functions. *Analytical Cellular Pathology*. 14(1):19-30, 1997, 14(1):19–30, 1997.
- [58] R. Andrews, A.B. Tickle, and J. Diederich. A review of techniques for extracting rules from trained artificial neural networks. In R. Dybowski and V. Gant, editors, *Clinical Applications of Artificial Neural Networks*, chapter 12. Cambridge University Press, Cambridge. In press.
- [59] D. Nauck, F. Klawonn, and R. Kruse. *Foundations of Neuro-Fuzzy Systems*. John Wiley, Chichester, 1997.