Accounting and Business Research. Vol. 26. No. 4. pp. 347-355. 1996

Neural Networks and Empirical Research in Accounting

Duarte Trigueiros and Richard Taffler*

Abstract—This article seeks to provide an overview of the potential role of neural network (connectionist) methodology in empirical accounting research. It highlights how the accounting task domain differs substantially from those for which neural network techniques were originally developed. A non-technical overview of neural network methodology is given, along with guidelines to help accounting researchers interested in applying these new tools to recognise the potential dangers and strengths underlying their use. An illustrative example is provided. The paper suggests research areas in accounting where neural network approaches could make a potential contribution. Explicit recommendations for prospective authors are made.

1. Introduction

Statistical modelling plays an important role in accounting research. This paper is concerned with the potential application of neural networks (connectionist models) in accounting research in the light of recent claims that such methodologies can out-perform traditional statistical approaches.

The majority of neural network studies in the business area to date have been classical financial forecasting applications.[1] Accounting applications are more limited and tend to be of a classification nature, applying neural network methodology in areas already well served by conventional statistical techniques, where the main concern is the comparative performance of the new methodology.[2]

In addition, much of this work appears in non-accounting and non-finance journals, is undertaken by computer scientists and engineers, and is in case study form (Hill et al., 1994, p. 11). Authors typically claim that neural network models out-perform conventional statistical techniques, although such comparative studies often do not use best practice in their statistical modelling (Chatfield, 1993).

This paper considers whether connectionist models are an appropriate tool for analysing accounting data and, if so, in which areas they can be most usefully applied. The next section provides a framework for comparing connectionist (pattern recognition) and accounting research task domains, and this is followed by a brief non-technical overview of neural network methodology. To illustrate the dangers of uncritical application of the technique, a review of applications in one widely addressed area, corporate failure prediction, in comparison with conventional multivariate approaches, is next undertaken. The paper then considers important methodological issues and reviews areas in empirical accounting research where neural networks may have the potential to contribute usefully. The concluding section gives advice to prospective authors.

2. A framework for comparison

Neural networks were originally developed to deal with problems in artificial intelligence such as speech, text and other pattern-recognition tasks, which conventional computing approaches were unable to solve. More recently, connectionist models have also been widely applied in such areas as defence and medicine, but in the same pattern recognition context. Are the distinguishing characteristics of the task domains where neural networks succeed a recommendation for using them in social science research?

*Duarte Trigueiros is senior lecturer in the Department of Business Studies, ISCTE, Lisbon. Richard Taffler is professor of accounting and finance at the City University Business School, London. Correspondence should be addressed to Professor Taffler, City University Business School, Frobisher Crescent, Barbican Centre, London EC2Y 8HB. This article was submitted in February 1995 and accepted in May 1996.

[1] See Hill et al. (1994); Trippi and Turban (1993); and Refenes (1995) for reviews. Finance-orientated studies of more potential relevance to accounting researchers are, for example, in the prediction of option prices (Baestaens et al., 1994, ch. 5), the modelling of arbitrage pricing theory stock returns (Refenes, 1995, ch. 7) and the prediction of stock price performance (Kryzanowski et al., 1993 and Yoon et al., 1993).

[2] Outside the well-explored failure prediction domain, neural network methodology has also been applied in the accounting area, for example, to the going concern qualification decision (Coats and Fant, 1993), bond rating (Singleton and Surkan, 1995), credit scoring (Jensen, 1992), audit litigation and audit opinion giving (Hansen et al., 1992), the analytical review process (Coakley and Brown, 1993) and the prediction of mergers (Sen et al., 1995).

Table 1
Five Comparative Characteristics of Pattern Recognition and Accounting Task Domains

Task domain	*Characteristics*				
	Complexity of the relationship	*Underlying theory describing the relationship*	*Number of explanatory variables*	*Variability explained (R^2)*	*Size of available samples*
Pattern recognition	Very complex (highly interactive and non-linear)	Well-established	Small: three or fewer	High	As large as required
Accounting	Typically simple	Often competing or weak	Often five or more	Generally low	Often limited

Table 1 lays out five important characteristics that differentiate pattern recognition and accounting relationships: complexity of the functional form, existence of underlying theory, proportion of the variability generally explained by the fitted model, minimum number of independent variables required to specify the relationship, and availability of large samples for model building.

The major obstacle in such problems as speech and text recognition is complexity, since the underlying theoretical constructs are well-established and variable explanatory power is large. In contrast, accounting data may be characterised by simple functional forms and missing (unknown or unmeasurable) variables, and theory may be weak. Explanatory power, although significant, is often low.

The last distinguishing aspect, the availability of large samples, is especially important. In pattern recognition, as in other experimental sciences, samples of required size can be generated artificially. Not so in the social sciences, where samples are in general small and case data cannot be simulated. This gravitates against trying to fit complex relationships between dependent and independent variables.

As such, pattern recognition and empirical accounting research may be viewed as quite distinct cognate areas. How can we usefully apply neural networks, developed to deal with problems in one discipline, to the other? Does the ability of neural network methodologies to model complexity, which is their main strength, entail danger when applied to simple relationships where variables may have low explanatory power?

3. Inside neural networks

There are many types of neural network; we focus here on the technique principally used in the accounting and finance literature, the 'multilayer perceptron' (MLP).[3]

Conventional linear or generalised linear modelling tools such as logistic regression or linear discriminant analysis derive a function

$$y = f(b_0 + b_1 x_1 + \dots + b_n x_n), \tag{1}$$

relating $x_1, \dots, x_n$ explanatory variables to an outcome y, minimising some error measure, where b_0, $b_1, \dots, b_n$ are the fitted coefficients and the operator f denotes the functional form. The multilayer perceptron differs from such statistical approaches only in the procedures used to obtain the functional form. Specifically:

- The MLP does not rely on distributional and other statistical assumptions in deriving model coefficients, but searches heuristically for the coefficient set that minimises total error.
- Also, instead of fitting the desired relationship by means of one, unique, linear or logistic function, MLP fits several intermediate models. As depicted in Figure 1, a given model may contain, for example, three generalised linear functions whose predicted values are fed into two other functions, which then provide the overall output value.

The latter characteristic is what distinguishes MLP from more conventional methodologies. No other modelling approach fits a relationship between dependent and independent variables by building intermediate functions and optimising the overall fit. This is an important development in model building. The heuristic providing such global optimisation is known as 'back-propagation of errors' (Rumelhart et al., 1986) and is a generalisation of the well-known hill-climbing heuristic for iterative unconstrained optimisation.

[3] Some of the principles underlying the MLP apply equally to other connectionist tools. Trigueiros (1994) provides an introduction to the self-organised map (Kohonen, 1984) and illustrates with an accounting application. Kryzanowski and Galler (1995) use Boltzmann machines in the analysis of small business financial statements.

Figure 1
A Neural Network – Each Node Represents a Separate Generalised Linear Function

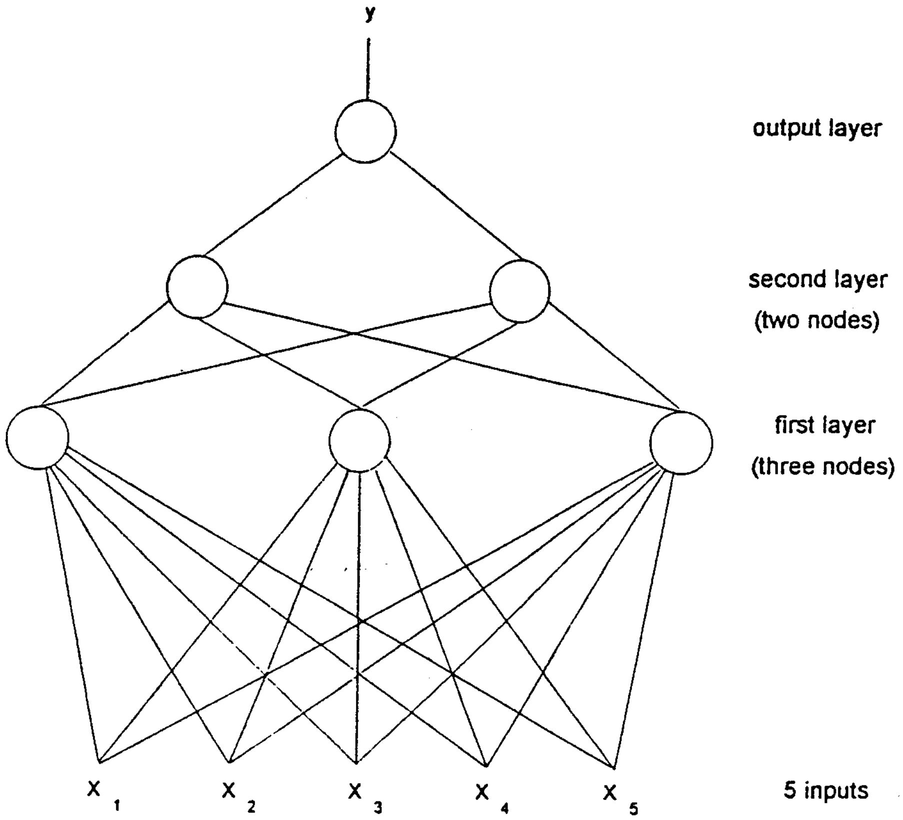

Neural network texts use the languages variously of telecommunications, neuro-biology and computer science. In the specialised literature, for example, explanatory variables are referred to as 'inputs' while values predicted by the model are known as 'outputs'. It is also usual to call the observations 'input patterns' or simply 'patterns'. Also, a given network will be characterised by 'nodes' forming 'layers' where, broadly speaking, each node is similar, typically, to an individual logistic regression.[4,5] The coefficients of these regressions are variously known as 'weights', 'connections', 'synaptic links' or a mixture of these. The appendix provides a glossary of relevant terms.

Other characteristics of neural network methodology are also relevant to accounting researchers; some are beneficial. some are limitations:

- a key strength of the MLP is its ability explicitly to handle variable interactions and other forms of non-linearity;
- since a neural network with enough nodes can approximate whatever functional form best fits the sample data (see Hill et al., 1994, p. 6), generalisation needs to be treated with care because of the tendency towards overfitting. Rigorous out-of-sample testing of any such model is thus even more important than with conventional statistical approaches;[6]
- building a neural network requires considerable computer power and skill in guiding the algorithm;

[4] Most neural network applications use sigmoidal, such as logistic or hyperbolic tangent, or similar smooth threshold functions to obtain the required non-linear formulation (the f in equation 1). The logistic formulation is given by $y = \frac{1}{1+e^{-X}}$ and the hyperbolic tangent by $y = \frac{e^{X}-e^{-X}}{e^{X}+e^{-X}}$ where $X = \sum b_i x_i$.

[5] Figure 1 illustrates a network with five inputs and three layers of 3, 2 and 1 nodes respectively.

[6] A rule-of-thumb generally used for avoiding overfitting is to restrict the number of coefficients to a maximum of 10% of the number of cases.

- neural network methodologies do not yet provide adequate significance and hypothesis tests;[7]
- neural networks are difficult, if not impossible, to interpret or explain conceptually.

Some of these require further elaboration.

3.1 Heuristic Search

Neural networks begin searching for minimum error by setting their coefficients or weights to random values. An observation is then sampled randomly without replacement from the data set and each coefficient in the network is modified by an arbitrary small value to reduce the error between the expected value and the actual value (the output of the network).[8] This is termed a 'presentation'. Next, another observation is randomly sampled from the remaining data set and the same procedure is repeated with regard to this second case, treated independently of the first. This process continues until the entire set of observations is exhausted. The whole set of presentations is then repeated. After a large number of these cycles, the network coefficients, initially random values, tend asymptotically towards describing the underlying relationship. Such an iterative procedure eventually leads to minimum error in some pre-defined sense, thus, in effect, yielding results similar to traditional statistical tools. In connectionist terminology, this optimisation process is termed 'learning' or 'training'.[9]

Heuristic search procedures make MLP less dependent for error minimisation on certain underlying statistical assumptions about a data set. Nonetheless, MLP suffers from exactly the same problems arising from influential or extreme cases as traditional statistical methods; this is often not recognised when acknowledging the distribution-free strengths of MLP.

3.2. Non-linearity

The principal benefit of adopting neural network methodology in practice lies in its ability to model complex interactions between independent variables. For example, in a classification type problem, MLP uses more than one boundary for separating groups. Since each node or intermediate function in a neural network is a partial classifier, it contributes one boundary to the overall classification. The final function utilises these boundaries in building a more complex frontier.

An important example is known as the 'exclusive OR' (XOR) classification problem which involves, in its simplest form, two groups of cases, eg., groups A and B, and two explanatory variables eg., x_1 and x_2. As illustrated in Figure 2, when both x_1 and x_2 are either large or small the group is B. When the same variables go in opposite directions, then the group is A.[10]

Whereas, in theory, only two independent variables should be sufficient to separate the two groups A and B, conventional classifiers such as linear or quadratic discriminant analysis are ineffective since a single boundary, either linear or non-linear, is inadequate.

The XOR classification problem is of paramount importance in pattern recognition applications, and the success of MLP here stems from its ability to solve it. Using MLP in classification task domains where there is no plausible reason to expect XOR interactions will substantially reduce its potential utility when compared with conventional statistical methodologies.

Another, although less important, benefit is the ability of a neural network to deal with piece-wise non-linear relationships explicitly and thus perform better than conventional polynomial models out of sample range (Hill et al., 1994, p. 6).

3.3. Statistical Considerations

Issues important in statistical modelling, such as sampling, and significance and hypothesis testing, are less relevant in pattern recognition. Consequently, statistical tests have yet to be developed beyond a rudimentary level in connectionist methodology:

> '...since the ANN [artificial neural network] model form is non-linear in the model coefficients, the normal probability model is not applicable. Consequently, ANNs do not have parametric statistical properties (eg., they do not have individual coefficient or model significance tests based on the t or F distributions)' (Gorr et al., 1994, p. 19).

The concern is that, as such significance tests are developed, some of the benefits of the neural network approach (eg., distribution-free optimisation) might become less evident.

3.4. Building and Interpreting Models

The building of MLP models is difficult to carry out, and the resulting models and their outputs are

[7] Although this is a serious problem for the social scientist, it is not necessarily a major drawback in the experimental sciences.

[8] It is this updating of weights that is known as back-propagation.

[9] The technique described in this paragraph is the most widely used, but by no means the only one available. It is termed 'stochastic learning' in contrast to 'batch learning', where coefficient updating is carried out only at the end of each cycle. See Azoff (1994, ch. 4) for a non-technical description of learning procedures used in the forecasting of time-series.

[10] The same problem can be generalised to multivariate, multi-group situations.

Figure 2
The XOR Classification Problem

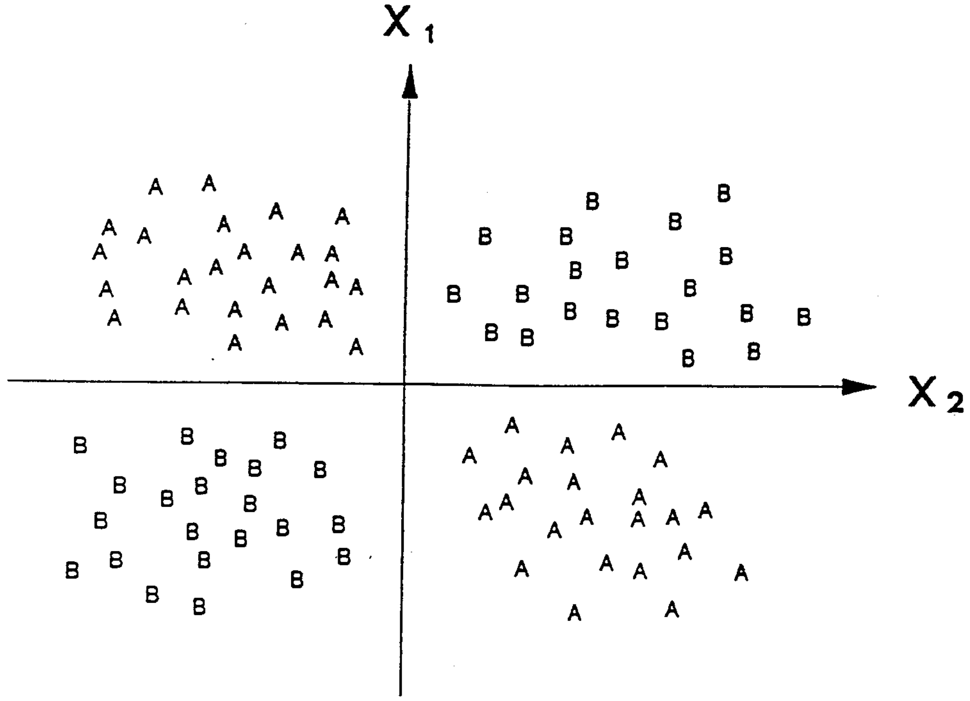

seldom directly interpretable.[11] MLP training requires extensive computing power, especially when the number of observations is large, and direct intervention by the model builder, which, to be effective, presupposes practice and experience.

The appropriate number of layers and nodes is often application-specific and, in practice, needs to be determined on a trial and error basis. There are also many rules-of-thumb, often picked up by experience, which are important for avoiding such problems as false minima or oscillations during training, for speeding up the iterative search for a global solution, and for obtaining parsimonious models through coefficient (weight) or node pruning.[12]

Once a neural network model is developed, only rudimentary methods can be applied to assist in model interpretation (see, for example, Sen et al., 1995, p. 337); however, in general, '...insights from the behaviour of individual model components explaining estimates or forecasts are difficult to obtain' (Gorr, 1994, p. 2).

Whereas neural network models are frequently criticised for their opacity, lack of interpretability should be seen in the light of the complexity of the problem. If the modelled relationship is itself complex, then there is no reason to believe that the fitted model will be transparent. However, the question remains as to whether MLP is too powerful an instrument when used for modelling the relatively simple relationships conventionally found in accounting.

> '...where relatively few explanatory measures are available for making predictions, simple models are often the best, and perhaps no amount of sophisticated methodology will make any improvement....Thus, in cases where there is no underlying structure in the available data, ANN is simply not going to perform any better than the simpler models' (Gorr et al., 1994, p. 19).

4. An application: predicting financial distress

The most widespread application of neural network methodology in the accounting domain to date has been in bankruptcy prediction models, with more than 25 such papers published at the

[11] Azoff (1994, ch. 4), Baestaens et al. (1994, ch. 1) and Refenes (1995, chs. 2, 3) provide a good survey of current knowledge on network building procedures.

[12] For example, as false or *local* optima may be picked up during network learning, it is advisable to repeat training with different sets of randomly determined initial coefficient values.

Table 2
Summary of the Characteristics of Representative Failure Prediction Models using Neural Networks

Paper	*Task*	*Selection of variables*	*Number of coefficients*	*Sample selection*	*Number of observations used in fitting model (failed: non-failed)*
Salchenberger, Cinar and Lash (1992)	S & L associations failure	Previous studies	5 × 3	Matched	100 + 100
Tam and Kiang (1992)	Bank bankruptcy	Previous studies	19 × 10	Matched	59 + 59
Sharda and Wilson (1993)	Corporate bankruptcy	Altman (1968)	5 × 10	Matched	65 + 64
Coats and Fant (1993)	Firm financial distress as perceived by auditor reports	Altman (1968)	5 × 8	Matched on a 2 to 1 basis	47 + 94
Rahimian et al. (1993)	Firm failure	Altman (1968)	5 × 5	Matched	38 + 36

time of writing. Such studies usually compare a neural network approach with traditional linear discriminant analysis (LDA) or logit methodologies, and the authors, almost without exception, report an increase in the rate of correct classification of firms, banks or savings and loans associations as failed and non-failed using MLP.

Typical recent studies are Salchenberger et al. (1992), Tam and Kiang (1992), Sharda and Wilson (1993), Coats and Fant (1993) and Rahimian et al. (1993). However, comparisons made between neural network and multivariate statistical methodologies are problematic, and inadequate attention is paid to the extant literature.[13] In each case, disproportionate effort is expended in fitting the MLP technique compared with applying the comparative statistical technique (Chatfield, 1993). Classification results, independent of methodology, are uniformly poorer than those reported in conventional earlier studies, perhaps highlighting the respective authors' lack of understanding of the task domain, which is essential for valid application of any model building methodology.[14]

Table 2 summarises the characteristics of these five studies. Only in the case of Salchenberger et al. (1992) are the number of coefficients fitted (15) less than 10% of the observations (see footnote 6 *supra*). The most extreme case is Tam and Kiang (1992), who use 19 highly collinear input variables with 10 nodes (190 coefficients fitted) and only 118 cases in model fitting! Numbers of coefficients derived compare with only four or five in the case of conventional LDA Z-Score models (eg., Altman, 1968; Taffler, 1983).

In comparison, a carefully undertaken large sample study such as that by Altman et al. (1994), which also is a real world application not a methodological 'test-bed', finds little or no difference in classification performance between neural networks and conventional multivariate statistical techniques.[15]

Their best neural network, which has nine inputs and, significantly, only five intermediate functions or nodes (45 coefficients estimated from a data set of 800 firms) performs no better than their equivalent 11 variable linear discriminant model in classifying the held-out cases.[16]

Altman et al. point to the long processing time required to fit neural network models, and the arduous trial and error process required to discover the best model structure, as well as stressing the trap of overfitting. In addition, derived weights are not readily interpretable as with discriminant or logit analysis. They also mention the problems for the financial analyst posed by illogical network behaviour.[17]

[13] For example, Rahimian et al. (1993, p. 161) dismiss traditional statistical procedures thus: '...hence the predictions of discriminant analysis or dummy regression analysis should be taken with a grain of salt.'

[14] In addition, in all the five papers, samples are selected on a matched basis leading to the likelihood of bias in any hold-out tests, and variable selection tends to be arbitrary.

[15] A parallel study, equally relevant from a methodological vantage point although in a different task domain, is Gorr et al. (1994), who develop comparative models for the prediction of student grade point average.

[16] Type I errors of 10.9% and only 4.9% and Type II errors of 6.4% and 9.7% respectively.

[17] Changes in the output variable are not monotonically related to small perturbations in input variables considered one at a time. This phenomenon is consistent with the existence of a degree of overfitting and sample bias in their derived model

5. Discussion

What contribution can neural network methodology make to accounting research and add to our understanding of accounting issues?

Neural networks are not a substitute for an understanding of the task environment and may best be applied in complex situations where there is no theory to assist the model builder (Gorr, 1994, p. 3). Such tools are only ever likely to dominate conventional statistical models when strong non-linearities, and, most importantly, interactions between independent variables, are present. Other criteria for effective model development, such as the assumption of stationarity, absence of multicollinearity and influential cases and, in particular, model parsimony, equally apply.

5.1. Methodological Issues

Empirical research in accounting typically takes the following form (Tomkins and Groves, 1983, p. 362):

i) Theories are formulated in terms of the relationships between categories, and based on a review of the existing academic literature.
ii) The theory is used to establish a research problem.
iii) The problem is resolved into hypotheses, and dependent and independent variables identified.
iv) Precise and highly-structured predetermined procedures for data collection are established. The data collected are usually in quantitative form.
v) The data are subjected to mathematical or statistical analysis, leading to an almost exclusively quantitative validation of the hypotheses being tested.

In scientific method, apparent performance of a developed model is not the sole objective of the research process. The goal is a better understanding of the underlying accounting issues of concern to the researcher (Chua, 1986, p. 608). The analytical tool used is not of intrinsic interest itself but only a means for elucidating the underlying phenomena.

Moreover, the Popperian doctrine of falsifiability (see, for example, Chua, 1986, p. 607) requires the ability to replicate experimental findings. Because of their strictly heuristic nature, and the extensive requirement for model-builder interaction with the optimising algorithm, neural network results are difficult to replicate, even on the same data set. The opacity of the fitted models can often also add to the problem of understanding underlying relationships.

5.2. Potential Useful Contributions

There may, nonetheless, be cases where neural network methodology can contribute to the accounting researcher's work, when validly applied: for instance, where non-linearity, in particular interactions of an XOR nature,[18] is an underlying facet of the relationship being modelled and, most importantly, where large samples are available to allow its manifestation. The takeover literature illustrates the potential for XOR interactions. The poor results of, for example, Palepu (1986) when predicting takeover targets may reflect the attempt to impose a single boundary between groups where two or more are required. If we believe, for instance, that companies are bid targets for combinations of different reasons, then we require a methodology that allows us to model this appropriately.

Drawing on the substantial literature applying connectionist approaches of a pattern recognition nature in financial forecasting, such as real-time market trading and technical analysis (eg., Baestaens, 1994, ch. 5 and Trippi and Turban, 1993, part 5), we may speculate that related methodology could also have potential application in forecasting accounting variable time-series such as cash flows or earnings. This may particularly be so if hidden interactions or non-linearities pertain in the underlying relationship (Tippett, 1990) such as with half-yearly or quarterly data (Hill et al., 1994, p. 8).

Another potentially useful application of neural network methodology could be in predicting stock returns from accounting and stockmarket data where complex relationships exist and theory is not always helpful. Extant studies assume monotonic linear or log-linear relationships between stock returns and firm factors[19] (eg., Ou and Penman, 1989; Chan and Chen, 1991; Fama and French, 1992; Holthausen and Larcker, 1992; and Lakonishok et al., 1994) whereas this is not necessarily so. Potentially complex interactions between predictor variables are either ignored or handled in a very limited manner. In addition, conventional methodologies used to forecast average returns (eg., the assumption of constant β) are not adequate (Ball and Kothari, 1989). The availability of large sample sizes also suggests connectionist methodologies could play a useful role here.

We may also speculate that, in contrast to the conventional habit of using stepwise regressions and principal component analysis to determine the appropriate independent variables for a neural network, the reverse is more logical. Neural network methodology could be useful for prospective analysis where there is no prior theory to guide model formulation prior to forming hypotheses,

[18] Such as the existence of nominal explanatory variables.
[19] Such as book/market, size, β, dividend yield, P/E, gearing and other accounting-based measures.

and structuring more rigorous analytical models.[20] For example, when studying industry homogeneity, Berry and Trigueiros (1993, p.121) use MLP to select a parsimonious set of variables to be used subsequently in their LDA models.

6. Recommendations to authors

Prospective authors interested in applying neural network methodology to accounting problems should concentrate on using these tools to enhance understanding rather than solely for empirical performance comparisons. The issues of contribution to theory and real advances in knowledge are paramount. Specifically:

- Before falling back on connectionist approaches, authors should first ascertain whether poor empirical performance of conventional statistical approaches is due to their inability to deal appropriately with the complexity of the underlying relationships being studied, or rather through lack of key predictors. An equally plausible reason may be that, given the set of independent variables available, there is nothing more that can be explained independent of methodology.
- Authors should go beyond case studies, simply describing, at best, their neural network methodologies and individual results, and provide enough information to permit replication of findings in related situations.
- Before submission, authors must ensure terminology and style are intelligible to the readership of accounting journals. Neural network papers are currently aimed at a different audience, ie, engineers, applied computer scientists or the operational research community.
- Neural networks exemplify the data mining problem. Therefore, when using connectionist modelling or similar powerful tools, authors should carefully guard against overfitting. The precepts of common sense and parsimony underlying good practice in statistical modelling apply also to neural networks.
- An understanding of the salient differences between accounting research and that in the experimental sciences is mandatory when adopting neural network techniques. For example, sampling issues, statistical inference (eg., significance tests) and hypothesis testing, are far less important in pattern recognition tasks.

This paper should not be taken as a criticism of neural network methodology per se, but only the manner in which it has frequently been applied to date. We believe that the principles underlying our comments are applicable, not only to neural networks, but also to any powerful modelling tool. Such novel techniques should not be adopted uncritically by accounting researchers without first fully understanding their potential contributions and dangers.

References

Azoff, E. (1994). *Neural Network Time Series Forecasting of Financial Markets* (New York: Wiley).

Altman, E. (1968). 'Financial Ratios, Discriminant Analysis and the Prediction of Corporate Bankruptcy'. *Journal of Finance*, Vol. 23, No. 4, pp. 589–609.

Altman, E., Marco, G. and Varetto, F. (1994). 'Corporate Distress Diagnosis: Comparisons Using Linear Discriminant Analysis and Neural Networks (the Italian Experience)'. *Journal of Banking and Finance*, Vol. 18, No. 3, pp. 505–529.

Ball, R. and Kothari, S.P. (1989). 'Non-stationary Expected Returns: Implications for Tests of Market Efficiency and Serial Correlation in Returns'. *Journal of Financial Economics*, Vol. 25, No. 1, pp. 51–74.

Baestaens, D., Bergh, W. and Wood, D. (1994). *Neural Network Solutions for Trading in Financial Markets* (London: Pitman).

Berry, R. and Trigueiros, D. (1993). 'Applying Neural Networks to the Extraction of Knowledge from Accounting Reports: A Classification Study', in Trippi, R. and Turban, E. (ed.), *Neural Networks in Finance and Investing* (Chicago: Probus), pp. 103–123.

Chan, K.C. and Chen, N.F. (1991). 'Structural and Return Characteristics of Small and Large Firms', *Journal of Finance*, Vol. 46, No. 4, pp. 1,467–1,484.

Chatfield, C. (1993). 'Neural Networks: Forecasting Breakthrough or Passing Fad?', *Internat. Journal of Forecasting*, Vol. 9, No. 1, pp. 1–3.

Chua, W. (1986). 'Radical Developments in Accounting Thought', *Accounting Review*, Vol. 61, No. 4, pp. 601–632.

Coakley, J. and Brown, C. (1993). 'Artificial Neural Networks Applied to Ratio Analysis in the Analytical Review Process', *Information Systems in Accounting, Finance & Management*, Vol. 2, No. 1, pp. 19–40.

Coats, P. and Fant, F. (1993). 'Recognizing Financial Distress Patterns Using a Neural Network Tool', *Financial Management*, Vol. 22, No. 3, pp. 142–165.

Fama, E.F. and French, K.R. (1992). 'The Cross-section of Expected Returns', *Journal of Finance*, Vol. 47, No. 2, pp. 427–465.

Glaser, B. and Strauss, A. (1967). *The Discovery of Grounded Theory: Strategies for Qualitative Research* (New York: Aldine Pub.).

Gorr, W. (1994), 'Research Prospective on Neural Network Forecasting.' *Internat. Journal of Forecasting*, Vol. 10, No. 1, pp. 1–4.

Gorr, W., Nagin, D. and Szczypula, J. (1994), 'Comparative Study of Artificial Neural Network and Statistical Models for Predicting Student Grade Point Averages', *Internat. Journal of Forecasting*, Vol. 10, No. 1, pp. 17–34.

Hansen, J., McDonald, J. and Stice, J. (1992), 'Artificial Intelligence and Generalised Qualitative Response Models: An Empirical Test on Two Audit Decision-Making Domains', *Decision Science*, Vol. 23, pp. 708–723.

Hill, T., Marquez, L., O'Connor, M. and Remus, W. (1994), 'Artificial Neural Network Models for Forecasting and Decision Making', *Internat. Journal of Forecasting*, Vol. 10, No. 1, pp. 5–15.

[20] In this context, it should be recollected that Glaser and Strauss (1967) strongly argue for the relevance of grounded theory, the generation of theory from data through induction, also to quantitative measures: '...*if quantitative data is handled systematically...the analyst will indeed find rich terrain for discovering and generating theory*' (p. 220, their italics). They also point out that: 'Statistical tests of significance of an association between variables are not necessary when the discovered associations...are used for suggesting hypotheses' (p. 200).

Holthausen, R.W. and Larcker, D.F. (1992), 'The Prediction of Stock Returns using Financial Statement Information', *Journal of Accounting and Economics*, Vol. 15, pp. 373–411.
Jensen, H. (1992), 'Using Neural Networks for Credit Scoring', *Managerial Finance*, Vol. 18, No. 6, pp. 15–26.
Kohonen, T. (1984), 'Self-organization and Associative Memory' (Berlin: Springer-Verlag).
Kryzanowski, L. and Galler, M. (1995), 'Analysis of Small-business Financial Statements', *Journal of Accounting, Auditing and Finance*, Vol. 10, No. 1, pp. 147–172.
Kryzanowski, L., Galler, M. and Wright, D.W. (1993), 'Using Artificial Neural Networks to Pick Stocks', *Financial Analysts Journal*, July/August, pp. 21–27.
Lakonishok, J., Shleifer, A. and Vishny, R.W. (1994), 'Contrarian Investment, Extrapolation and Risk', *Journal of Finance*, Vol. 49, No. 5, pp. 1,541–1,578.
Ou, J. and Penman, S. (1989), 'Financial Statement Analysis and the Prediction of Stock Returns', *Journal of Accounting and Economics*, Vol. 11, pp. 295–329.
Palepu, K. (1986), 'Predicting Takeover Targets, a Methodological and Empirical Analysis', *Journal of Accounting and Economics*, No. 8, pp. 3–35.
Rahimian, E., Singh, S., Thammachote, T. and Virmani, R. (1993), 'Bankruptcy Prediction by Neural Network', in Trippi, R. and Turban, E. (ed.), *Neural Networks in Finance and Investing* (Chicago: Probus), pp. 159–171.
Refenes, P. (1995), *Neural Networks in the Capital Markets* (London: Wiley).
Rumelhart, D., Hinton, G. and Williams, R. (1986), 'Learning Internal Representations by Error Propagation', *Parallel Distributed Processing*, Vol. 1 (The MIT Press).
Salchenberger, L., Cinar, E. and Lash, N. (1992), 'Neural Networks: A New Tool for Predicting Thrift Failures', *Decision Sciences*, Vol. 23, No. 4, pp. 899–916.
Sen, T., Tech, V. and Sen, N. (1995), 'Predicting Corporate Mergers', in Refenes, P. (ed.), *Neural Networks in the Capital Markets* (New York: Wiley), pp. 325–340.
Sharda, R. and Wilson, R. (1993), 'Performance Comparison Issues in Neural Network Experiments for Classfication Problems', Nunamaker, E. and Sprague, R. (ed.), *Organisational Systems and Technology: Proceedings of the 26 Hawaii International Conference on System Sciences (IV)* (IEEE Computer Society Press), pp. 649–657.
Singleton, J. and Surkan, A. (1995), 'Neural Networks for Bond Rating Improved by Multiple Hidden Layers', in Refenes, P. (ed.), *Neural Networks in the Capital Markets* (London: Wiley) pp. 301–307.
Taffler, R. (1983), 'The Assessment of Company Solvency and Performance Using a Statistical Model', *Accounting and Business Research*, Vol. 13, No. 52, pp. 295–307.
Tam, K. and Kiang, M. (1992), 'Managerial Applications of Neural Networks: The Case of Bank Failure Predictions', *Management Sciences*, Vol. 38, No. 7, pp. 926–947.
Tippett, M. (1990), 'An Induced Theory of Financial Ratios', *Accounting and Business Research*, Vol. 21, No. 81, pp. 77–85.
Tomkins, C. and Groves, R. (1983), 'The Everyday Accountant and Researching his Reality', *Accounting, Organizations and Society*, Vol. 8, No. 4, pp. 361–374.
Trigueiros, D. (1994), 'Incorporating Complementary Ratios in the Analysis of Financial Statements', *Accounting Management and Information Technologies*, Vol. 4, No. 3, pp. 149–162.
Trippi, R. and Turban, E. (1993), *Neural Networks in Finance and Investing* (Chicago: Probus).
Yoon, Y., Swales, G. and Margavio, T. (1993), 'A Comparison of Discriminant Analysis Versus Artificial Neural Networks', *Journal of the Operational Research Society*, Vol. 44, No. 1, pp. 51–60.

Appendix

Glossary

Neural network terminology	**Statistical modelling terminology**
Neural network	Model
Synapses, weights, connectivities, etc.	Coefficients of the model
Inputs	Independent variables
Outputs	Dependent variables
Outcome or target	Expected value
Node	Logistic regression
Hidden layer	Intermediate set of logistic regressions
Learning	Coefficient estimation
Supervised learning	Regression, discriminant analysis, etc.
Unsupervised learning	Principal components and cluster analyses
Architecture	Model description (eg., number of nodes and layers)
Convergence	In-sample performance
Generalisation	Out-of-sample performance

Reproduced (in modified form) with thanks to Dr Paul Refenes, London Business School.