VALIDATING A CONNECTIONIST MODEL OF FINANCIAL DIAGNOSIS

PER E. PEDERSEN
Department of Economics and Business Administration
Agder College
Grooseveien 36
4890 Grimstad, Norway
E-mail: Per.Pedersen@hia.no
http://ulke.grm.hia.no/~perpe

Several studies have demonstrated a superiority of neural networks as models for financial diagnosis. It has been proposed that the ability of these models to represent intermediate abstractions is a main reason for their superior performance. It has also been proposed that the represented intermediate abstractions resemble the diagnostic concepts used by skilled diagnosticians, and thus, that they have cognitive relevance. In this paper, we investigate these propositions applying an experimental methodology to obtain valid data sets of financial diagnoses. A multilayered perceptron model is developed and validated using cross validated performance measures and analyses of error term distributions and internal representations. The hidden units of the connectionist model represent intermediate abstractions explaining the model's superior performance, but the cognitive relevance of these intermediate abstractions is not obvious.

1 Introduction

Financial diagnosis is the classification task performed when a subject makes a judgment of the financial situation of the firm based upon information from the financial statement (Methlie, 1987). This task is performed in several contexts, such as bankruptcy prediction, going concern judgment and loan decision contexts. The diagnosis task constitutes an important basis for prediction and judgment in all these contexts and has many similarities across the different contexts.

Several authors have proposed that neural networks outperform traditional statistical techniques in predicting the environmental outcome of these contexts (e.g., Tam, 1991; Tam & Kiang, 1992; Wilson & Sharda, 1994; Yoon et al., 1993). Some researchers have argued that the superior performance of the neural network models can be explained by their ability to represent intermediate features or abstractions of relevance to the task (e.g., Raghupathi et al., 1991; Srivastava, 1993). Some researchers have also proposed that these abstractions resemble those used by skilled diagnosticians (e.g., Berry & Trigueiros, 1993; Surkan & Singleton, 1990;

Singleton & Surkan, 1995). Thus, three propositions have been made regarding the properties of neural networks as predictive models for financial diagnosis. First, performance of these models are better than comparable benchmarks of traditional models. Second, the ability to represent and utilize intermediate abstractions are the main reason for this superior performance. Third, these abstractions have what we will term "cognitive relevance". Internal representations are considered to have cognitive relevance if they resemble intermediate abstractions, such as strategies, heuristics, rules, variables, concepts, prototypes or classes, previously documented to be relevant in either prescriptive or behavioral theory of financial diagnosis.

As referred above, the first proposition has been tested and confirmed by several researchers in predictive studies, but the last two propositions have received less attention. There may be several reasons for this lack of attention. First, testing these propositions implies a shift in focus from prediction to knowledge description and cognition. Second, what is considered relevant intermediate abstractions in these task contexts is not clear (Jones, 1987, p. 136), so that validation is difficult. Third, it has not been established if knowledge of the intermediate abstractions used by diagnosticians will improve the predictions of a model, since the predictive ability of diagnosticians in these task contexts is much debated (e.g., Libby, 1975; Casey, 1980, Chalos, 1985; but see Hopwood et al. 1994). Finally, the interest in analyzing internal representations found in early neural network studies (e.g., Gorman & Sejnowski 1988; Hanson & Burr, 1990) have been displaced by focus on testing predictive ability and introducing new models and algorithms. Despite these reasons, propositions on neural network models' internal representations should not go unattended.

2 Theory and model

The financial diagnosis task has been studied with different approaches. Traditional studies in accounting and finance have taken a predictive approach, and have used a wide variety of statistical and other models to predict the environmental outcome in different contexts of the financial diagnosis task (see e.g., Altman et al., 1981). Neural network models add a new technique to the "toolbox" of these researchers. Behavioral studies in accounting and finance have taken one of two different approaches. Judgment modeling studies have used models similar to those used in the predictive approach when modeling human diagnosticians' judgments of the outcomes in different task contexts of financial diagnosis, (e.g., Libby, 1975; Casey, 1980, Chalos, 1985). Cognitive studies are either experimental or descriptive in orientation. Experimental cognitive studies use cognitive theory to predict the effect on diagnosticians judgments of manipulating independent variables, such as information content and form. Consequently, the result of these studies is knowledge of how general cognitive theory applies to the financial diagnosis task. Descriptive cognitive studies follow the information processing theory tradition of Newell and Simon (1972) and apply protocol analysis to the verbal utterances of human diagnosticians during diagnosis. A main result of these studies is the formulation of productions implemented as rules in expert systems (e.g., Bouwman et al., 1987; Biggs et al., 1993).

All these approaches have contributed to the identification of diagnostic knowledge and intermediate abstractions relevant to the financial diagnosis task. Typically, these abstractions are formalized in diagnostic concepts. Table 1 shows an overview of some of the most widely applied diagnostic concepts used in different financial diagnosis task contexts.

Table 1. Diagnostic concepts in selected studies of financial diagnosis.

Concepts: Studies:	Leverage	Profitability	Liquidity	Debt coverage	Asset balance	Cash position	Size	Interest coverage	Capital turnover	Task context[1]
Libby, 1975		X	X		X	X				BP
Abdel-khalik & El-Sheshai, 1980			X	X						BP
Chalos, 1985	X	X	X		X	X				BP
Simnett & Trotman, 1989	X	X	X	X	X			X		BP
Kida, 1980	X	X	X			X			X	G
Hopwood et al., 1994	X	X	X		X	X	X			G
Rodgers, 1991	X	X	X							L
Altman, 1968	X	X			X				X	BP
Ohlson, 1980	X	X	X	X	X		X			BP
Frydman et. al, 1985	X			X	X	X		X		BP
Zavgren & Friedman, 1988	X	X	X			X			X	BP
Mutchler, 1985	X	X	X	X						G
Srinvasan & Kim, 1987	X	X	X				X			L
Kaplan & Urwitz, 1979	X	X		X			X	X		B
Ziebart & Reiter, 1992	X	X					X	X		B

If we accept the importance of intermediate abstractions in the financial diagnosis task, a model of the task can be suggested. Instead of a simple classification model in which a diagnosis is given directly from an evaluation of the stimulus dimensions of the financial statement, a somewhat more complex classification model is suggested. The next step in adding complexity to a classification model is by introducing intermediate abstractions, often termed abstract stimulus dimensions (Smith & Medin, 1981). These intermediate abstractions are formed by evaluating

[1] BP = Bankruptcy prediction, B = Bond rating, G = Going concern judgement and L= Loan decision.

or transforming the original stimulus dimensions of the financial diagnosis. Often it is assumed that several original stimulus dimensions of the financial statement can be combined to form an evaluation of an intermediate abstraction, such as when information on contribution margins and return on total assets is combined to form an evaluation of the profitability of a firm. If model complexity is increased, a hierarchy of intermediate abstractions can be introduced, in which new intermediate abstractions are formed on the basis of evaluations of other such abstractions. However, the simplest form only assuming one set of intermediate abstractions is often used. In traditional models, the intermediate abstractions are represented by concepts believed to be of diagnostic relevance.

A theoretical basis for identifying such diagnostic concepts can be found in prescriptive theory of financial analysis. However, their main basis is empirical, resulting from studies applying principal components analysis to financial statement information in order to find the "underlying dimensions" of this material (Pinches et al., 1973; Gombola & Ketz, 1983). Whether these "underlying dimensions" are diagnostic stimulus dimensions is, however, not clear.

A model introducing a set of represented abstractions intermediating original stimulus dimensions and diagnostic response is very similar to the simplest multilayered neural network models including one layer of hidden units representing a set of intermediate abstractions. Even though many different neural network models can be applied, and these can form different types of intermediate abstractions, the most widely known multilayered perceptron including one layer of hidden units is applied here as a model of the financial diagnosis task. Since financial diagnosis is a cognitive task with a behavioral outcome, we term the model a connectionist model of financial diagnosis.

3 Research design

To provide a data set of valid financial diagnoses, a controlled experiment was set up. Since information on small and medium sized firms' market values is often lacking, diagnosis depends even more upon financial statement information for such firms. A random sample of 75 firms was selected from an established small firms register maintained at the Norwegian School of Economics and Business Administration. Financial statement information from two consecutive years was collected and balance statements, income statements, funds flow statements and selected ratios were presented in a booklet to 108 subjects participating in a higher level financial analysis course for auditors. 89.9 % of the subjects had prior auditing experience and were familiar with the financial diagnosis task. The rest of the subjects had other accounting experience. Thus, the subjects were considered skilled diagnosticians.

To control for individual variation in diagnostic behavior, all subjects received a booklet of three financial statements for diagnosis. The procedure of distribution was randomized, so that on an average 4.3 skilled subjects diagnosed each firm. The subjects used a response form to characterize the profitability, financing, liquidity, and leverage of each firm on predefined 5-point Likert scales. A similar evaluation and a verbal characterization of the general financial situation of the firm were also given by the subjects.

From the data, a composite judge measure of the subjects' evaluation of the general financial situation was computed as the average value of their response on the general financial situation indicators. The use of composite judge measures has been recommended in financial diagnosis studies to reduce individual variation and increase diagnostic predictive value (Libby, 1981). In addition, the composite judge measure transforms the response of a classification task into an approximately interval scaled variable with good distribution properties without changing the classificatory character of the task itself.

Each subject also indicated the most important cues of the financial statements. Of the 108 subjects, 97.2 % indicated the use of one or more cues in any diagnostic area. Of these, 83.8 % indicated that cue values of two consecutive years, or the relationship between them, were used. Since the cue values of two consecutive years were highly correlated, the majority of the subjects indicated sensitivity to correlated stimulus dimensions. 55.4 % of all cues indicated were cues representing financial ratios, thus financial ratios seemed to represent the most important stimulus dimensions of the task. The main reason for the subjects using these ratios is probably that they are approximately firm size independent. By selecting 16 financial ratios with cue values from two consecutive years as preliminary input variables in the connectionist model, the most important stimulus dimensions were represented by variables with size independence and good distribution properties.

4 Simulation design

Clearly, 75 data sets of financial diagnoses are too few observations to properly set the weight values of a connectionist model with 32 input variables and one response variable. Since the number of valid diagnoses were restricted to 75, we reduced the number of input variables in the model, using sensitivity based selection of a smaller input variable set. By applying sensitivity analysis procedures similar to Moody and Utans (1995), a constrained set of 12 input variables giving the largest increase in training error when excluded, was selected. Not surprisingly, these input variables represented values from two consecutive years of the six most important financial cues as indicated by the diagnosticians.

A traditional multilayered perceptron model was set up applying the original principles of Rumelhart et al. (1986). The traditional asymmetric sigmoid output functions was used. Inputs and outputs were scaled to the [0,1] scale, but no other representational transformations were made. Consequently, the model had local variables representations at the input and output layers.

For model estimation, the original epoch-based learning of Rumelhart et al. (1986) was applied, setting ? of the hidden layer to 0.5 and ? of the output layer to 0.4. In addition, a small momentum term ? of 0.1 was used. For each simulation, initial weights were randomly selected from a uniform distribution in the range [-0.2, 0.2] for the hidden layer and in the range [-0.7, 0.7] for the output layer.

As a basis for performance evaluation, n-fold cross validated mean square error (MSE) was used This performance measure have been recommended by several authors (e.g., White, 1990). To control model complexity, cross validated mean square errors were computed for models including 0, 2, 4, 6, 8, 10, 12 and 14 hidden units. Thus, the procedure was similar to a constructive algorithm, selecting the best model based upon n-fold cross validated performance measures.

Six benchmark models were developed in order to evaluate the performance of the connectionist model. First, two benchmark models were produced using a recommended procedure in the accounting and finance literature (e.g., Libby, 1975; Zavgren & Friedman, 1988). This procedure recommends using ordinary least squares regression analysis on a selected set of principal components of the original financial cues to avoid multicolinearity problems. Two such models were estimated using nine and five principal components respectively. Next, two benchmark models were produced using the original 32 variables selected by the subjects in an ordinary least squares regression and in a stepwise regression procedure selecting a restricted set of these 32 variables. The two final benchmark results were produced using the 12 input variables of the connectionist model in an ordinary least squares regression and a stepwise regression procedure selecting an even more restricted set of variables. The performance results of the benchmarks were produced using n-fold cross validation. Consequently, the mean square errors are averages of the performance of 75 models for each benchmark. The performance results of the six benchmarks models for the financial diagnosis data are shown in table 2.

Table 2. Performance results for the benchmark models.

Benchmark / Results	MSE	Corr. target	Corr. dist.
9-factor regression	0.232	0.059	0.252
5-factor regression	0.357	0.110	0.382
32-variables regression	0.435	0.179	0.010
32-var. stepw. regr.	0.275	0.147	0.095
12-variables regression	0.221	0.169	0.149
12-var. stepw. regr.	0.212	0.177	0.128

Table 2 shows the cross validated mean square errors and two measures illustrating the distribution of the models' error terms. Both correlations of the error terms with the target output values, and similar correlations of the error terms with distance from mean target value are shown. The first measure illustrates if model errors are larger for good or bad firms, whereas the second measure illustrates if the model is sensitive to extreme target values in general. These models should provide good benchmarks for evaluating the performance of the connectionist model.

5 Analysis and results

Performance results of the connectionist model are shown in table 3. The number of hidden units is indicated in the model name. Cross validated MSE is shown for every 5000 learning iterations.

Table 3. Cross validated MSE for the connectionist model

Iterations:	5000	10000	15000	20000	25000	30000
Model:						
HID0	0.182	0.182	0.185	0.186	0.186	0.187
HID2	0.232	0.185	0.171	0.159	0.158	0.156
HID4	0.175	0.160	0.151	0.147	0.147	0.145
HID6	0.180	0.174	0.163	0.157	0.156	0.159
HID8	0.168	0.173	0.161	0.161	0.161	0.162
HID10	0.172	0.162	0.157	0.155	0.156	0.161
HID12	0.178	0.166	0.157	0.160	0.160	0.160
HID14	0.175	0.169	0.165	0.163	0.159	0.162

As expected, MSE reached a flat minimum during learning and increased with learning beyond this minimum due to overtraining. Even though performance was best for the model with two hidden units, the larger difference in MSE was between the models with hidden units and the model without such units.

Simple t-tests of the differences between means can be used to evaluate of the propositions made earlier in this paper. First, a t-test of the difference in MSE between the best connectionist model stopped at the optimal number of training iterations (HID4) and the best benchmark model showed that the difference was significant and in favor of the connectionist model at ? =0.05 (t=1.96, d.f.=74). This test indicated that the connectionist model was better at modeling the diagnostic outcome produced by our subjects. It was interesting to note that the t-value was relatively low despite the large difference in MSE. This indicated a large and unfortunate variance in the error terms of the best benchmark model.

The distributions of the error terms could be investigated by correlating the error terms with target values and distances from mean target values. For the best connectionist model, the correlation of the error terms with target values was 0.048, and the correlation of the error terms with distance from mean targets was -0.055. These measures indicated that the error terms had a uniform distribution over the range of target values. For the best benchmark model, the corresponding values were 0.177 and 0.128, indicating that errors were somewhat larger for the "healthier" and "extreme" firms. This unfortunate distribution of errors was eliminated in the connectionist model.

Further, a t-test of the difference in MSE between the best connectionist models with and without hidden units showed a significant difference in favor of the model with hidden units at ? =0.05 (t=2.27, d.f.=74). This test supported the proposition that the improved performance could be explained by the ability of the more complex connectionist model to represent and utilize intermediate abstractions of relevance to the diagnosis.

It has been proposed that the intermediate abstractions represented by the hidden units have cognitive relevance. The most obvious presumption in the financial diagnosis area is that the hidden units are local representations of variables that can be interpreted as diagnostic concepts. The most obvious diagnostic concepts are profitability, financing, liquidity and leverage. Principal components analysis of the 12 input variables of the connectionist model showed that the factors with the highest eigenvalues could be interpreted as representing such concepts. Thus, it could be *suggested* that similar dimensions or "factors" were represented by the four hidden units of the best connectionist model.

To investigate this suggestion, analysis of the connectionist model representations was necessary. Experiments indicated that the final weights of the hidden units were very sensitive to initial weights. Consequently, 10 different models were developed using all cases of the data sets for learning, different initial weights, and by stopping the learning at the previously found optimal number of learning iterations. The MSE of these models averaged 0.112, and consequently, some overfit was present due to all cases being used in model development. Except from the somewhat lower MSE, the error terms had distributions very similar to the cross validated error terms, indicating that the 10 models were representative "versions" of the previously cross validated model. When analyzing the weights further, we found that the main reason for the differences in weight patterns were opposite signs of the weight patterns and different absolute weight values due to very different bias weights. After adjusting for opposite signs and differences in bias weights, a cluster analysis was performed using the weights from the input units to the hidden units of all hidden units in the ten models simultaneously.

This analysis revealed that the weight patterns were actually quite similar in the ten models. Some differences were still present due to the very flat error surface, but the cluster analysis was useful in detecting groups of hidden units with similar incoming weight patterns. All models had at least two very local hidden units. The absolute values of these units' weights were large, making them account for most of the models' response variation. Some models had more than two local hidden units, but in those cases, at least two of the local hidden units had almost similar weights.

Thus, it could be concluded that the two functionally different local hidden units were most responsible for model response variation. In most of the models, the rest of the hidden units were more distributed, and only accounted for small variation in model response.

Since the local hidden units were so few, it could be suggested that the hidden units represented other dimensions of the financial cues than the diagnostic concepts profitability, liquidity, financing and leverage. If the hidden units had represented these concepts, four local hidden units with very different weight patterns should have been found. Hinton diagrams of two representative weight patterns from the two local hidden units of one of the models are shown in figure 1.

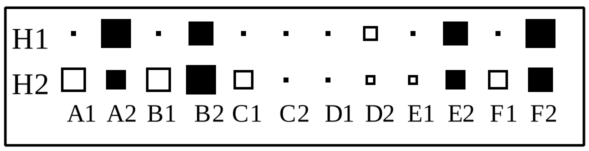

Figure 1. Weight pattern of two local hidden units.

In figure 1, the weights from six financial cues A - F to two local hidden units (H1 and H2) are shown. For each financial cue, two input units represent values of the first and second of two consecutive years. The financial cues are from left to right; operating margin (A), net income to total assets (B), net income to equity (C), average interest rate (D), acid test (E) and equity to total assets (F).

This weight pattern indicated very little specialization of each hidden unit on financial cues representing different diagnostic concepts. Instead, it seemed that the upper hidden unit (H1) used the second year values of a broad set of financial cues representing profitability, financing, liquidity and leverage. Thus, this hidden unit detected the level of the financial cues of the latest year. Thus, it was termed "level-oriented".

The lower hidden unit (H2) used both values of a broad set of financial cues, mainly representing profitability, liquidity and leverage. However, the weight pattern was interesting. It had negative weights from cues of the first year and positive weights from cues of the second year. Consequently, this unit was activated by the change in the level of a broad set of financial cues. The hidden unit actually computed the general trend in the financial cues from one year to the next. Consequently, this hidden unit was termed "change-oriented".

Further analysis of the outputs of these hidden units showed that the "change-oriented" hidden units had outputs that were approximately linear functions of the change in financial cues. The "level-oriented" hidden units had outputs that were highly non-linear functions of the level of the latest financial cues. This non-linearity gave under-proportional outputs for extremely high cue values.

To form the final diagnosis of the financial situation, the outputs of the two local hidden units were combined in an approximately linear function. Thus, it could be concluded that the diagnosis of the financial situation was performed by combining the response of two local hidden units representing the level of and the change in a broad set of financial cues.

6 Conclusion and discussion

To evaluate if the superior performance of neural networks for financial diagnosis could be explained by these models' ability to develop relevant internal representations, a behavioral approach was applied here. By setting up a behavioral experiment, a data set of valid financial diagnoses performed by skilled subjects under controlled conditions was provided. Based upon indications of the skilled subjects and sensitivity analysis, a set of important financial cues was selected as inputs to a multilayered perceptron model. Since the model was based upon behavioral data, we termed it a connectionist model of financial diagnosis. To evaluate the performance of the model, six different linear benchmarks were used.

Three findings were made. First, the performance of the connectionist model was significantly better than the benchmark models' performance when evaluated by the recommended n-fold cross validated mean square error. By comparing the error terms of the connectionist model and the best benchmark model, it was found that the error term distribution of the connectionist model was favorable. Consequently, we concluded that the superior performance of neural networks proposed for predictive models was also valid for models based upon behavioral financial diagnosis data.

Second, it was found that the connectionist model with hidden units significantly outperformed the model without hidden units developed with comparable parameter settings. This led us to conclude that the main reason for the superior performance of the connectionist model was the hidden units and not the non-linearity of the output function, nor the method of parameter estimation (gradient descent). This finding was in correspondence with previously untested propositions made by several researchers applying neural networks to financial diagnosis.

Third, simple analysis of the internal representations of the connectionist models revealed no direct correspondence between the intermediate abstractions represented in the connectionist models and diagnostic concepts, such as profitability, financing, liquidity or leverage. The hidden units of the connectionist models seemed to represent the level and change in a broad set of financial cues. This was rather surprising, since principal components analysis of the same financial cues had revealed factors corresponding well to traditional diagnostic concepts. The underlying dimensions of level and change in the financial cues used by the connectionist models could be explained in two ways. First, the connectionist model may have implemented a functional heuristic not used by the skilled diagnosticians, but which still produced a behavior similar to that of the diagnosticians. Second, the diagnosticians may actually have used this functional heuristic, performing the diagnosis by an evaluation of the general level of a broad set of financial cues and an evaluation of the change in a similar broad set of cues. Some experiments were performed constraining the weight pattern of the connectionist model so that only representations of diagnostic concepts could be developed by the hidden units. In this case, performance deteriorated significantly. This indicated that the internal representations of the connectionist model had cognitive relevance, even though the representations were not as initially expected. A likely explanation of this finding was that the diagnosticians actually implemented a much simpler heuristic during diagnosis than prescriptive theory of financial diagnosis recommended.

There are several threats to the validity of these conclusions, such as lacking validity of the firms used as cases for diagnosis, or lack of skill among the diagnosticians. However, the firms were randomly selected, the diagnosticians had experience in financial diagnosis, and first of all, the diagnostic data were provided using a controlled experiment and randomization procedures not traditionally applied to tests of connectionist model performance. Several methods were used in evaluating model performance, all indicating superiority in favor of the connectionist model.

Our suggestion is that validation of neural network performance may not only be improved using multiple data sets and benchmarks, but as in any other empirical study, by using and reporting experimental procedures to demonstrate the validity of the data. In addition, evaluation of connectionist model performance can be improved by combining analyses of model performance. Even though cross validation improves the evaluation, analysis of the distributions of the models' error terms should be performed in a way similar to outlier analysis in linear models. Finally, analysis of the internal representations should be applied to compare the representations of the model to intermediate abstractions believed to be relevant in the application domain of the model. If the intermediate abstractions differ from these, their validity should be discussed and evaluated. Combining traditional performance measures, analysis of error term distributions and evaluation of internal representations constitutes a three-dimensional validation of a connectionist model.

References

Abdel-khalik, A.R., & El-Sheshai, K.M. (1980). Information choice and utilization in an experiment on default prediction. *Journal of Accounting Research*, autumn 1980 , 325-342.

Altman, E.I. (1968). Financial ratios, discriminant analysis and the prediction of corporate bankruptcy. *The Journal of Finance*, September 1968 , 589-609.

Altman, E.I., Avery, R.B., Eisenbeis, R.A. & Sinkey, J.F. (1981). *Application of classification techniques in business, banking and finance*. Greenwich, CT: The JAI Press.
Berry, R., & Trigueiros, D. (1993). Applying neural networks to the extraction of knowledge from accounting reports: A classification study. In Trippi, R.R. & Turban, E. (eds.), *Neural networks in finance and investing* (p. 103-123). Chicago, IL: Probus Publishing.
Biggs, S.F., Selfridge, M., & Krupka, G.R. (1993). A computational model of auditor knowledge and reasoning processes in the going-concern judgement. *Auditing: A Journal of Practice & Theory,* 12(supplement), 82-99.
Bouwman, M.J., Frishkoff, P.A., & Frishkoff, P. (1987). How do financial analysts make decisions? A process model of the investment screening decision. *Accounting, Organizations and Society*, 12, 1-29.
Chalos, P. (1985). Financial distress: A comparative study of individual, model and committee assessments. *Journal of Accounting Research*, 23, 527-543.
Frydman, H., Altman, E.I., & Kao, D. (1985). Introducing recursive partitioning for financial classification: The case of financial distress. *The Journal of Finance*, 40, 269-291.
Gombola, M.J., & Ketz, J.E. (1983). Financial ratio patterns in retail and manufacturing organizations. *Financial Management*, 12(2), 45-56.
Gorman, R.P., & Sejnowski, T.J. (1988). Analysis of hidden units in a layered network trained to classify sonar targets. *Neural Networks*, 1, 75-89.
Hanson, S.J., & Burr, D.J. (1990). What connectionist models learn: Learning and representation in connectionist networks. *Behavioral and Brain Sciences*, 13, 471-518.
Hopwood, W., McKeown, J.C., & Mutchler, J.F. (1994). A reexamination of auditor versus model accuracy within the context of the going-concern opinion decision. *Contemporary Accounting Research*, 10, 409-431.
Jones, F.L. (1987). Current techniques in bankruptcy prediction. *Journal of Accounting Literature*, 6, 131-164.
Kaplan, R.S. & Urwitz, G. (1979). Statistical models of bond ratings: A methodological inquiry, *Journal of Business*, 52, 231-261.
Kida, T. (1980). An investigation into auditors' continuity and related qualification judgements, *Journal of Accounting Research*, 18, 506-523.
Libby, R. (1975). Accounting ratios and the prediction of failure: Some behavioral evidence. *Journal of Accounting Research*, spring 1975, 150-161.
Libby R. (1981). *Accounting and human information processing: Theory and applications*. Englewood Cliffs, NJ: Prentice Hall.
Methlie, L.B. (1987). On knowledge based decision support systems for financial diagnosis. In Holsapple, C.W., & Whinston, A.B. (Eds.), *Decision support systems: Theory and application*. Berlin: Springer Verlag.
Moody, J.E., & Utans, J. (1995). Architecture selection strategies for neural networks: Application to corporate bond rating prediction. In Refenes, A.P. (Ed.), *Neural networks in the capital markets* (pp. 277-300). New York: Wiley.

Mutchler, J.F. (1985). A multivariate analysis of auditor's going-concern opinion decision. *Journal of Accounting Research*, 23, 668-682.
Newell, A., & Simon, H.A. (1972) *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
Ohlson, J.A. (1980). Financial ratios and the probabilistic prediction of bankruptcy. *Journal of Accounting Research*, 18, 109-131.
Pinches, G.E., Mingo, K.A., & Caruthers, J.K. (1973). The stability of financial patterns in industrial organizations. *Journal of Finance*, 28, 389-396.
Raghupathi, W., Schkade, L.L., & Raju, B.S. (1991). A neural network approach to bankruptcy prediction. Reprinted in Trippi, R.R., & Turban, E. (Eds.), *Neural networks in finance and investing* (pp. 141-158). Chicago, IL: Probus Publishing.
Rodgers, W. (1991). How do loan officers make their decisions about credit risks? A study of parallel distributed processing. *Journal of Economic Psychology*, 12, 243-265.
Rumelhart, D.E., Hinton, G.E., & Williams, R.J. (1986). Learning internal representations by error propagation. In Rumelhart, D.E., & McClelland, J.L. (Eds.), *Parallel distributed processing; Explorations in the microstructure of cognition*. Cambridge, MA: MIT Press.
Simnett, R., & Trotman, K.T. (1989). Auditor versus model: Information choice and information processing. *The Accounting Review*, 64, 514-528.
Singleton, J.C., & Surkan, A.J. (1995). Bond rating with neural networks. In Refenes, A.P. (Ed.), *Neural networks in the capital markets* (pp. 301-307). New York: Wiley.
Smith, E.E., & Medin, D. (1981). *Categories and concepts*. Cambridge, MA: Harvard University Press.
Srinvasan, V., & Kim, Y.H. (1987). Credit granting: A comparative analysis of classification procedures. *The Journal of Finance*, 42, 665-683.
Srivastava, R.P. (1992). Automating judgmental decisions using neural networks: A model for processing business loan applications. *Proceedings of the 20the Annual ACM Computer Science Conference* (pp. 351-357), New York: ACM
Surkan, A.J., & Singleton, J.C. (1990). Neural networks for bond rating improved by multiple hidden layers. In Trippi, R.R., & Turban, E. (Eds.), *Neural networks in finance and investing*. Chicago, IL.: Probus Publishing.
Tam, K.Y. (1991). Neural network models and the prediction of bank bankruptcy. *OMEGA International Journal of Management Science*, 19, 429-445.
Tam, K.Y., & Kiang, M.Y. (1992). Managerial applications of neural networks: The case of bank failure predictions. *Management Science*, 38, 926-947.
White, H. (1990). Connectionist nonparametric regression: Multilayer feedforward networks can learn arbitrary mappings. *Neural Networks*, 3, 535-549.
Wilson, R., & Sharda, R. (1994). Bankruptcy prediction using neural networks. *Decision Support Systems*, 11, 545-557.
Yoon, Y., Swales, G., & Margavio, T.M. (1993). A comparison of discriminant analysis versus artificial neural networks. *Journal of the Operations Research Society*, 44, 51-60.

Zavgren, C.V., & Friedman, G.E. (1988). Are bankruptcy prediction models worthwhile? An application in securities analysis. *Management International Review,* 28, 34-44.
Ziebart, D.A., & Reiter, S.A. (1992). Bond rating, bond yields and financial information. *Contemporary Accounting Research,* 9, 252-282.