CONNECTIONIST MODELS OF FINANCIAL DIAGNOSIS

by

PER E. PEDERSEN

Dissertation presented to the Norwegian School of Economics and Business Administration in partial fulfilment of the requirements for the degree Doctor Oeconomiae.

October, 1995

Acknowledgments

While undertaking the present research I have greatly benefited from the encouragement and valuable insights of a number of persons. Special thanks are due to the members of my dissertation committee: Professor Kjell Grønhaug and professor Leif B. Methlie at the Norwegian School of Economics and Business Administration, and professor Svein Nordbotten at the University of Bergen.

I owe thanks to professor Knut Boye and professor Gunnar A. Dahl at the Norwegian School of Economics and Business Administration, who provided me with access to the students participating in the present investigation. I also wish to thank professor Charles B. Stabell at the Norwegian School of Management for interesting discussions during an early phase of this project, professor Andrew J. I. Jones at the University of Oslo for careful evaluation of my masters thesis, colleagues at the Norwegian School of Economics and Business Administration for a stimulating environment during my stays in Bergen, and the staff at the library of Agder College in Grimstad for assistance in providing much of the background material for this thesis.

This research has benefited from financial support by the Norwegian School of Management, The Norwegian School of Economics and Business Administration, and Agder College.

Narestø
October, 1995 Per E. Pedersen

Contents

Acknowledgments i
List of tables vi
List of figures ix
Nomenclature xi
Abstract xiv

PART I - INTRODUCTION 1

Chapter 1. Introduction 3
1.1 Perspective 3
1.2 Problem 6
1.3 Purpose 8
1.4 Organisation of the dissertation 9

PART II - THEORY 13

Chapter 2. Financial diagnosis theory 15
2.1 The judgement modelling approach 20
2.2 The cognitive approach 31
2.3 The predictive approach 42
2.4 Other contributions of relevance to financial diagnosis 51
2.5 Summary of findings on the financial diagnosis task 53

Chapter 3. Cognitive theory 63
3.1 Classification theory 69
3.1.1 Definitional theory 71
3.1.2 Prototype theory 75
3.1.3 Exemplar theory 81
3.2 Connectionist theory 86
3.2.1 General connectionist theory 89
3.2.1.1 Functional components 89
3.2.1.2 Operating principles 92
3.2.1.3 General connectionist and backpropagation models 94
3.2.2 Connectionism and artificial neural networks 100
3.2.3 Connectionist modelling and environmental constraints 105
3.2.4 Connectionist models of categorisation and classification 110

Chapter 4. Connectionist models of financial diagnosis 119
4.1 Neural networks in financial diagnosis 119
4.2 A connectionist model of financial diagnosis 129

PART III - METHOD 137

Chapter 5. Research design 139
5.1 The stimulus material 140
5.2 Subjects 144
5.3 Treatments and procedures 145
5.4 Measurement and properties of measures. 147

Chapter 6. Simulation design 157
6.1 Methodology of connectionist simulations 157
6.2 Simulation environment parameters and design tools 158
6.3 Generalisation measures 162
6.4 Benchmarks 165

PART IV - SIMULATIONS AND RESULTS 169

Chapter 7. A connectionist model of classificatory response 171
7.1 Performance results 172
7.2 Analysis of model representations 174
7.3. Conclusions 179

Chapter 8. A connectionist model of continuous response 181
8.1 Performance results 181
8.2 Analysis of model representations 191
8.2.1 Common hidden unit analysis 197
8.2.2 Analysis of non-common hidden units 199
8.3 Conclusions 210

Chapter 9. A constrained connectionist model of continuous response 217
9.1 Performance results 219
9.2 Analysis of model representations 230
9.3 Constraining complexity by cue importance 247
9.4 Conclusions 249

PART V - DISCUSSION 257

Chapter 10. Discussion 259
10.1 Summary and conclusions 259
10.2 Discussion 263
10.3 Implications 269
10.3.1 Theoretical and methodological implications 270
10.3.2 Practical implications 271
10.4 Suggestions for further research 273

References 277

Appendix A. Summary of applications of neural networks to problems in business administration. 305
Appendix B. Introductory text, stimulus material, and response form. 308
Appendix C. Companies with complete financial statements and their code. 311
Appendix D. Ratio formulas 313
Appendix E. Summary statistics of independent variables 314
Appendix F. Correlation matrix of independent variables 315
Appendix G. Factor loadings of first benchmark (A) 319
Appendix H. Factor loadings of second benchmark (B) 320
Appendix I. Summary statistics of dependent variables. 321
Appendix J. Manipulation check for variables level and trend 322
Appendix K. Typical relationships between learning and cross validated errors 323

List of tables

Table 2.1 Task context and perspective combinations of financial diagnosis 18
Table 2.2 Research questions in different approaches to the study of financial diagnosis 19
Table 2.3 Selected judgement modelling studies of financial diagnosis 21
Table 2.4 Selected experimental cognitive studies of financial diagnosis 32
Table 2.5 Selected descriptive cognitive studies of financial diagnosis 38
Table 2.6 Selected predictive studies of financial diagnosis 45
Table 2.7 Concepts operationalised by independent variables in selected studies of financial diagnosis. 55
Table 3.1. Selected early connectionist models 95
Table 4.1 Applications of connectionist and artificial neural network models to financial diagnosis. 121
Table 5.1 Summary statistics of firms in the stimulus sample 143
Table 5.2 Summary statistics of sample subjects 145
Table 5.3. The ten most frequently indicated cues used on different diagnostic areas 149
Table 5.4 Measurement statistics of the level measure of the financial situation 153
Table 5.5 Measurement statistics of the trend measure of the financial situation 153
Table 5.6 Most frequently used linguistic terms ordered by average level value 154
Table 6.1 Performance results of the first (9-factor) benchmark (A) 166
Table 6.2 Performance results for second (5-factor) benchmark (B) 167
Table 7.1. Mean squared error (MSE) for diagnosis of bankruptcy 172
Table 7.2. Correct classifications of bankruptcy and non-bankruptcy 173
Table 7.3. The effects of jogging input values 5 % in the bankruptcy classification model 176
Table 8.1. Mean squared error (MSE) of the level diagnosis 182
Table 8.2. T-tests of best connectionist model vs. best stepwise regression model for the level diagnosis 183
Table 8.3. T-tests of best connectionist model vs. best benchmark model (A) of chapter 6 for the level diagnosis 183
Table 8.4. T-tests of connectionist model with 2 hidden units vs. connectionist model without hidden units for the level diagnosis 183
Table 8.5. Correlations of SE and target for the level diagnosis 184
Table 8.6. Correlations of SE and distance from mean target for the level,diagnosis.. 184
Table 8.7. Mean squared error (MSE) of the trend diagnosis 186
Table 8.8. T-tests of best connectionist model vs. stepwise regression for trend diagnosis 186
Table 8.9. Correlations of SE and target for the trend diagnosis 187

Table 8.10. Correlations of SE and distance from mean target for the trend diagnosis 188
Table 8.11. Mean squared error (MSE) of level diagnosis in a combined model 189
Table 8.12. Mean squared error (MSE) of trend diagnosis in a combined model 189
Table 8.13. Correlations of SE and distance from mean target for the trend diagnosis in a combined model 190
Table 8.14. Results of 10 combined models with full learning sample and randomised initial weights 191
Table 8.15. Statistics illustrating the differences in representations between the 10 versions of each combined model 194
Table 8.16. Mean weight values, standard deviation and t-value of test of $\mu=0$ for the weights between input units and the common hidden units 198
Table 8.17. Correlations of the hidden unit outputs with targets and difference between targets 201
Table 8.18. Mean weight values, standard deviation and t-value of test of $\mu=0$ for the weights between input units and 10 trend-oriented hidden units 205
Table 8.19. Mean weight values, standard deviation and t-value of test of $\mu=0$ for the weights between input units and 13 level-oriented hidden units 206
Table 8.20. Mean weight values, standard deviation and t-value of test of $\mu=0$ for the weights between input units and 24 "difference" units 208
Table 8.21. Mean weight values, standard deviation and t-value of test of $\mu=0$ for the weights between input units and 9 units in the last cluster 209
Table 8.22. Mean squared error (MSE) of the level and trend diagnoses in a restricted model 212
Table 8.23. The restricted model's correlations of hidden unit outputs with composite judge diagnoses of the four diagnostic areas "profitability", "financing", "liquidity" and "leverage" 213
Table 9.1. The effects of jogging input values 5 % in the combined model 218
Table 9.2. Number of weights in the different combined models 218
Table 9.3. Mean squared error (MSE) of the level diagnosis in a constrained model.. 219
Table 9.4. T-tests of the best level model of chapter 8 (HID2) vs. best constrained connectionist model 220
Table 9.5. Correlations of SE and target for the level diagnosis in a constrained model 221
Table 9.6. Correlations of SE and distance from mean target for the level diagnosis in a constrained model 221
Table 9.7. Mean squared error (MSE) of the trend diagnosis in a constrained model . 222
Table 9.8. T-tests of the best trend model of chapter 8 (HID6) vs. best constrained connectionist model 223

Table 9.9. Correlations of SE and target for the trend diagnosis in a constrained model 223
Table 9.10. Correlations of SE and distance from mean target for the trend diagnosis in a constrained model 223
Table 9.11. Mean squared error (MSE) of the level diagnosis in a constrained and combined model 224
Table 9.12. T-tests of the best combined model of chapter 8 (HID2) vs. the best constrained and combined connectionist model (level diagnosis) 225
Table 9.13. Mean squared error (MSE) of the trend diagnosis in a constrained and combined model 225
Table 9.14. T-tests of the best combined model of chapter 8 (HID10) vs. the best constrained and combined connectionist model (trend diagnosis) 226
Table 9.15. Correlations of SE and distance from mean target for the trend diagnosis in a constrained and combined model 227
Table 9.16 Cases classified as outliers in models with 32 and 12 inputs 228
Table 9.17 Analysis of model outliers 229
Table 9.18. Results of 10 constrained and combined models with full learning sample and randomised initial weights 230
Table 9.19. Statistics illustrating the differences in representations 231
Table 9.20 Factor loadings from principal components analysis of the input variables 244
Table 9.22 Five most important cues of each diagnostic area 248
Table 9.23. Mean squared error (MSE) of the level diagnosis in a combined subject selected model 249
Table 9.24. Mean squared error (MSE) of the trend diagnosis in a combined subject selected model 249
Table 9.25. Mean squared error (MSE) of the level and trend diagnoses in a constrained and restricted model 252
Table 9.26. The constrained and restricted model's correlations of hidden unit outputs with composite judge diagnosis of the four diagnostic areas "profitability", "financing", "liquidity" and "leverage" 253

List of figures

Figure 2.1 Lens model of Brunswik (1952) 15
Figure 3.1 Decision bounds of definitional theory 73
Figure 3.2. Standard connectionist model 78
Figure 3.3. Decision bounds of prototype theory 79
Figure 3.4. Decision bounds of exemplar theory 84
Figure 3.5. The effects on decision bound of changing exemplar distribution in a category 84
Figure 3.6 The ALCOVE model 85
Figure 3.7 Backpropagation network with one hidden layer 96
Figure 3.8. Single layered perceptron 100
Figure 3.9 A distributed (coarse coded) representation of two stimulus dimensions .. 106
Figure 3.10 Operationalisation by model 109
Figure 3.11 The isoactivation contours of ALCOVE 115
Figure 3.12 Isoactivation contours of backpropagation 117
Figure 4.1. A connectionist model of financial diagnosis with example stimulus dimensions, intermediate abstractions and response classes 131
Figure 4.2. Examples of alternative intermediate abstractions with varying degree of locality 134
Figure 7.1 Hinton diagram of connections between the hidden and output layer of 10 bankruptcy classification models with 2 hidden units 177
Figure 7.2. Hinton diagram of the connections between the input layer and hidden layer of 10 bankruptcy classification models with 2 hidden units 178
Figure 8.1. Hinton diagram of the weights between hidden layer and output layer of 10 versions of the combined models 193
Figure 8.2. Cluster analysis dendrogram of input to hidden weights in all the 30 versions of the combined connectionist model 196
Figure 8.3. Cluster analysis dendrogram of unscaled response of the output units 200
Figure 8.4. Hidden unit outputs of four hidden units as a function of level and trend. 202
Figure 8.5. Hinton diagram of all weights in the three analysed combined models 204
Figure 8.6. Hinton diagram of weights in a restricted connectionist model 213
Figure 9.1. Cluster analysis dendrogram of input to hidden weights in all the 30 versions of the constrained and combined connectionist model 232
Figure 9.2 Hinton diagram of the weights between hidden and output units for the 10 versions of the constrained and combined models 234
Figure 9.3 Cluster analysis dendrogram and Hinton diagram of the constrained and combined models with two hidden units 235

Figure 9.4 Cluster analysis dendrogram and Hinton diagram of the constrained and combined models with three hidden units 237
Figure 9.5 Cluster analysis dendrogram and Hinton diagram of the constrained and combined models with four hidden units 239
Figure 9.6 Plot of hidden unit outputs (y-axis) vs. linearly predicted values 241
Figure 9.7 Linearities and nonlinearities in model C3 243
Figure 9.8 In vivo clustering of hidden unit outputs in the C3 model 246
Figure 9.9 Hinton diagram of weights in a constrained and restricted connectionist model 252

Nomenclature

Cue abbreviations

ACCPAY	accounts payable
ACID	acid test
AIR	average interest rate
APT	accounts payable period
ART	accounts receivable period (collection period)
ASSTURN	assets turnover
BER	equity ratio
CAPASS	capital assets
CASH	cash
CHLIK	change in liquidity
CHOSTREC	change in other short term receivables
CHWORK	change in working capital
CGROWTH	change in costs
CONTPR	contribution margin
CURR	current ratio
CURRASS	current assets
EQUITY	equity
FREERES	accumulated retained earnings
ICOV	interest coverage
ITURN	inventory turnover time
LTINV	long term inventory financing
LTL	long term liabilities
OOPREV	other operating revenues
OPMARG	operating margin
OPROF	operating profits
OSTREC	other short term receivables
PLBEI	profit and loss before extraordinary items
PROMARG	profit margin
RES	deferred credits
ROE	return on equity
ROI	return on assets
SGROWTH	change in sales
STL	short term liabilities
TOTCAP	total assets/capital

General abbreviations

ART	adaptive resonance theory
ALCOVE	attentional learning coverage model
BSB	brain-state in box model
d.f.	degrees of freedom
FIFO	first in first out
HIDi	backpropagation model with i hidden units
ID3	the ID3 recursive partitioning algorithm
KNN	k-nearest neighbour algorithm
LIFO	last in first out
LVQ	learning vector quantisation artificial neural network
MDS	multidimensional scaling
MSE	mean squared error
OLS	ordinary least squares
PDP	parallel distributed processing
RBF	radial basis function artificial neural network
SE	squared error

Symbols

C_j	class j in a classification task
d_{pa}	distance in psychological space between pattern p and exemplar
or	
	prototype a.
E_p	error of a network output when pattern p is presented
F, f and g	functions
h_{ji}	centre of the receptive field of unit j on feature or stimulus
	dimension i in a network
I_{pi}	value of a cue of pattern p on feature or stimulus dimension i
M	number of classes in a classification task
n	number of features or stimulus dimensions in a classification
task,	
	and number of cues in a lens model
net_{pj}	signal coming into unit j when pattern p is presented to a
network	
N	number of patterns in a classification task,
	and number of observations in a sample/statistical test
o_{pi}	output of unit i when pattern p is presented to a network
$P(C_j \| p)$	conditional probability of pattern p being classified in class j
r_{ei}	cue validity coefficients in a lens model

r_{si}	cue utilisation coefficients in a lens model
$s(p,a)$	similarity of pattern or item p to the item, exemplar or prototype
a	
s_i	similarity between exemplars or items on a particular feature or stimulus dimension i
t	the test observator t
t_{pj}	target of unit j when pattern p is presented to a network
w_{ji}	weight of a connection from unit i to unit j in a network
X_i	value of cue i in a lens model
Y_e	criterion variable in a lens model
Y_s	judgement variable in a lens model
α	significance level in a statistical test, and Chronbach's α and momentum term in a backpropagation model, and attention parameter in ALCOVE
δ_{pj}	error gradient of unit j when pattern p is presented to a backpropagation model
Δw_{ji}	a small change in w_{ji} of a network
η	learning rate parameter in a backpropagation model
η_h	learning rate parameter in the hidden layer of a backpropagation model
η_o	learning rate parameter in the output layer of a backpropagation model
σ_{SE}	standard deviation of squared errors

Abstract

Financial diagnosis is when a subject makes a judgement of the financial situation of the firm based upon information from the financial statement. This task is performed in several contexts, such as bankruptcy prediction, going concern judgement and loan decision contexts.

Three approaches to financial diagnosis are found in the literature; a judgement modelling, a cognitive, and a predictive approach. A review of these approaches constitutes a task analysis of financial diagnosis. A somewhat surprising finding in the review is that even though several of the approaches apply a classification conception of the financial diagnosis task, cognitive classification theory has not been used to any extent to explain subjects' diagnostic behaviour while performing financial diagnosis. This is very different from other diagnostic tasks, which have been extensively studied from a cognitive classification perspective. From this finding, we conclude that cognitive classification theory can be used to increase our understanding of the financial diagnosis task in general and of the role of less investigated concepts in cognitive accounting, such as pattern recognition, pattern matching and prototypes.

To provide the basis for applying a cognitive classification perspective on the financial diagnosis task, a presentation of definitional, prototype and exemplar theory of classification is given. Both prototype and exemplar theories have recently been implemented in connectionist models, and these models are considered among "the leading candidates" in contemporary classification theory. Based upon the task analysis of the financial diagnosis task, a connectionist classification model is selected and applied to the task. The backpropagation model of Rumelhart, Hinton and Williams (1986) is considered to have the ability to develop internal representations functional in performing complex classification tasks, such as financial diagnosis. From the model, three propositions are made. The first proposition, P1, states that connectionist models of financial diagnosis should show better fit than benchmarks of linear models. The second proposition, P2, states that the improved fit could primarily be explained by the ability of the connectionist models to build internal representations, and the third proposition, P3, states that these internal representations should have cognitive relevance.

To evaluate these propositions, a financial diagnosis experiment is reported. 108 subjects participated in the diagnosis of 75 randomly selected small and medium sized firms. Full financial statements and selected ratios of two consecutive years were used as stimulus material, and several measures of diagnostic response were collected. The treatment plan resulted in 324 diagnoses of the 75 firms, averaging 4.32 diagnoses per firm. To create the stimulus-response pairs representing learning and test samples of the connectionist model, both simple and composite judge measures of the subjects' diagnoses are designed.

A simulation design is developed that accommodates resampling methods and cross validated measures to evaluate the performance of the connectionist model. Furthermore, several benchmarks are developed using traditional methods of the judgement modelling approach to financial diagnosis.

The propositions are evaluated using three simulations with varying stimulus and response representations. The first simulation uses a stimulus representation consisting of 17 selected financial ratios. Diagnostic response is measured by a bankruptcy classification variable. The second simulation uses the same stimulus representation as the first, but diagnostic responses are measured by composite judge assessments of level and trend diagnoses of the financial situation. In the third simulation, the diagnostic response representations of the second simulation are used, but sensitivity based measures are used to constrain the stimuli to six financial ratios. Generally, model fit is improved from simulation one through simulation three.

Strong support is found for proposition P1. The connectionist models show significantly better fit than traditional benchmarks when evaluated by cross validated average squared error. In particular, the model with constrained stimulus representations and composite judge diagnoses shows favourable performance. For the connectionist models showing significantly better fit than the benchmarks, tests are made to evaluate proposition P2. In these tests, significantly better fit is found for the connectionist models with hidden units than for the models without hidden units. Because all simulation parameters are similar in the two model types, it can be concluded that the difference in performance is explained by the internal representations of the hidden units. This finding support proposition P2. Evaluation of proposition P3 is done by representational analysis. The representations built by the hidden units of the connectionist models were expected to consist of derived stimulus dimensions reflecting different diagnostic areas, such as "profitability", "financing" or "liquidity". However, completely different and much more complex representations are built by the hidden units. A direct interpretation of these units as representing concepts, variables or prototypes is difficult. However, interpretation of the molar behaviour of the connectionist models is possible using both rule-based and prototype-based terms. A rule-plus-exception interpretation is given of some of the models, while other models are best described as computing similarity to prototypical firms, such as the "bad, but promising" or the "good, but with alarming trend" firms.

Two implications from these results are particularly interesting. First, connectionism offers a way to develop cognitive models of financial diagnosis that show good fit to behavioural data. Second, connectionist models offer a way to unify judgement modelling and cognitive approaches to financial diagnosis because cognitive models are developed with methods similar to those traditionally applied in judgement modelling studies.

PART I - INTRODUCTION

Chapter 1. Introduction

Understanding how analysts characterise and classify firms based upon information from the financial statement is relevant in several task contexts, and is the main subject of this thesis. Financial analysts may characterise a firm as a risky investment. Loan officers may decide to reject a loan. An auditor may be reluctant to characterise the firm as a going concern, and a rating agency may change the rating of the firm, all based upon the same information.

The term "financial diagnosis" was first introduced by Methlie (1987), as a general term to describe these characterisations across task contexts. The term "diagnosis" gives strong connotations to a medical diagnosis, and refers to the financial diagnosis as something more than an ad hoc characterisation of the financial situation of the firm. These connotations can lead to the assumption that firms may have "diseases" threatening their existence, and that the identification of such potential "diseases" is possible by investigating the "manifestations" or "symptoms" identifiable in the financial statement.

The outcome of a financial diagnosis has economic consequences for both the firm and the analyst. A specific characterisation may lead to changes in financial costs, stock prices or even future contracts for the firm, and may result in changes in income, costs and the reputation of the analyst. Consequently, understanding the financial diagnosis task and the way this is performed is of great relevance to management, accounting and finance.

1.1 Perspective

There are several ways in which a task can be conceived (e.g. Mintzberg, Raisinghani & Theoret, 1976; Simon, 1979). In addition to response time, the dimensionality and properties of stimuli and responses are used to identify different types of tasks (Rouse, Hammer & Lewis, 1989). Different aspects are focused depending on how the task is conceived. In a recent introduction to cognitive science (Osherson & Smith, 1990), a taxonomy of thought processes of progressively more complex forms is presented, with problem solving as the most complex form. This taxonomy can be used to illustrate how different aspects of the financial diagnosis task are focused with different task conceptions.

Considered as a *problem solving task*, the time duration from stimulus presentation to final response is considerable, and it is assumed that several subtasks are performed by subgoaling and intermediate solutions. Consequently, identification of the subgoals may be focused.

A particular form of problem solving task conception is that the financial diagnosis takes the form of a *hypothesis testing process*, in which the diagnostician formulates a hypothesis early in the process and gathers information to test this hypothesis. With this task conception, identification of the hypothesis and the efficiency of the hypothesis testing strategy may be focused.

A task conception assuming a somewhat less complex form of thought is that the financial diagnosis takes the form of a *choice*. In a choice task, the final decision depends upon a preference function relating the decision to the perceived utility of the consequences of different choice alternatives. With this task conception, estimation of preferences and ordering of choices may be focused.

Considered as a *prediction task*, the diagnosis centres around the trends of financial items, and the possible consequences a prolonged trend of the same form may have in the future. Deviations of human diagnoses from forecasts of formal models may be focused with this task conception.

Considered as a *classification task,* we may conceive financial diagnosis as mapping the N-dimensional space of financial items onto meaningful classes of firms. The classes are clusters of firms which have a diagnosis in common. This conception is closely related to a categorisation of the presented stimulus. With this conception, the relevant classes and the representations necessary to structure the classification may be focused.

Considered as a *pattern recognition task*, the stimulus is perceived as a pattern similar to a previously perceived pattern with an identified diagnosis . The present stimulus is given the same diagnosis. With this approach, time duration from stimulus to response is short, and focus may be on visual features relevant to recognition. Consequently, very little of the task is open to cognitive investigation, and the task is treated as a perceptual task more than a cognitive task.

From a cognitive perspective, financial diagnosis has often been treated as a problem solving task. Medical diagnosis has been extensively investigated as a problem solving task (see Elstein, Shulman & Sprafka, 1990), but approaches treating diagnosis as classification are also found (Brooks, Norman & Allen, 1991). As an example of this conception, consider the following definition by Kirkebøen (1993) :

"He (the diagnostician) compares the patient's pattern of the symptoms with the patterns usually associated with a given disease. For any disease there is a class of patterns of symptoms. These classes are characterised as diagnoses. The determination of a diagnosis for a particular patient is the fit of the observed pattern of symptoms with the general pattern of symptoms for the disease. This way diagnosis is equivalent to performing a classification.......What the clinician does when he gives the patient a diagnosis, is to place the patient within a category of diseases" (Kirkebøen, 1993, p. 167, translated from Norwegian)

A similar view can be found in Chandrasekaran and Goel (1988):

"Medical problem solving thus may be organized first as classifying patients' symptoms onto disease categories, i.e., diagnosis as classification, and then indexing the therapeutic actions by the disease categories." (Chandrasekaran & Goel, 1988, p. 417)

This definition gave rise to the term "classificatory diagnosis" (Chandrasekaran and Goel, 1988) as a conception of diagnosis as classification. A similar conception of *financial* diagnosis has been put forward by Methlie (1994):

Financial analysis is a form of diagnostic problem solving. To diagnose is the act or process leading to detection of a fault or defect of the studied object (in medical terminology: a disease) on the basis of observed symptoms. This process is clinical in nature, which means that each case must be treated as unique. Diagnostic knowledge is organised around classes of phenomena. When we have decided what class the phenomena belongs to, we can treat the problems by using knowledge of the class' attributes. The way class membership is found, is central to diagnostic problem solving. (Methlie 1994, p. 336 (translated from Norwegian))

The conception of diagnosis as classification does not necessarily imply that classification takes the form of a *simple* mapping of stimulus to response. When the classification is complex or when the number of categories is large, it must be assumed that the classification is performed using intermediate information processing steps and intermediate abstractions, possibly by subclassifications between stimulus presentation and response (Chandrasekaran & Goel, 1988). In addition, detecting relevant parts of the stimulus may be part of the classification task, possibly performed by specialised feature detectors. We assume that identification of relevant features in the stimulus and the use of intermediate abstractions are important parts of the diagnostician's knowledge of financial diagnosis.

Identification of the intermediate information processing steps and intermediate abstractions in diagnosis receive different attention depending upon how the diagnosis task is conceived. From a traditional information processing perspective, the task is conceived as a problem solving task, and the intermediate information processing *steps* are often focused. Consequently, what is *done* by the diagnostician is revealed with this perspective. From a classification perspective, the intermediate abstractions necessary to perform the classification will be focused. Consequently, the *knowledge* required of the diagnostician is focused in this perspective.

1.2 Problem

Traditional studies of financial diagnosis take one of three approaches; a cognitive, a judgement modelling, or a predictive approach. The *cognitive approach* traditionally focuses on the information processing behaviour of the analyst. Based on the information processing theory of cognition, interviews or protocols are recorded and analysed to model the cognitive behaviour of the analyst. Studies with a *descriptive* cognitive orientation (e.g. Anderson, 1988; Biggs, 1984; Bouwman, 1983; Bouwman, Frishkoff & Frishkoff, 1987) use the full apparatus of models and methodology of traditional information processing theory (Newell & Simon, 1972). Studies with an *experimental* cognitive orientation (e.g. Libby & Frederick, 1990; Trotman & Sng, 1989) use the information processing theory to formulate hypotheses on, for example, information search, knowledge representation or experience effects, and use an experimental design to test these hypotheses.

A *judgement modelling* [1] (Ashton, 1981; Libby, 1975) approach focuses on the stimulus-response pairs of the analyst, and uses traditional linear models[2] to model the relationship between quantitative accounting information and the classification response of the analyst. This last approach is not necessarily intended to explain the "real" mode of information processing used to form judgements (Ashton, 1981, p. 13), but a model of the stimulus to response mapping is developed. The cognitive and the judgement modelling approaches are in many ways extensions of the cognitive versus stimulus-response debate in psychology (see Dennett, 1978).

In addition to these two behavioural traditions, financial diagnosis is studied as part of several task contexts in accounting and finance, focusing on developing a model with a purely *predictive purpose* (see Altman, Avery, Eisenbeis & Sinkey, 1981).

---

[1] Also referred to as the behavioural approach (e.g. Bedard, 1989)
[2] Traditionally, linear discriminant analysis or linear regression analysis.

When both an economic criterion variable and a human judgement of the variable exist, the judgement modelling approach above and the predictive approach can be combined in the lens model of Brunswik (1952).

Even though the judgement modelling approach has several methodological advantages, the underlying cognitive theory of analysts' information processing is underspecified. In most studies, weights of a linear model are the only "internal representation" required to explain cognitive behaviour. The dominating theory for explaining such cognitive behaviour has for a long time been the information processing theory of Newell and Simon (1972) and its more knowledge intensive successors (e.g. Newell, 1990). However, the separation of cognition from verbalisation (Ericsson & Simon, 1984; Nisbett & Wilson, 1978; Nisbett & Ross, 1980), is one among several methodological problems in this theory. In addition to the methodological criticisms, this theory has been met with general arguments raised by researchers in philosophy (e.g. Dreyfus, 1972) and linguistics (e.g. Lakoff, 1987). In cognitive science, what has been termed as an "anti rule movement" has proposed alternative explanations for cognitive processing previously assumed only to be explainable by information processing theory (see Smith, Langston & Nisbett, 1992).

In the applied field of cognitive and behavioural accounting there has been a growing need for theories paying more attention to pattern recognition and pattern matching (e.g. Bedard & Biggs, 1991; Bouwman et al., 1987), schematic organisation of memory (e.g. Choo & Trotman, 1991), and analogical reasoning (e.g. Biggs, Messier & Hansen, 1987). Cognitive theories focusing on these phenomena, such as classification theory, have received little attention in cognitive and behavioural accounting. Even though the judgement modelling and predictive approaches have treated the financial diagnosis task as a categorisation or classification task, cognitive studies have been preoccupied with using the hypothesis testing and problem solving approach offered by information processing theory. This may have been due to the assumption that classification theories in cognitive psychology provide little room for the use of intermediate representations and abstracted features in cognitive processing. This may well have been right, but recent progress in the area has opened new avenues of research.

Recently, a new collection of cognitive theories under the term "connectionism" (see Rumelhart & McClelland, 1986; Smolensky, 1988) has been explored in several areas. It offers an orientation where the stimulus-response pairs of the subjects are in focus, but where a cognitive model of the representations and processing necessary to map stimulus to response is developed with connectionist methodology. In financial diagnosis, this mapping takes the form of a classification. To simplify, connectionism uses methodological instruments and principles similar to the judgement modelling approach to model cognitive representations and processes of the individual performing the task (Seidenberg, 1993). With a possibility to develop cognitive models with methods free from many of the limitations of information processing methodology, connectionism offers a way to unify theoretical approaches in cognitive and behavioural accounting.

Several authors (e.g. Bedard & Biggs, 1991) have suggested that initial financial diagnosis may be seen as a pattern recognition problem, in which the analyst forms an opinion based on the recognition of patterns of cues in the financial statement. At the heart of pattern recognition lies the idea of an organisation of memory that facilitates recognition by matching represented to observed patterns (Rumelhart, Smolensky, McClelland & Hinton, 1986). In exemplar theories, classification consists of a measurement of the similarity of new and known preclassified patterns, and a classification based upon this similarity measure (Estes, 1994). Information processing theory has been criticised for its inability to explain similarity based pattern recognition (Dreyfus & Dreyfus, 1987; Winograd & Flores, 1986) and classification (Estes, 1994), and connectionist theory has been suggested as an alternative (Smolensky, 1988).

In this study, we investigate the properties of connectionist classification models of financial diagnosis. In particular, three questions are raised. First, what properties do connectionist classification models have as models of financial diagnosis? Second, how do connectionist models fit financial diagnostic behaviour when compared to traditional models, and third, how does the capacity of some connectionist models to develop internal representations apply to the financial diagnosis task?

1.3 Purpose

In this dissertation, two aims are of primary relevance. The first is to investigate how these new *theoretical* perspectives in cognitive science can be applied to financial diagnosis to increase our understanding of the task. The second aim is to investigate *empirically* how the perspectives can be used and evaluated as models of financial diagnostic behaviour. The two purposes need further elaboration.

Since classification theory has proven relevant in explaining human behaviour in other diagnostic tasks, this study aims at investigating the relevance of a classification conception of financial diagnosis. In particular, we are interested in investigating how the most recent classification theory, connectionist theory, can contribute to behavioural and cognitive explanation. Since connectionist models can develop cognitive models using judgement modelling methodology, they may be suited to unify different approaches in cognitive accounting applied to the financial diagnosis task. To provide a basis for the application of connectionist classification theory, the most important properties of the theory must be explained and clarified.

Even though many of the theoretical conclusions in this study may be relevant to other cognitive accounting tasks, our focus is on financial diagnosis. The financial diagnosis task was chosen as a representative cognitive accounting task for four reasons. First, input to the task is quantitative financial information, characteristic of many cognitive accounting tasks. Second, the task is much investigated both from a behavioural and a predictive perspective. Third, financial diagnosis is a task of economic significance to firms and analysts. Fourth, the task has been investigated in a previous research project at the Norwegian School of Economics and Business Administration (Methlie, 1993, 1994), and our project is based upon the knowledge generated in that project, and hopefully, adds further knowledge to it.

As our second aim, we wish to investigate *empirically* how connectionist classification theory can be used to model financial diagnostic behaviour. Since connectionist modelling still is in its youth, new methodological principles must be developed and adapted to fit this application task. We intend to develop and adapt methodological principles that make connectionist modelling practically applicable to cognitive accounting tasks. As a cognitive accounting task investigated to a considerable extent, financial diagnosis research provides benchmark models making comparisons of connectionist models with traditional models possible. By using such benchmarks, we suggest that despite the exploratory nature of connectionist modelling, formulation of models and evaluation of derived propositions are possible[1].

In addition to evaluating connectionist models' fit to financial diagnostic behaviour, we wish to investigate if connectionist models' internal representations offer an additional source of information for model evaluation. Analysis of connectionist models' representations may offer a new way to link empirical models and theoretical principles not possible within previous approaches to financial diagnosis.

1.4 Organisation of the dissertation

The remaining parts of this dissertation is organised as follows. Part II consists of three chapters. In chapter 2, a *task analysis* of the financial diagnosis task is performed by reviewing relevant research within each of the judgement modelling, the cognitive and the predictive approaches to financial diagnosis. Within each approach, research on different task

[1] These are methodological principles traditionally applied to research in the "context of justification".

contexts of financial diagnosis is reviewed. Based upon this task analysis, a classification conception of the financial diagnosis task is chosen.

In chapter 3, relevant cognitive theory is introduced. To provide the necessary theoretical basis for applying a classification conception of the financial diagnosis task, the definitional, prototype and exemplar *theories of cognitive classification* are introduced in section 3.1. In addition, some examples of cognitive classification models are given. Among the most recent cognitive classification models are the *connectionist* models of classification. These models have several desirable properties as models of financial diagnosis, such as the ability to develop internal representations. However, connectionist models are relatively new, and consequently, an in-depth presentation of their theoretical basis is considered necessary. This presentation is given in section 3.2. Section 3.2 ends with a review of the most well-known connectionist models of classification and a discussion of their relevance as models of financial diagnosis.

In chapter 4, the task analysis of chapter 2 and the theoretical basis given in chapter 3 is used to propose a *connectionist model of financial diagnosis*. Furthermore, three *propositions* that can be evaluated empirically are made.

In part III, the *method* used to evaluate the propositions of chapter 4 is presented in two chapters. An empirical evaluation of our connectionist model requires two operations. First, a set of valid stimulus-response pairs must be provided. Second, a set of simulations must be designed where the connectionist model "learns" to map the stimulus to response. In chapter 5, the experimental *research design* used to provide the stimulus-response pairs of financial diagnoses is presented. In chapter 6, the methodological aspects of the *simulation design* used in this study are reported.

Different operationalisations of the stimulus-response pairs representing valid financial diagnosis are applied. In part IV, the *results* of three simulations using our connectionist model of financial diagnosis and three different stimulus-response operationalisations are reported.

In chapter 7, an operationalisation corresponding to a bankruptcy classification context of the financial diagnosis task is used. In chapter 8 and 9, composite judge operationalisations of financial diagnoses are simulated using two different stimulus operationalisations. For each of the three simulations, the three propositions of chapter 4 are evaluated in reports of model performance and by representational analyses. At the end of each chapter, main conclusions resulting from the model simulations are summarised.

In part V, the main conclusions of this study are discussed. Factors limiting the validity of our conclusions are discussed, and improvements are suggested. Many of these improvements gives suggestions for further research on the application of connectionist theory and models to financial diagnosis.

PART II - THEORY

Chapter 2. Financial diagnosis theory

As referred to in chapter 1, part I of this dissertation introduces and reviews two areas of research. Relevant theory and empirical research on the financial diagnosis task are reviewed in this chapter. To understand financial diagnosis as a classification task, cognitive classification theory is introduced. Alternative theories of classification are introduced in chapter 3, and both general connectionist and connectionist classification theories are presented. Due to the novelty of our connectionist classification perspective on the financial diagnosis, these theories are given an in-depth presentation. In chapter 4, selected research of the two areas are merged into a connectionist model of financial diagnosis.

This chapter reviews selected studies of relevance to financial diagnosis, and shows how financial diagnosis can be defined as a task that is largely similar across task contexts. We summarise the main findings on the task as viewed from three perspectives; a judgmental, a cognitive, and a predictive perspective. This chapter does not intend to summarise or review the vast literature of experimental and descriptive cognitive accounting research. This has previously been done with focus on comprehensive review (Ashton, 1982; Libby, 1981), on special research questions or disciplines (Bedard & Chi, 1993; Bonner & Pennington, 1991; Ho & Rodgers, 1993), on particular research methodology (Klersey & Mock, 1989; Rodgers, 1991b), on unifying perspectives (Peters, 1993) or with intentions to suggest fruitful research directions for the future (Gibbins & Jamal, 1993; Libby & Luft, 1993).

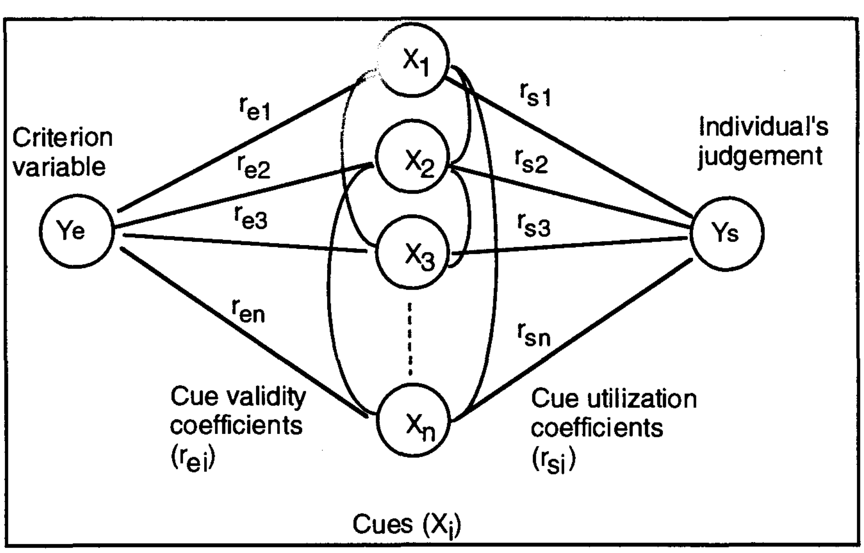

Figure 2.1 Lens model of Brunswik (1952)(From Ashton, 1982).

The lens model (Brunswik, 1952; see Ashton, 1982) can be used to illustrate different theoretical approaches to the study of financial diagnosis[1], Figure 2.1 shows the two parts of

[1] And similar financial judgement and accounting judgement situations

the lens model. The cues, termed $X_1, X_2, \ldots, X_n$, are used to predict either an individual judgement, termed Y_s, or a criterion variable, termed Y_e. Different types of cues can be used, but traditionally, the cues are collected from the financial statement of a firm. An example of a judgement variable may be subjects' individual judgements of a firm being bankrupt. An example of a criterion variable may be a variable representing whether the firm is bankrupt or not. The lens model illustrates the two models that can be developed using the same cues. If the model is used to predict judgements, in principle, a *cognitive model* is developed. If the model is used to predict the criterion variable, in principle, an *economic model* is developed. In the original lens model, the criterion variable is used to evaluate the accuracy of individuals' judgements and the predictive accuracy of both models. The original lens model uses the same kind of model to predict both judgements and criterion variables. The standard model is a linear weighting model equivalent to what is found in traditional regression or discriminant models. Weights in the model of the criterion variable are termed cue validity coefficients (r_{ei}), and weights in the model of the judgement variable are termed cue utilisation coefficients (r_{si}).

The lens model can be used as a framework for introducing the three approaches to financial diagnosis found in the literature. A *judgement modelling* approach traditionally uses all aspects of the lens model in its development of both a model of the individuals' judgement and a model of the criterion variable. Next, evaluations of both models' and subjects' accuracy are performed.

Cognitive approaches focus on the right hand side of the lens model. In an experimental cognitive approach, the researcher formulates hypotheses on variables of relevance to the judgmental process, and tests these. In a descriptive cognitive approach, the researcher uses information processing theory and methodology to describe the information processing necessary to perform the judgement of the right hand side of the lens model. Traditionally, no evaluations of the predictive accuracy of judgements and models are performed.

In a *predictive approach* , the researcher concentrates on modelling the left hand side of the lens model without reference to cognitive theory, but purely based upon an economic theory of the process leading to the event measured by the criterion variable.

To illustrate the three perspectives, the prediction of bankruptcy is representative. An analyst may evaluate whether a company is likely to go bankrupt or not, and this judgement can be modelled at the right hand side of the lens model. However, it is easy to confirm whether the company actually went bankrupt, and this event may be predicted using the same indicators as in the model of the judgement. The first of these models is a model of the judgement process, and the second is a model of the economic process leading to bankruptcy. However, the explicitness in the literature regarding whether the first model is a model of the judgement process or not varies:

> *"These, like all models, are abstractions and do not purport to represent "real" mental processes" (Libby, 1981, p. 22).*

In principle, an economic theory should underlie the economic model on the left hand side of the lens model, and a cognitive theory should underlie the cognitive model of the right hand side. Whether a cognitive theory underlies the model of the individual judgement, varies with tradition, discipline and task context. To conclude, the right hand side of the lens model has been modelled with cognitive theory, with economic theory, or with a combination of both as basis. The left hand side is always modelled with an economic theory as basis.

A second aspect illustrated by the lens model is the view that a financial diagnosis task is a mapping of cues to judgmental variables. Traditionally the task is performed by mapping the high dimensional stimulus space to the lower dimensional response space.

With these aspects of the financial diagnosis task introduced, we can define the financial diagnosis task in the following way:

> A situation where the subject makes a judgement of the financial situation of the firm based upon information from the financial statement.

When used as a reference for selecting relevant empirical research, the definition provides four criteria for a study of financial diagnosis to be of relevance. First, a human decision maker must express an opinion on the financial situation of the firm. Next, financial statement cues or information must be provided. Third, financial analysis or diagnosis must at least be part of the task studied. Fourth, there must be a focus on judgmental or behavioural aspects of the task. Several studies of relevance do not satisfy the fourth criteria, but incorporate an economic theory of how the opinion should be formed. These studies use a dependent variable that is expressed by a human decision maker, but formulate their theory on how the opinion is formed on an economic theory rather than on a cognitive theory. In the context of the lens model, these studies operate purely on the left hand side of the model, but their economic theory is relevant to the knowledge assumed represented in a cognitive model of the right hand side.

The traditional view is that the context and purpose of the financial diagnosis is a functional way to classify how it is performed. This view is found in some traditional textbooks. Foster (1986) treats financial analysis within the context of asset pricing, capital market and equity applications, corporate restructuring, debt rating, distress analysis, and loan decisions. However, main elements of the task, such as for example the importance of financial statement cues, are similar across contexts. A perspective, treating financial diagnosis as similar across task contexts, is found in other textbooks on financial statement analysis (e.g. Hawkins, 1986), and research has been performed supporting this as a useful perspective (e.g. Barnes & Huan, 1993).

Different task contexts of the financial diagnosis task are found within auditing, accounting and finance. In table 2.1 some of the relevant task contexts are presented[1].

Task context	Judgement	Cognitive	Predictive
Analytical review (ratio analysis)	Nelson, 1993	Bedard & Biggs, 1991	
Bankruptcy prediction	Simnett & Trotman, 1989	Chewning & Harrell, 1990	Ohlson, 1980
Bond rating		Danos, Holt & Imhoff, 1984	Kaplan & Urwitz, 1979
Earnings forecasting	Houghton & Woodliff, 1987	Biggs, 1984	
Going concern judgement	Hopwood, Mc-Keown & Mutchler, 1994	Biggs, Selfridge & Krupka, 1993	Koh & Killough, 1990
Investment screening		Bouwman, et al. 1987	
Loan decision	Rodgers, 1991	Danos, Holt & Imhoff, 1989	Doukas, 1986
Risk assessment	Mear & Firth, 1988	Holt & Morrow, 1992	

Table 2.1 Task context and perspective combinations of financial diagnosis .

Not all approaches to each of the task contexts of table 2.1 are equally relevant. Thus, some of the approaches are without example studies. Some combinations of contexts and approaches may never be found. One example is studies in bond rating, where it is unlikely that process tracing methods will ever be allowed in studying bond raters' diagnostic behaviour.

An important similarity across task contexts is that despite differences in original disciplines and research traditions, the financial statement contains the most relevant cues in performing the task. In auditing, parts of the analytical review and going concern judgements are task contexts in which financial diagnosis is performed. Of particular relevance to this study is research on the task context classified as "preliminary going concern evaluation" in Bonner and Pennington's (1991) classification of audit tasks. In accounting, the bankruptcy prediction task is closely related to the going concern judgements of auditors, but the cues, context and

[1] Relatively recent studies are given as example studies to show that research within all three approaches still is relevant.

purpose of the classifications may differ. In banking, the loan decision shares many of the characteristics of the bankruptcy prediction task when there is doubt about granting a loan. In finance, closely related but "reversed" task contexts are the investment screening and risk assessment tasks, both relying heavily on information from the financial statement. Parts of all these tasks meet the criteria in our definition of a financial diagnosis task given above. Our proposal is that the financial diagnosis part of the tasks listed in table 2.1 may have more similarities across disciplines than other tasks have within one specific discipline. Similar propositions have been made by Bonner and Pennington (1991) for auditing tasks, and by Gibbins and Jamal (1993) for several accounting tasks.

Methodologically, there are also three approaches to the study of financial diagnosis. Early studies adopted the lens model orientation illustrated in figure 2.1 (e.g. Libby, 1975). Studies formulating a cognitive model of the stimulus-response mapping can take one of two methodological orientations. With an experimental orientation, hypotheses about the cognitive representations and processes intermediating stimulus and response are formulated, and attempts are made to set up an experimental design to test these hypotheses. A process orientation has a more descriptive purpose, and uses protocol analysis or other process tracing methods to investigate the cognitive processes of the subjects during performance of the task.

Table 2.2 illustrates the differences in research focus between the judgmental modelling, the cognitive, and the predictive approaches to financial diagnosis by listing some of the most relevant research questions pursued within each of the approaches.

Judgement modelling	Cognitive experimental	Cognitive descriptive	Predictive
Analyst and model accuracy	Experience effects	Description of problem solving behaviour	Cue predictability
Cue usage, utilisation and selection	Information load effects	Cue usage, utilisation and selection	Properties of cues (e.g. probability distributions)
Agreement (consensus)	Information format effects	Description of representational forms	Test or development of conceptual model
Stability and consistency	Effects of representation and organisation of knowledge	Description of reasoning and search strategies	
Environmental predictability	Effects of reasoning and search strategies	Goal organisation and subgoaling	
Analysts self-insight	Effects of process verbalisation		

Table 2.2 Research questions in different approaches to the study of financial diagnosis

In addition to studies of direct relevance to the financial diagnosis task, several contributions are found within the cognitive accounting literature studying related tasks that may be of relevance to our study. One example is studies of auditing tasks relying explicitly on information processing theory that formulate and test hypotheses on subjects organisation of knowledge (Biggs & Wild, 1985). Selected findings from such studies are reviewed in section 2.4

A special area of research difficult to place within our framework is the research utilising machine learning algorithms for rule induction (e.g. Frydman, Altman and Kao, 1985; Hansen, Koehler, Messier and Mutchler, 1993). Studies in this area that apply rule induction to the right hand side of the lens model, will be treated as part of the judgement modelling approach. Studies using rule induction methodology primarily as a predictive method, are treated as studies with a predictive approach. This separation of contributions is similar to how we treat studies applying the same statistical method with a judgement modelling or a predictive purpose.

In section 2.1, selected studies within the judgement modelling approach are reviewed, followed by a review of selected contributions in the cognitive approach in section 2.2. Selected predictive studies relevant to the financial diagnosis task are reviewed in section 2.3. In section 2.5, a summary of supported standard assumptions on the financial diagnosis task is presented along with a summary of some of the most relevant areas for further research of relevance to this study.

The review of each approach is organised as a simplified problem solving process, reviewing findings related to the task first, information search and usage second, reasoning processes third, representation and knowledge organisation fourth, and outcome of the diagnosis task last.

2.1 The judgement modelling approach

Studies within the judgement modelling approach to financial diagnosis vary in their application of the complete lens model of Brunswik (1952) shown in figure 2.1. Early studies are easily classified as lens model studies, but later studies have concentrated more on the right hand side of the lens model (e.g. Holt & Carrol 1980), departed from the use of quantitative cues only (e.g. Schepanski, 1983), and presumed a conceptual model underlying the right hand side of the model (e.g. Rodgers, 1991a). Thus, a judgement modelling study is no longer equivalent to a standard lens model application. An overview of judgement modelling studies relevant to financial diagnosis is shown in table 2.3, with the studies organised in chronological order within each task context. In addition, the focus of each study is indicated.

Reference	Task context	Focus
Libby, 1975	Bankruptcy prediction	Accuracy of subjects and linear model
Abdel-khalik & El-Sheshai, 1980	Bankruptcy prediction	Effect of using self-selected cues
Casey, 1980a	Bankruptcy prediction	Replication of Libby (1975)
Zimmer, 1980	Bankruptcy prediction	Replication of Libby (1975)
Houghton, 1984	Bankruptcy prediction	Variations in age of data
Houghton & Sengupta, 1984	Bankruptcy prediction	Variations in prior probabilities
Chalos, 1985	Bankruptcy prediction	Comparison with committee assessments
Messier & Hansen, 1988	Bankruptcy prediction	Recursive partitioning
Selling & Shank, 1989	Bankruptcy prediction	Comparison with process tracing
Simnett & Trotman, 1989	Bankruptcy prediction	Optimal cue selection
Houghton & Woodliff, 1987	Earnings forecasting	Model differences for success and failure prediction
Kida, 1980	Going-concern judgement	Standard lens model application
Hansen et al., 1993	Going-concern judgement	Recursive partitioning comparison
Hopwood et al., 1994	Going-concern judgement	Realistic prior probabilities
Holt & Carroll, 1980	Loan decision	Standard lens model application
Dietrich & Kaplan, 1982	Loan decision	Limited lens model application
Schepanski, 1983	Loan decision	Critique of linear models
Chalos & Pickard, 1985	Loan decision	Comparison with committee assessments
Rodgers & Housel, 1987	Loan decision	Conceptual model LISREL
Rodgers & Johnson, 1988	Loan decision	Conceptual model LISREL
Rodgers, 1991a	Loan decision	Conceptual model LISREL
Mear & Firth, 1988	Risk assessment	Right hand side predictability
Mear & Firth, 1987a	Risk assessment	Cue usage and self-insight
Mear & Firth, 1987b	Risk assessment	Different accuracy measures
Wright, 1977	Stock recommendations	Lens model right hand side application

Table 2.3 Selected judgement modelling studies of financial diagnosis

There are mainly four task contexts found among the judgement modelling studies shown in table 2.3. Research within the bankruptcy prediction task context has shown an additive knowledge accumulation with replications, refinements and alternative use of methods, making findings easily comparable. Within the other three tasks contexts, the findings are much more difficult to unify.

Within the bankruptcy prediction task context, the *cues presented* to the subjects are few, and correspond to cues found useful in predictive studies. Traditionally all cues are *ratio cues* relating two or more traditional items of the financial statement to each other. However, some of the studies provide the subjects with financial statements excluding calculated ratios, whereas other studies include them. Obviously, task content changes when computations are required by the subjects. In Chalos' (1985) study of both individual and group judgements, *full financial statements* were provided. In the studies of students' and loan officers' loan decisions by Rodgers (Rodgers, 1991a; Rodgers & Housel, 1987, Rodgers & Johnson, 1988), both financial statement information and financial ratios were provided.

Some of the studies use statistical techniques to select relevant cues to be presented. One example is the early study of Libby (1975), in which factor analysis was used to select presented ratios. The selected ratios were assumed to represent "profitability", "activity", "liquidity", "asset balance" and "cash position"[1]. Similar to studies within the predictive approach, a conceptual model or theory is rarely used to justify the presented ratios. In the study of Kida (1980), the presented ratios were based upon significance of the ratios in a linear model. This selection procedure will favour the use of a linear model even by the cues presented[2]. Traditionally, the cues presented represent financial cues of more than one period.

An alternative way of presenting cues to the subjects is by letting the subjects select the cues from a menu. This approach was used by Abdel-khalik and El-Sheshai (1980), allowing the subjects to purchase the cues. Also in Selling and Schanks' study (1989), the cues were selected by the subjects, but the cues had no price, and a maximum of seven cues could be selected. A similar approach to cue selection was followed by Chalos and Pickard (1985) in a loan decision task. The effects of variations in presented cues have received little attention within judgement modelling research.

[1] In general, theoretical concepts, terms used for internal representations, and latent variables are placed in double quotation marks.

[2] In principle, the same bias is introduced when linear methods are used to select cues in a predictive model.

The *cases* [1] used in the judgement modelling studies are often selected from corresponding studies with a predictive approach (Libby, 1975; Hansen et al., 1993). The traditional sample is composed of cases belonging to one of two classes with equal prior probabilities of occurrence. The bankruptcy prediction studies are representative. The subjects are presented distressed and non-distressed cases with equal probabilities of occurrence. Some studies have applied more realistic probabilities of occurrence (Houghton & Sengupta, 1984), or adjusted results for equal and unrealistic probabilities (Hopwood et al., 1994). Some studies have also examined the effects of providing the subjects with information of the prior probabilities (e.g. Abdel-khalik & El-Sheshai, 1980, Houghton, 1984). Similar case selection strategies are found among studies of other relevant task contexts.

Cue usage is measured in some of the studies by analysis of standardised betas in linear models of the subjects' judgement, but mainly it is assumed that cues presented are the cues used. Libby (1975) used the following five ratios[2]: Net income/total assets, current assets/sales, current assets/current liabilities, current assets/total assets and cash/total assets. These ratios were also used by Casey (1980a). In addition to the net income/total assets ratio used by Libby (1975), Zimmer (1980) used dividend/earnings, debt/cash flow, long-term debt/equity and a quick assets ratio. These ratios were assumed to represent the theoretical concepts of "profitability", "dividend policy", "debt coverage", "long term solvency" and "short term solvency", respectively. The cues selected by Zimmer (1980) have also been used in more recent studies (Houghton, 1984; Houghton & Sengupta, 1984; Houghton & Woodliff, 1987) even though other concepts were assumed to be measured by the same cues[3].

Kida (1980) studied the financial diagnosis task within a going concern context, but used ratios very similar to Libby's (1975). He used net income/total assets, net worth/total debt, quick assets/current liabilities, sales/total assets and cash/total assets, measuring the theoretical concepts of "profitability", "leverage", "liquidity", "capital intensiveness" and "cash position", respectively. In other task contexts, *special financial ratios* of particular relevance to the task context were used. One example is the use of market related financial statement information in the risk assessment studies of Mear and Firth (1987a, 1987b, 1988). Holt and Carroll (1980), stressed the importance of *trend indicators*. They used both earnings and sales trend ratios calculated over three years. Chalos and Pickard (1985) provided their subjects with trends of traditional ratios such as net income/total assets, total debt/net worth, working capital/sales and the acid ratio. Trend ratios have also been used in studies of the bankruptcy prediction, risk assessment, and stock prediction tasks (e.g. Abdel-khalik & El-

[1] The firms from which the financial statement information is collected.
[2] Throughout this chapter, ratios are presented with terms used by the original authors.
[3] Houghton and Woodliff (1987) assume their cues measure "income", "liquidity", "dividend policy", "cash flow" and "leverage".

Sheshai, 1980, Mear & Firth, 1987, Wright, 1977). Recent studies comparing linear and nonlinear methods have also incorporated trend ratios (e.g. Hansen et al., 1993).

Judgement modelling studies vary in the extent to which they use cues *outside* the traditional financial statement. Early judgement modelling studies primarily used financial statement cues only. Examples of cues from sources other than the financial statement that have been provided and tested for cue usage, are market related cues, such as the ratios market equity/liabilities or book equity/market equity (Simnett & Trotman, 1989). Other market related information typically used in risk assessment tasks, are cues measuring systematic risk, such as beta (Mear & Firth, 1988; Wright, 1977). In addition, customer history information has been used in the loan decision task. Customer history is often represented by past loan history and overdues (Holt & Carroll, 1980). Industry related cues have been used both in the loan decision task and in risk assessment tasks. Two examples are the industry stability and trend ratios used by Holt and Carroll (1980), and the expected industry results index used by Mear and Firth (1988).

Some studies have transformed the quantitative financial statement cues into *qualitative* information. The transformation has been founded on results in the cognitive approach to financial diagnosis, indicating that the subjects transform cues to qualitative information before the judgement is performed (e.g. Bouwman et al., 1987). Schepanski (1983) followed this approach in a study comparing linear models to alternative nonlinear models. In these cases, only a few qualitative categories are used for each cue, such as "high" and "low" for levels or "up" and "down" for trends.

The accuracy of *outcomes* is of relevance to both individual judgements and models of individual judgements. Judgement and model accuracy measures vary with the task studied. If the dependent variable is nominal or ordinal, classification errors are reported. If the variable is interval, some distance or correlational measure can be used. Classification accuracy is less sensitive to variations in a distance measure when there are few categories. This makes comparisons of accuracy across task contexts difficult. A classification accuracy measure is also less sensitive to variation in the performance of alternative models. Consequently, care should be taken in interpreting classification error differences as strong support for model differences. In a judgement modelling approach, the effect of using various performance measures was investigated by Mear and Firth (1987). They found that, although their results were consistent with prior research when measured by correlational performance measures, the use of other error measures produced conflicting indications of analyst performance. Since the reported error measures vary with judgement variable, and thus, task context, results on accuracy and predictability in judgement modelling studies will be reviewed separately for each task context.

The accuracy of a model of the criterion variable is termed *environmental predictability*. In the *bankruptcy prediction* studies reported in table 2.3, environmental predictability is high. Both in lens model studies, and in predictive studies, classification errors less than 10 % are common (Altman et al., 1981; Ashton, 1982). The *accuracy* of judgements are generally found to be lower than the environmental predictability (Ashton, 1982). Recently, the relationship between environmental predictability and accuracy was investigated by Hopwood et al. (1994). They showed that environmental predictability was lower and accuracy was higher than previously assumed when accounting for prior probabilities, misclassification cost and separation of stressed and non-stressed companies.

The standard propositions basically confirmed by judgement modelling studies in *bankruptcy prediction*, are that the predictability of the task is high. Individual judgements are better than random assignment, but individuals are significantly outperformed by the model of the criterion variable, and by the composite judge[1]. Further, they are slightly outperformed by a linear model of their judgements. These results are assumedly explained by individuals' inferior selection of cues, not by inferior cue weighting. Judgements are thought to be consistent, and small differences among individuals are found. Students work well as surrogates for more experienced or professional subjects.

The standard propositions are the result of a number of judgement modelling studies. Libby (1975) found an environmental predictability of 88 %. The decision rule applied by individuals was highly linear and had an accuracy of 74 %. The composite judge outperformed the individuals with an *accuracy* of 82 %. Stability, consistency and consensus were high. These results summarise most of the hypotheses later tested in the judgement modelling approach. In a replication of Libby's study, Casey (1980) found considerably lower accuracy of both individuals and composite judges. Furthermore, the composite judges did not significantly outperform the individuals. Casey (1980) suggested the difference in findings may be explained by non-disclosure of prior probabilities. Zimmer (1980) also replicated Libby's study, finding results very consistent with the results of Libby (1975), and supporting the effect of disclosing priors. Abdel-khalik and El-Sheshai (1980) investigated how self-selection of cues and disclosure of priors affected the judgements. They found that very few cues were selected, and that earnings trend, current ratio and cash flow/total debt were most frequently used. Predictability was found to be high, and the disclosure of priors did not improve accuracy. Both model and human selection of cues were compared, leading Abdel-khalik and El-Sheshai (1980) to conclude that the inferiority of subjects could be explained by wrong selection of cues, and not by cue weighting. Houghton (1984) also tested

[1] Composite judge diagnoses are computed as the average response of a committee of judges.

the effect of disclosing priors, and found accuracy to improve when subjects were given information of the prior probability of failure. In addition, he tested how the age of data affected the subjects' accuracy, finding that accuracy decreased with the age of data. Houghton (1984) suggested that the effect of priors may not only be affected by disclosure of prior information. In Houghton and Sengupta (1984), the effect of more realistic probabilities of failure was tested. It was found that accuracy improved when probability of failure was half the probability of non-failure. However, the interpretation of this finding should be done with care, since all studies show that the most frequent error is the false classification of failed cases as non-failed.

The hypothesis stating that the performance of interacting and composite judges is better than that of individuals was tested by Chalos (1985). He found the following ranking of accuracy: Interacting committee>composite judge>model of the criterion variable>model of individual judgement>individual[1]. However, the difference in performance between the interacting committee and the composite judge was not statistically significant. The hypothesis that individuals' inferiority is caused by inoptimal cue selection proposed by Abdel-khalik and El-Sheshai (1980), was tested by Simnett and Trotman (1989). They found sub-optimal selection of cues, and failure to improve accuracy when subject-selected cues and model weights were used. Both findings support the hypothesis of Abdel-khalik and El-Sheshai (1980). When provided with the model-selected cues, however, the subjects' performance deteriorated. Simnett and Trotman (1989) suggested that information processing was not a limiting factor when subjects selected their own cues, but became a limiting factor when subjects were provided with model selected cues. Consequently, most of the standard propositions referred above regarding the bankruptcy prediction task still stand, even though their support is less conclusive than often cited.

The capacity of a linear model to capture the judgements of human subjects has been questioned within the bankruptcy prediction task context. Selling and Schank (1989) set up an interesting experiment registering the subjects' cue usage during the task. Next, they modelled the subjects' judgements both in a linear model and in a decision tree . Both the linear and the decision tree models were able to approximate the subjects' judgement with high accuracy, but the cue importance measures of the two models showed very little convergence. A major weakness of this study was that the cues were disguised. The development of a decision tree in the Selling and Schank (1989) study, is closely related to the general application of recursive partitioning models within the judgement modelling approach. While the development of the decision tree of Selling and Schank (1989) was driven by the subjects cue selection sequence, recursive partitioning uses mathematical algorithms to produce the

[1] The symbol ">" represents "better than".

decision tree. Messier and Hansen (1988) developed a decision tree using a rule induction method attributable to Quinlan (1986), and compared their model with the results of Abdel-Khalik and El-Sheshai (1980), and Libby, Trotman and Zimmer (1987). The rules induced by their system were few and simple. Their model had a significantly better accuracy on a holdout sample, suggesting that the model outperformed the previous models, individuals and composite judge committees. However, no discussion was found in the study on the realism of the resulting production system as a model of the subjects' represented knowledge.

The judgement modelling literature on *earnings forecasting* is closely related to the time-series approach to the same task context found within the predictive orientation[1]. An exception is the study of Houghton and Woodliff (1987) in which the earnings forecast output was transformed to a failure and success prediction output. When viewed as a failure prediction task, the results of this study were very similar to the bankruptcy prediction task results reviewed above. When viewed as a success prediction task, however, the task seemed much more difficult for the subjects to perform. This study also confirmed that students were adequate surrogates for more experienced subjects in the earnings forecasting task context.

Many similarities are found between the bankruptcy prediction task and the *going concern* judgements of auditors. One of the earliest studies investigating this judgement, was the study of Kida (1980). His findings were very similar to the findings on the bankruptcy prediction task context reviewed above, but when the subjects should decide whether to *qualify* for going concern problems or not, the difference between the two tasks contexts became obvious. Subjects decided not to qualify for 24.6 % of the cases predicted as distressed firms. This result indicated that several other considerations were made by the subjects when deciding to qualify or not. The same result has been confirmed in cognitive experimental studies (e.g. Barnes and Huan, 1993), and also explains the finding of Altman and McGough (1974), that even though a model predicted going concern problems for 82 % of the cases, auditors only qualified in 44 %. In a recent study, Hopwood et al. (1994) questioned both the standard assumptions of the bankruptcy prediction task context, and their relevance to the going concern context. They showed that when misclassification costs were considered, and obviously non-stressed firms were eliminated from the analysis, both models and auditors were quite poor predictors of failure. However, no traditional accuracy measures were provided in this study.

A considerable overlap exists between the bankruptcy prediction task context and the loan default or loan classification tasks. In some cases, it is not clear how these contexts should be separated (see e.g. Chalos, 1985), since the relationship between loan default and bankruptcy

[1] The approach typically focuses predictions made by models using time series of one or very few financial statement items, and thus, is of less relevance to this study given our definition of the financial diagnosis task.

is so close. In a loan officer classification study, Holt and Carroll (1980) used a standard five valued ordinal classification variable, intended to measure *loan risk.* Classification results of this study were not comparable to failure prediction results because of differences in the scale of the dependent variables. However, a simple linear model predicted 77 % of the evaluations of the loan officers correctly when the output was transformed to a dichotomous scale. This model contained two financial statement cues and four exogenous variables. One of these was a variable measuring the classification derived from last years' review. The two most significant financial variables were an interest coverage ratio and a debt/equity trend ratio. Holt and Carroll (1980) used artificial cases built by combining information from real cases to cover more variations in the combination of cue values. One important comment from the debriefing interview was that this procedure was discovered by the loan officers, indicating that their ability to detect unrealistic cases is high. Dietrich and Kaplan (1982) used real cases and a similar classification variable to Holt and Carroll's (1980). In their analysis, the original classification variable was transformed to a four valued ordinal scale dependent variable and to a dichotomous dependent variable. The purpose of the transformations was to allow comparisons with predictive studies of failure detection. Dietrich and Kaplan (1982) found that probit analysis accuracy was somewhat better than regression analysis in models of the loan classifications. The models used three financial statement cues: Debt/equity, funds-flow/fixed-commitments, and a sales trend indicator. The models were compared to the predictive models of Altman (1968) and Wilcox (1973), and were found to perform better than these models in predicting high risk loans. The study by Dietrich and Kaplan (1982) used empirical classifications from actual loan reports as a dependent variable, while model accuracy was compared to previous studies with a predictive approach. The judgmental character of the dependent variable used by Dietrich and Kaplan (1982) positions the study in a judgement modelling tradition, and comparisons with actual loan defaults could have been valuable. Chalos and Pickard (1985) used a more traditional lens model approach in a loan decision context, and some of the results of the study have been reviewed above and in Chalos (1985). However, the dependent variable used by Chalos and Pickard (1985) was very similar to a bankruptcy prediction variable, but was applied in a loan default context. This makes their findings more relevant to the bankruptcy prediction task context than to the loan decision.

Schepanski (1983) used the loan decision task context in a study comparing linear and nonlinear models. The traditional additive linear model was compared to four nonlinear models suggested in other information processing studies in psychology[1]. Qualitative cues were used, and a factorial design was set up. Schepanski (1983) found that the linear and nonlinear models were comparable on accuracy, but several of the nonlinear models better

---

[1] The alternative models of Schepanski (1983) were a multiplicative model, a constant-weight averaging model, a geometric averaging model and a range model (see Schepanski, 1983, pp. 584-585).

explained the effects of using constrained cue sets found in subdesigns of the full factorial design. Cue usage in the nonlinear models was quite different from that assumed in the linear models, confirming the findings within the bankruptcy prediction task context referred to above, that cue usage is sensitive to model. One finding in Schepanski's (1983) study was that a main assumption of a linear model was incorrect; namely that the additive effect of a cue is independent of the presence of other cues. The effect of a cue varied with the value of other cues in its context. One example was that greater effect was placed on a cue when the information set was small than when it was large. Schepanski (1983) hypothesised that paying attention to the size of the information set was a reasonable strategy if cues were correlated, something that is often eliminated in traditional judgement modelling studies. Thus, Schepanski's study indicated that evaluating predictive accuracy of a model must be combined with a qualitative investigation of model behaviour under varying conditions. The interaction effects found in Schepanski's (1983) study have been searched without success in later studies. As an example, Brown and Solomon (1991) suggested that the lack of similar findings may be partly due to how interaction terms are implemented in linear models, and partly due to the task formulations used.

In the judgement modelling studies referred to so far, the *conceptual models* of the right hand side of the lens model have been very simple[1]. In the cognitive approach to financial diagnosis, information processing theory is applied as a theory of the cognitive processes underlying the right hand side of the lens model, and this theory is used to formulate hypotheses of such a conceptual model. Rodgers and Housel (1987) were among the first to apply similar principles within a judgement modelling approach. Using a study of the loan decisions of students and loan officers, their results were reported in several articles (Rodgers & Housel, 1987, Rodgers & Johnson, 1988, Rodgers, 1991a). The conceptual model used in these studies was developed from findings in several studies of the cognitive descriptive approach . One of the most evident assumptions was that the financial diagnosis task could be divided into a perceptual and a judgmental phase (Bouwman, 1982[2]). Based upon this assumption, Rodgers and Housel (1987) compared the performance of experienced and less experienced subjects. They found that their conceptual model, implemented in LISREL (Jøreskog & Sørbom, 1988), fitted the cue usage and decision outcome data well. In Rodgers and Housel (1987), the conceptual model was not explicitly investigated. Instead, a more traditional comparison of the differences between experienced and less experienced subjects was performed. The authors also tested the differences between data driven and conceptually driven subjects' performance. Generally, performance results were comparable to research reviewed above, and no obvious difference was found between experienced and less

[1] Assuming an additive linear model, transforming the linear additive model by a sigmoid function, or applying the simple nonlinear models in Schepanski (1983).
[2] Bouwman (1982) applies the concepts familiarisation and analysis to these phases.

experienced subjects' performance. In Rodgers (1991a), only data from the experienced group was used to test the differences in behaviour and conceptual model between data driven and conceptually driven subjects. Differences were found both in behaviour and conceptual models, but two limitations should be noticed. In the student group used in Rodgers and Housel (1987), the difference in behaviour between the two groups was opposite from what had previously been found in the experienced group, making the total effect of different search strategies on behaviour insignificant. Next, the differences in conceptual models were tested by comparing the number of significant regression weights in the conceptual model of each group. However, this is not equal to testing the significance of the differences in regression weights of the two models. Consequently, a significant difference in the two conceptual models can not be inferred from Rodgers' (1991a) tests. Despite these weaknesses, the studies represent an important contribution to the formulation of a conceptual model of the loan decision process. An important finding of these studies was the support for a separation of the financial diagnosis task into a perceptual and a judgmental phase.

In a series of articles, the results of a study performed by Mear and Firth (1987a, 1987b, 1988), on the risk assessments of an equity security performed by 38 security analysts, were reported. The dependent variables of this study were an estimate of the expected risk and the expected 12-month return on the security. The scale of the dependent variables allowed the use of regression analysis to model the judgements, and a wide variety of error measures could be applied. In Mear and First (1987a), the cue usage and self-insight of subjects were reported. Subjects reported using indicators measuring systematic risk, proprietorship, profitability and liquidity in their risk judgements. The first three of these indicators were also the most significant in the linear model of the subjects' performance. Further tests of self-insight were performed, leading Mear and Firth (1987a) to conclude that the subjects had only moderate self-insight. Used in a linear model, the subjective weights performed worse than the model weights. However, when introducing unit weights, the model performed significantly worse than with the subjects' weights. In Mear and Firth (1987b), the use of alternative error measures in the study was reported. Using a traditional correlational error measure, results consistent with other studies of the relative performance of analysts and linear models were found. However, conflicting results were found when using other error measures. Of special interest were the differences between the relative performance of subjects and the linear model of the subjects. Mear and Firth's (1987b), study indicated that this difference may previously have been overstated. In Mear and Firth (1988), the hypothesis of Farelly, Ferris and Reichenstein (1985) that accounting variables of the financial statement contained information sufficient to predict perceived risk, was tested. The results were found to be consistent with this hypothesis. Only one market related independent variable was included in this study, making the risk judged by the subjects highly dependent upon financial statement information.

There are very few bond rating judgement modelling studies. Some exceptions are found in the municipal bond rating literature. Raman (1981), and Lewis, Patton and Green (1988), both applied the lens model approach to the study of analysts' rating of cities. However, the cues provided in these studies were so different from cues of traditional corporate financial diagnosis that the studies had only limited relevance to our task contexts.

Many of the *principles* of the judgement modelling approach have been used within more cognitively oriented studies. One example is cognitive studies using the cue importance measures of the lens model as a measure of the cognitive relevance of a cue (Chewning & Harrell, 1990). In the Chewning and Harrell (1990) study, the inconsistency of such an application was particularly obvious, since an inverted U-shape effect of information load was assumed, while simultaneously, a linear model estimates additive effects of information cues. Another example of how judgement modelling studies have influenced other approaches to financial diagnosis, is found in the experimental setup of cognitive studies examining the effect of variables outside the traditional lens model. These experiments are traditionally set up by using two or more parallel judgement modelling experiments. By comparing the results of these, a conclusion is drawn on the effect of the experimental variable (Iselin, 1993). Such experiments depend entirely on the applicability of a linear model. These examples were provided to illustrate the significance of judgement modelling studies to the way cognitive research is performed within behavioural and cognitive accounting.

2.2 The cognitive approach

In the late 1970's interest in a cognitive approach to information processing behaviour in accounting and finance started to grow. With the classical study of Libby (1975), an interest in [illegible] *differences* in cognitive behaviour between subjects with, for example, different experience or information load was started. This development led to an experimental orientation in the cognitive approach, in which the researcher used information processing based theory (see e.g. Hogarth, 1987) to formulate hypotheses of differences in information processing behaviour explained by variables outside the traditional lens model. Moriarty's (1979) study on the effect of alternative presentation forms was among the first to apply this approach. Simultaneously, research building on the traditional verbal protocol research in cognitive science using the full scale information processing theory and methodology, also appeared in cognitive accounting. Among the first of these studies was Bouwman's (1982) study of information processing behaviour during financial analysis. Since then, research with both orientations has continued within behavioural and cognitive accounting. Some differences between relevant research questions in the experimental and the descriptive orientations are illustrated in table 2.2.

First, we attend to relevant research within the experimental cognitive approach to financial diagnosis. Next, we review selected research within the descriptive cognitive approach. Selected contributions to financial diagnosis within the experimental cognitive approach are listed in chronological order within each task context in table 2.4. Since the studies are experimental, two columns indicating the main independent and dependent variables in each study are shown in table 2.4[1].

Reference	Task context	Independent var.	Dependent var.
Moriarty, 1979	Bankruptcy prediction	Presentation mode	Accuracy
Casey, 1980b	Bankruptcy prediction	Information load	Accuracy and time
Trotman & Sng, 1989	Bankruptcy prediction	Hypotheses framing and prior expectation	Information choice
Chewning & Harrell, 1990	Bankruptcy prediction	Information load	Cue utilisation and decision quality
Iselin, 1991	Bankruptcy prediction	Interacting vs. composite group	Performance
Iselin, 1993	Bankruptcy prediction	Information load	Decision quality
Danos et al., 1984	Bond rating	Experience	Adjustment of initial hypothesis
Kida, 1984	Going-concern judgement	Confirmatory strategy	Search and use of information
Choo & Trotman, 1991	Going-concern judgement	Knowledge differences	Recall
Ricchiute, 1992	Going-concern judgement	Working paper order	Performance and confidence
Barnes & Huan, 1993	Going concern judgement	Mitigating information	Performance adjustment
McGee et al., 1978	Investment screening	Personality	Performance, confidence and information use
Danos, Holt & Imhoff, 1989	Loan decision	Subsequent information	Decision and con-fidence adjustment
Rodgers, 1992	Loan decision	Perceptual strategy	Performance
Anderson, 1985	Offer price prediction	Verbalisation	Accuracy and time
Libby & Frederick, 1990	Ratio analysis	Experience	Explanations
Nelson, 1993	Ratio analysis	Knowledge and learned error frequency	Performance and frequency knowledge
Enis, 1988	Return prediction	Current-valued data	Accuracy and consensus
Holt & Morrow, 1992	Risk assess-ment	Type of experience	Conformance to Bayes theorem

Table 2.4 Selected experimental cognitive studies of financial diagnosis

[1] The author's terms are used of the dependent and independent variables.

One of the first studies applying a cognitive experimental approach to the study of financial diagnosis, was the study of McGee, Shields and Birnberg (1978). They tested the effects of personality differences on cue usage and decision outcome. Generally, it was found that personality variables did not appear to be useful in describing, understanding or predicting human information processing behaviour. Even though McGee et al.'s personality variables had some weaknesses, their experimental design was typical for this approach. Soon after, Moriarty (1979) studied the effect of different *information forms*. He found that subjects receiving cues represented by Chernoff faces outperformed subjects receiving quantitative information. The independent variables of these two first studies have later received little attention in financial diagnosis research.

An independent variable receiving considerable interest, has been *information load* . It has been investigated by several authors, and also closely relates to the *amount of information used.* First, Casey (1980b) found increasing accuracy with an increase in information load up to a certain level. With a greater information load, accuracy did not improve. The subjects with the greatest information load showed no superior accuracy, but used significantly more time to perform the task. In this study, information load was manipulated by introducing notes to the financial statement information. Generally, an inverted U-shape accuracy effect of information load was hypothesised. However, the information overload necessary to create reduced accuracy may not have been obtained in Casey's (1980b) study. Information overload was found by Chewning and Harrell (1990) in a similar task to have a negative effect on outcome consistency, agreement and consensus. The level at which information overload occurred was measured by finding the information load giving fewer significant coefficients in a linear model of the subjects' judgements. This measure of information load is not without difficulties, and the study represents one of the later examples of how cue usage measures of the lens model are still used in cognitive research in accounting[1].

Iselin (1993) refined the information overload concept by separating the concepts "information load" and "data load". His findings supported the hypothesis that accuracy was reduced at a high level of data load, often characterised as information overload. The data load supposed to induce information overload was very different in the studies referred to above. In the Chewning and Harrell (1990) study, eight cues were assumed to induce information overload, while 57 cues were assumed to be necessary in the Iselin (1993) study, and 15 ratios plus full income and balance statements were assumed necessary in the Casey (1980b) study.

[1] Recall how linear models assume additive effects of the cues, while the study actually search for nonlinear effects of cues or cue values.

Few have studied effects of different *information contents*. One exception is the study of Enis (1988), testing the effect of using current valued cues. He found that only sophisticated investors, utilising a small amount of current valued information, improved their accuracy and consensus. The overall effects of the current valued data were negative. The effects were very similar for accuracy and consensus, supporting the hypothesis that consensus measures can substitute measures of accuracy. Enis (1988) also used cue usage measures from the judgement modelling literature to control for the effect of information load.

Differences between data driven and conceptually driven subjects in a loan decision task were investigated by Rodgers (1992). The data driven subjects outperformed the conceptually driven, but the difference in accuracy was not considerable[1].

Different *search strategies* may alter the way sequentially ordered information affects diagnosis. Ricchiute (1992) tested the effect of presentation order, and found that information presented in a causal order of relevance to the going concern task led to better decisions. First, this result did not suggest that better diagnosticians reorganised the information in causal order, but simply indicated that the casual order provided a better diagnosis. Second, it raised the question if presentation order interacted with representations so that, for example, subjects with a "schematic organisation" of memory would benefit from the causal order. This question could not be answered in Ricchiute's (1992) study, because only subjects hypothesised to have a schematic organisation of memory participated[2]. The effect of introducing subsequent information is closely related to information order. Danos, Holt and Imhoff (1989) found that loan officers reached a high level of confidence early in the loan decision task, but despite the early confidence, they adjusted it in the correct direction when subsequent information was presented. A similar result was found by Barnes and Huan (1993), in a going concern judgement task. The subjects agreed on cue usage and classification of going concern status, and adjusted their decisions with mitigation information. However, in both studies, the subjects adjusted their decision more in the direction of the low classification error cost alternative[3]. The findings on agreement and confidence in the initial part of the task in these studies, suggest that this part of the task is *cognitively separable* , and that high accuracy and consensus on this task is achievable.

The effect of introducing subsequent information is also closely related to the *reasoning strategy* followed by the subjects. Following a confirmatory strategy could lead to less attention to subsequent disconfirming information. A search for the use of confirmatory

---

[1] See also the discussion of other aspects of Rodgers' (1992) study in section 2.1.

[2] The subjects were 100 partners.

[3] Lowest classification error cost is supposed for not granting a loan and for not qualifying in the going concern judgement.

strategies was done by Kida (1984), in a study of going concern judgements. He found only weak support for the hypothesis that a confirmatory strategy was followed in this task, and even less support for the hypothesis that this would lead to confirmatory bias. Kida (1984) suggested that this may be explained by the non-sequential character of the task. Trotman and Sng (1989) replicated and extended the experiment by Kida (1984) in a failure prediction context, and found further support for Kida's conclusion. A confirmatory bias in the direction of failure in general was found, but this bias could not be manipulated by different framing of the initial hypotheses, and was characterised as a general "failure bias". Information on cue diagnosticity was found to reduce this general bias. The bias could be explained by subjects' consideration of misclassification cost and failure to conform to Bayes theorem. The last proposition was tested by Holt and Morrow (1992) in a risk assessment task context. They compared the ability of lenders and auditors to conform to Bayes theorem, but found no difference between the two groups. However, auditors learned to conform to the theorem with experience, but lenders did not. This finding was used to support the hypothesis that "there are more incentives for avoiding risk in the bank lending environment than in the auditing environment" (Holt & Morrow, 1992, p. 549). The finding could alternatively be explained by subjects' consideration of misclassification cost as causing the "failure bias" in these environments.

In other cognitive accounting tasks, *representation*, *memory organisation* , and *knowledge* differences have recently received considerable attention (e.g. Bonner, 1990; Brown & Solomon, 1991; Frederick, 1991). In the financial diagnosis task contexts, a similar attention has, unfortunately, not been found. Choo and Trotman (1991) investigated the knowledge representations of experienced and inexperienced subjects in a going concern judgement and recall experiment. Based upon schema theory, they hypothesised that experienced subjects should recall more atypical than typical cues, because the atypical cues were in conflict with the proposed schematic memory. They also hypothesised that this would affect judgements, leading to better judgements for the subjects with the proposed schematic organisation of memory. Both these hypotheses were supported. The relationship between a schematic organisation of memory and the recall of atypical items was *hypothesised*[1] in this study, and was based on a "schema-plus-tag" relationship (Graesser, 1981), with validity only to immediate tests of memory (Ellis & Hunt, 1993, p. 248).

Several studies of knowledge representation with an experimental cognitive approach have used *experience* as an independent variable. Experience has been considered an operationalisation of several concepts, such as, knowledge differences, level of expertise and professionality, and, as seen above, the presence of a schematic organisation of memory.

[1] A schematic organisation of the memory of the experienced subjects was also hypothesised, not tested.

Experience effects have been summarised across cognitive accounting tasks in several reviews (e.g. Bedard, 1989; Bedard & Chi, 1993; Bonner & Pennington, 1991; Choo, 1989). Following the findings of Bonner (1990), that experience effects are highly task dependent, we concentrate on studies focusing on the financial diagnosis task.

The hypothesis that experienced subjects have a schematic organisation of memory was investigated by Choo and Trotman (1991). Their study is representative of the research strategy applied to the test of experience effects and knowledge differences within the experimental cognitive approach. These studies use information processing theory to postulate that the experienced have some knowledge or knowledge organisation, not shared by the inexperienced. If this is the case, they further hypothesise the effect of that particular knowledge organisation on behaviour. Next, the behaviour is observed, and effects are explained. However, there are two links in this research strategy. The first is the link between experience and knowledge or knowledge organisation. The second link is between knowledge or knowledge organisation and behaviour. Because knowledge organisation is not directly observable, the first link must be firmly established in theory or in other empirical research. Unfortunately, the first link is not obvious (see e.g. Ellis and Hunt, 1993, pp. 246-247), and also suffers from the use of experience as operationalisation of a wide variety of theoretical constructs in information processing theory. Another example of the same strategy is found in the study of bond raters by Danos et al. (1984). They proposed that the "bond raters' training and review process, coupled with their repeated exposure to forecasts, foster the development of sharply defined schemata" (Danos et al., 1984, p. 549). Next, they proposed that these schemata "can facilitate the recognition and use of disconfirming evidence. Therefore, we posited that bond raters would overcome the common response bias of ignoring disconfirming evidence" (Danos et al., 1984, p. 550). The last of these hypotheses was tested and confirmed without the presumption of a schematic organisation of memory as intervening variable[1] (Barnes & Huan, 1993). Not surprisingly, the same hypothesis was confirmed by Danos et al. (1984).

Consistent with general findings in cognitive psychology literature on expert behaviour (see e.g. Ellis & Hunt, 1993, p. 283), Libby and Frederick (1990) found that the experienced subjects generated more accurate explanatory hypotheses, had more accurate knowledge of error occurrence rates, selected more commonly occurring explanations, and categorised their knowledge differently than the inexperienced subjects. However, they are careful in their use of intervening variables as explanations of these findings. Experienced subjects' knowledge of frequencies and the effects of change in this knowledge was investigated by Nelson (1993), in a simplified analytical review task. He found that by introducing a distracting task related to

[1] The authors have later replicated their study without use of the intervening variable, suggesting they agree with our proposition (Danos, Holt & Imhoff, 1989).

error frequencies, subjects learned these error frequencies. As expected, this learning affected the error frequencies of novice subjects in the analytical review task, but the experienced subjects did not alter their knowledge of error frequencies as a consequence of the distracting task.

Several reviews of the differences between expertise behaviour in the judgement modelling and cognitive experimental literature are found (e.g. Bedard, 1989; Cho, 1989). Bonner and Pennington (1991) suggested that the "mixed findings" of the two approaches could be task related. In their review of accounting tasks and expertise performance, they characterised the findings regarding expertise performance as "mixed" in all financial diagnosis task contexts, while many of the findings in other tasks, such as for example, internal control evaluation, were much more consistent and indicated superior expertise performance. These findings suggest that the financial diagnosis task may be a task where experience effects are not very evident.

Accuracy and *consensus* are the two aspects of judgement performance most often studied. In cognitive experimental studies, these aspects are studied as dependent variables[1]. Other aspects of the outcome have also been investigated. The subjects' *confidence* in their own decisions have been used (Danos et al., 1984, Danos et al., 1989). Cue *utilisation* has also been used in some studies (e.g. Chewning & Harrell, 1990), relying on lens model measures of cue usage. Investigation of the *properties* of these dependent variables is mainly performed in judgement modelling research, and in studies within the descriptive cognitive approach.

The descriptive cognitive approach began with studies applying protocol analysis methodology (Ericsson & Simon, 1984) to financial diagnosis tasks. Examples of essential research questions within the approach are illustrated in table 2.2. Some of the most relevant studies within this approach to financial diagnosis are listed in chronological order in table 2.5. In addition to an indication of the task context used in these studies, a column indicating the main focus of the descriptive studies is provided.

In descriptive cognitive studies, the task context is traditionally more realistic than in both judgement modelling and experimental cognitive studies. This realism is evident in the stimuli presented, the context of the tasks, and the responses expected from subjects performing the task. However, the number of subjects and cases are often limited, and internal validity is often focused at the expense of external validity. Despite these limitations,

[1] See the column indicating the dependent variables of table 2.4.

the descriptive orientation contains some of the most important studies of information processing behaviour in the financial diagnosis task.

All *tasks* studied in this approach require task specific knowledge for their solution. Newell and Simon (1972) assume that the greatest variation in problem solving behaviour is explained by the task structure. Their suggestion is that the less structured a task is, the larger differences in individual information processing behaviour can be observed.

Reference	Task context	Descriptive focus
Bouwman, 1982	Evaluate position of firm	Expert/novice differences
Bouwman, 1983	Evaluate position of firm	Comprehensive descriptive model
Biggs, 1984	Earnings forecasting	Information search and usage
Campbell, 1984	Loan decision	The use of four financial statement items
Biggs et al., 1985	Loan decision	Effects of task size and similarity of alternatives
Methlie 1993, 1994	Loan decision	Knowledge representation and computational model
Bouwman et al., 1987	Investment screening	Comprehensive descriptive model
Anderson, 1988	Prospectus evaluation	Expert/novice differences
Blocher & Cooper, 1988	Analytical review task[1]	Individual problem solving behaviour differences
Bedard & Biggs, 1991	Analytical review task[2]	Search for pattern recognition and hypothesis formation behaviour
Biggs et al., 1993	Going-concern judgement	Knowledge representation and computational model

Table 2.5 Selected descriptive cognitive studies of financial diagnosis

The effects of task size and similarity within the same task context was investigated by Biggs, Bedard, Gaber and Linsmeier (1985). In accordance with an information economics approach, loan officers used non-compensatory decision strategies when the size[3] of the task was large, and when similarity[4] was high. In this study only pre-evaluated qualitative information was used, and thus, generalisation was somewhat limited.

[1] The inventory account evaluation task studied by Blocher and Cooper (1988) is included in this review because they used a comprehensive stimulus material including large amounts of financial statement cues relevant to financial diagnosis in general.

[2] The inventory account evaluation task studied by Bedard and Biggs (1991) is included in this review because they used quantitative financial statement cues relevant to financial diagnosis as stimulus material.

[3] Task size was operationalised as the number of cues presented to subjects.

[4] Task similarity was operationalised as similarity in the cue value of two stimulus patterns.

A descriptive model of the financial analysis task was developed by Bouwman (1982, 1983). His sequential model divided the problem solving process into phases of problem detection, integration of findings, knowledge updating and final diagnosis. This model was refined in Bouwman et al. (1987), and the problem solving process was now divided into two main phases of familiarisation and analysis. In the familiarisation phase an information search was performed, and the *search strategy* of the analyst guided this search. Of the search strategies available, sequential search was most typical (Bouwman, 1983), but when explanation and analysis of findings were needed, the experts relied on a directed search (Bouwman, 1983). The sequential search was then used as a "safeguard" (Bouwman et al., 1987). These findings were confirmed by Anderson (1988). Another guiding instrument of the information search, especially among the experts, were "checklists" (Bouwman et al., 1987). These "checklists" were knowledge guiding both search and reasoning, and was used during the familiarising sequential search to detect interesting findings.

Biggs (1984) concentrated on the *type of information* used by experienced financial analysts in an earnings power assessment task. The report most widely searched was the income statement, with a relative percentage of 67.3 %. This was also confirmed in Anderson's (1988) study. Another finding reported by Biggs (1984), was the similarity of search behaviour among the analysts. He found that operating performance indicators and trend ratios were most often calculated. Blocher and Cooper (1988) also reported a wide range of ratio and trend information cues searched and used by auditors, but the most surprising observation was that the subjects did not use the most predictive ratios in this particular task. Bouwman et al. (1987) reported that 5 to 10 items in their comprehensive stimulus material represented 25 % of total cue usage. During familiarisation, the income statement was the most widely used report, and a ratio report not included in the earlier studies by Bouwman, was the second most used. However, during reasoning and analysis, non-financial items were most widely used.

Anderson (1988) expected professionals to search for smaller *amounts of information*, but could not confirm this hypothesis. However, the professionals *used* a smaller amount of information. Anderson (1988) also investigated the manipulation and weighting behaviour of the subjects, and found the subjects weighting negative cues more heavily than positive. This finding was contradictory to the finding reported in judgement modelling studies, that subjects overweight positive information (Slovic, Fischhoff & Lichtenstein, 1977).

In the Govindarajan (1980) content analysis, the use of earnings versus cash flow information was compared. The study concluded that earnings information was considered more important than cash flow information. Campbell (1984) tested loan officers' usefulness of four cues of "Big GAAP[1]" financial statements; earnings per share, deferred taxes, leases and inflation adjusted information. She found that loan officers did not utilise this information when making a loan decision.

Reasoning processes are seldom reported explicitly in financial diagnosis studies. Bouwman (1983) documented the subjects' use of *qualitative reasoning*. Quantitative information was transformed to a qualitative form to characterise both the level and trend of financial cues. This finding was common to all of Bouwman's (1982, 1983) subjects. The qualitative characterisation of financial information was the result of the familiarisation or examination phase. Only qualitative information that deviated considerably from what was expected, was remembered and used in the analysis phase.

The formulation of *hypotheses* to guide the reasoning, was typically found among experienced subjects both in financial analysis (Bouwman, 1982) and in auditing (Blocher & Cooper, 1988). This was not found among the inexperienced subjects, where a sequential attention to observations dominated the reasoning process (Bouwman, 1982). The reasoning process of these subjects was characterised as data driven and forward chained. Not formulating the relevant hypothesis was found by Bedard and Biggs (1991) to be the main reason of subjects' error in an analytical review task.

Most studies suggest some kind of schematic *representational form* of knowledge. In the early study of Bouwman (1982), the schematic structure was termed a "checklist", and this representational form was considered unique to the experts in the study. An example of such a "checklist" could be a list of common problems which transforms diagnostic *reasoning* into diagnostic *recognition* (Bouwman, 1984, p. 327). In Danos et al. (1984), differences between experts and non-experts were hypothesised to be explained by "common and sharply defined knowledge structures or schemata" (Danos et al., 1984 , p. 563). Later, Biggs and Wild (1985) suggested a representational form that could guide the recognition of known patterns in a time series. In Meservy et al. (1986), "frames of reference" was suggested as an alternative to production rules. In Biggs et al. (1987) "ad hoc schemata" were suggested as the representational form. These suggestions were supported by the finding of Bouwman et al. (1987), that "financial templates" was the most important representational form. These "financial templates were proposed to be "complex structures that contain a variety of knowledge: industry-specific standards of what is acceptable, "pictures" of typical company behaviour, typical problems for that kind of company, or industry, and "ready-made" evaluations of the attractiveness as an investment" (Bouwman et al., 1987, p. 26). The "financial templates" were activated early in the problem solving process, and guided

[1] Generally Accepted Accounting Principles.

reasoning. The "templates" also functioned as "recognition devices" during the familiarisation phase of the problem solving task, and reduced the analysis phase to one of synthesising recognised patterns. Bouwman et al. (1987) states:

> *It replaces a reasoning process by a recognition process, which is much faster and requires less effort. (Bouwman et al., 1987, p. 22)*

Pattern recognition was specifically searched by Bedard and Biggs (1991) in an analytical review task including only quantitative financial statement cues. They reported that the best performing subjects showed "pattern recognition abilities". Bad performance results were explained by three errors; acquisition error, pattern recognition error and hypothesis generation error. Thus, the Bedard and Biggs' (1991) model assumed that subjects first identified relevant discrepancies. These were recognised as a pattern, and the pattern was linked to a formulated hypothesis of the pattern being the cause of the error. Bedard and Biggs showed that 14 of the 21 subjects recognised the pattern, but did not formulate the hypothesis to link it to. They used this finding to suggest that many of the subjects had pattern recognition abilities, but could not connect the recognised pattern to the correct hypothesis. However, Bedard and Biggs did not investigate the representational form necessary for the pattern recognition to take place.

The importance of pattern recognition during the familiarisation phase of financial diagnosis was also stressed by Biggs et al. (1993[1]), in their attempt to build a computational model of the going concern decision of auditors. They proposed two initial phases very similar to the familiarisation and analysis phases referred to above. The primary purpose of the first phase was problem recognition. Several categories of *knowledge* were assumed to be necessary to perform problem recognition. Procedural knowledge guided the reasoning process. Financial knowledge was both specific and general, and was used to recognise patterns in the financial information indicating going concern problems. Event knowledge was case specific, and was used to link detected problems to causes. A prepositional network representation was proposed for most of the financial and event knowledge. However, the problem recognition itself was implemented by applying standard rule-based tests of the financial cues against level standards, or by comparing cues over a sequence of consecutive years.

One of the few descriptive cognitive studies *formalising* the financial knowledge presumed used by financial diagnosticians, and suggesting a method that could be used to turn verbal protocol data into representations of such knowledge, was a study by Methlie (1993, 1994). The main contribution in this study was the development of a method termed conceptual

---

[1] See also Selfridge, Biggs & Krupka, 1992.

analysis, that was used to identify and structure the concepts and attributes of concepts used by financial diagnosticians in a hierarchical conceptual structure. When conceptual analysis was applied to verbal protocols of an expert performing credit evaluation, two important findings were made. First, attributes of a concept were of two types. Evaluative attributes were evaluated against an internal standard, while comparative attributes were evaluated by comparisons of two or more attributes. An example of an evaluative attribute is the level of a financial ratio, while an example of a comparative attribute is the trend in two or more values of a financial ratio. This finding emphasises the importance of the two concepts "level" and "trend", both presumed functional in the evaluation of the same type of financial statement information. The second important contribution of Methlie's (1993, 1994) study was the structure of the conceptual hierarchy of the financial diagnosis expert. The hierarchy consisted of four major diagnostic concepts or areas representing "profitability", "financial structure", "financial leverage" and "liquidity". The expert's opinion on these diagnostic areas was formed by evaluative and comparative investigation of a selected set of financial ratios, and the diagnostic area opinions were merged in a main financial diagnosis. As expected, the diagnostic areas found by Methlie (1993, 1994) corresponded well to theoretical concepts identified in predictive studies, and were believed to be diagnostic of a firm's financial situation.

Tests of response *accuracy* are limited to situations where comparisons with actual results are possible, and consequently, not often focused in descriptive studies. One exception was a study by Anderson (1988), who tested the accuracy of professionals and non-professionals. He found greater accuracy among professionals in a security pricing task than among non-professional subjects. Surprisingly low accuracy was found by Blocher and Cooper (1988), in an analytical review task. Low accuracy was also found by Bedard and Biggs (1991), in their fairly similar analytical review task.

2.3 The predictive approach

The predictive approach to financial diagnosis consists of hundreds of published studies of the task contexts listed in table 2.1. Foster's (1986) empirically oriented textbook refers to five areas of research were financial diagnosis is involved; in the evaluation of securities, in corporate restructuring, in debt rating, in distress prediction, and in the loan decision. Analysts performing the security evaluation task use capital market information and other information collected from sources other than the financial statement. The corporate restructuring task context is not unambiguously related to financial diagnosis, and an acquisition can be motivated both in financially relevant and financially less relevant factors. Consequently, we focus on some of the most well known contributions primarily on the three task contexts; bankruptcy or distress prediction, bond or debt rating, and loan decisions. More comprehensive reviews of the predictive approach are found for bankruptcy prediction studies in Jones (1987) or Altman (1983), for bond rating studies in Belkaoui (1983), and for loan decision studies in Rosenberg and Gleit (1994). A comprehensive review of the contributions in all task contexts up to 1980 is found in Altman et al. (1981).

Providing a comprehensive review of this research is far beyond the scope of this section. However, referring to the lens model of figure 2.1, the predictive approach is of relevance for two reasons. First, the relevant cues of a financial diagnosis should correspond to relevant cues of financial prediction in the task contexts referred to above[1]. Next, an economic theory of financial diagnosis should underlie the left hand side of the model, and may be implemented as part of the knowledge represented in a cognitive model of financial diagnosis. Consequently, these two aspects of predictive studies are of particular relevance. After a brief review of some of the more relevant contributions within the predictive approach, we concentrate on the treatment of these two aspects.

A selection of some of the more important studies with a predictive approach is shown in table 2.6. The studies are grouped, and ordered chronologically within each task context. In addition, two columns indicating the methods applied in the predictive models and the main focus of each study, are provided.

From the selection in table 2.6, we see that the contributions to the three task contexts have had a very similar development. Originally, simple models were applied to the task, later followed by more sophisticated models. For example, discriminant analysis applications were typically followed by quadratic discriminant analysis, logit and probit analysis, and last, nonparametric techniques, such as recursive partitioning, have been applied.

Methodologically, these studies have an orientation quite different from traditional theory driven research. Most of the studies are data and method driven, starting with a potential set of predictive cues. Next, the most predictive cues are selected based upon prediction or classification accuracy in a sample. Validation of the model is traditionally performed on a holdout sample, or with forms of cross validation (Lachenbruch & Mickey, 1968; Stone, 1974).

[1] At least for skilled financial diagnosticians.

Reference	Task context	Method	Focus
Beaver, 1966	Bankruptcy prediction	Univariate analysis	Distributions of failed and other
Altman, 1968	Bankruptcy prediction	Discriminant analysis	Development of Z-model
Wilcox, 1976	Bankruptcy prediction	Conceptual model	Theoretical model
Altman, Haldeman & Narayana, 1977	Bankruptcy prediction	Discriminant analysis	Improvement of Z-model
Ohlson, 1980	Bankruptcy prediction	Logistic regression	More realistic assumptions
Mensah, 1983	Bankruptcy prediction	Discriminant analysis	Price level adjustment
Frydman et al., 1985	Bankruptcy prediction	Recursive partitioning	Nonparametric model
Gentry, Newbold & Whitford, 1985	Bankruptcy prediction	Logit analysis	Cash flow data
Zavgren, 1985	Bankruptcy prediction	Logistic regression	Economic interpretation
Karels & Prakash, 1987	Bankruptcy prediction	Discriminant analysis	Test for normality assumptions
Zavgren & Friedman, 1988	Bankruptcy prediction	Logistic regression	Underlying derived factors
Gilbert, Menon & Schwartz, 1990	Bankruptcy prediction	Logistic regression	Only stressed cases
Laitinen, 1991	Bankruptcy prediction	Discriminant analysis	Conceptual model
Falbo, 1991	Bankruptcy prediction	Discriminant analysis	Level, trend and stability
Molinero & Ezzamel, 1991	Bankruptcy prediction	Multidimensional scaling	Nonparametric groupings
Horrigan, 1966	Bond rating	Regression analysis	One function for each class
Pinches & Mingo, 1973	Bond rating	Discriminant analysis	Classification with 5 classes
Kaplan & Urwitz, 1979	Bond rating	Probit analysis	Relating linear score to rating
Bhandari, Soldofsky & Boe, 1979	Bond rating	Discriminant analysis	Rating changes
Reiter & Emery, 1991	Bond rating	Conjoint analysis	Methodological aspects
Ziebart & Reiter, 1992	Bond rating	Conceptual model	Theoretical model test
Buta, 1994	Bond rating	Recursive partitioning	Rules induced
Mutchler, 1985	Going-concern judgement	Discriminant analysis	Introducing mitigation factors
Koh & Killough, 1990	Going-concern judgement	Discriminant analysis	Client focus
Hansen et al., 1993	Going-concern judgement	Recursive partitioning	Comparison with traditional models

cont...

cont...

Reference	Task context	Method	Focus
Orgler, 1970	Loan decision	Regression analysis	Dummy variables
Edmister, 1971	Loan decision	Discriminant analysis	Dummy variables
Chesser, 1974	Loan decision	Logistic regression	Non-compliance
Doukas, 1986	Loan decision	Discriminant analysis	Comparison with behaviour
Srinvasan & Kim, 1987	Loan decision	Recursive partitioning	Behavioural dependent variable
Shaw & Gentry, 1990	Loan decision	Recursive partitioning	Rules induced

Table 2.6 Selected predictive studies of financial diagnosis

In *bankruptcy prediction*, Beaver (1966) constrained the set of potentially relevant cues by an investigation of the univariate predictive accuracy of 30 cues covering the concepts of "cash-flow", "profitability", "leverage", "liquidity[1]" and "turnover". He found that the ratio cash flow/total debt classified 87 % of the holdout cases correctly, one year prior to failure. Altman (1968) used discriminant analysis to develop his first Z-score model using five ratios: Working capital /total assets, retained earnings/total assets, earnings before interest and taxes/total assets, market value of equity/book value of debts, and sales/total assets. The model predicted 83.5 % of the holdout cases correctly.

Later Altman et al. (1977) improved the original Z-score model by including other variables, now intending to represent "profitability", "earnings stability", "interest coverage", "cumulative profitability", "liquidity", "capitalisation", and "size". In addition, this model incorporated considerations of prior probabilities, unequal covariance matrices, and costs of misclassification. New developments in classification methods have continuously been adopted by bankruptcy prediction researchers, and implemented in their models. Ohlson (1980) used logistic regression to overcome some of the methodological critiques of discriminant analysis. Nonparametric methods, such as recursive partitioning (Frydman et al., 1985) and multidimensional scaling (Molinero & Ezzamel, 1991) have been applied with some success.

New variables have been included in bankruptcy prediction models both with and without a firm theoretical foundation. Mensah (1983) found that *price level* adjusted ratios could improve a logit model when misclassification costs were considered. Gentry et al. (1985) found that the use of *cash flow* based variables had only limited success in improving current prediction models. Karels and Prakash (1987) suggested tests of *multinormality* should guide

[1] Both liquidity related to total assets and to current debt.

the selection of independent variables. The studies referred to above were all representative examples of the different ways to improve the initial prediction models by Beaver (1966) and Altman (1968)[1]. Recently, several authors have questioned the ability of bankruptcy models to separate failure from non-failure in a set of *stressed* firms. Both Gilbert et al. (1990) and Hopwood et al. (1994) found rather unimpressive results when this alteration of the task was performed.

Karels and Prakash (1987) were among the first to *compare* the different cues suggested in the bankruptcy prediction literature. In their summary (Karels & Prakash, 1987, p. 578), the six ratios found significant in more than one bankruptcy prediction study were; cash flow/total debt, quick assets/total assets, current assets/total assets, net income/total assets, operating income/total assets and total debt/total assets. Two of the ratios represent "profitability", two represent "liquidity", one represents "leverage" and one represents "debt coverage". Aspects other than the *level* of a ratio have also been considered. *Stability* of the ratios was introduced as an important independent variable in the Z-model of Altman et al. (1977). Explicit consideration of the *trend* of the ratios is less evident, but exceptions have been found (Falbo, 1991).

Studies founding their selection of independent variables and model implementations on a conceptual basis, have followed two approaches. The model can be based upon the "principle independent *factors* of the financial statement" (e.g. Zavgren & Friedman, 1988 p. 35), or upon some *theory* of the process of bankruptcy (e.g. Wilcox, 1973). The first of these approaches is based upon empirical research on the pattern of financial statement data (e.g. Gombola & Ketz, 1983; Pinches, Mingo & Caruthers, 1973). Zavgren and Friedman (1988) applied the seven *factors* of Pinches et al. (1973): "Inventory turnover", "receivables turnover", "cash position", "short-term liquidity", "return on investment", "financial leverage" and "capital turnover" to select relevant cues. However, the way these factors were related to bankruptcy remained largely unexplained. A similar approach was also applied in the judgement modelling approach, where e.g. Libby (1975) used factor analysis to guide the selection of cues from Deakin's (1972) 14-variables set. Even if the approach is built on patterns of variation, it is not obvious that these patterns of variation are of relevance to financial diagnosis. This general problem of factor analysis was stressed by Jones and Sibson (1987):

[1] 8 strategies for improvement are listed in Laitinen, 1991.

> *It (principal components analysis) relies for its success on the tendency for large variation also to be the interestingly structured variation, a connexion which is not logically necessary, and often fails to hold in practice,....(Jones and Sibson, 1987, p. 2)*

Studies applying a *theory -based* conceptual model in their selection of cues and model implementation are few. Wilcox (1976) used a gambler's ruin approach to develop a theoretical basis for the indication of financial risk, and found performance results of his model comparable to statistically derived models. A different conceptual basis for the prediction of failure was given by Argenti (1976), suggesting different failure processes. He suggested three failure processes were relevant: Some firms never rose above poor performance. Other firms experienced growth and rapid decline. The third category consisted of firms experiencing good and stable performance, with a sudden partial collapse. The partial collapse further initiated rapid decline. Laitinen (1991) combined the two approaches referred to above to build a conceptual model of the failure process. He discussed the relevance of traditional factors found in financial statement cues, and added a *growth* factor. Next, he used factor analysis to detect patterns of variation over *time in failed firms*. The results showed three failure processes represented by "the chronic failure firm", "the revenue financing firm" and "the acute failure firm". Unfortunately, his model was not validated on a holdout sample. Despite these efforts, the status of theoretical models in bankruptcy prediction can still be characterised by Jones' (1987) description[1]: "Theory has played a limited role in guiding empirical research projects" (Jones, 1987, p. 136)[2].

The going-concern judgement task context is very similar to the bankruptcy prediction context, but has not received similar interest in the predictive literature. One reason may be that the auditors' decision to qualify for a going concern problem includes an evaluation of contrary and mitigating information, after problem identification has been done. Mutchler (1985) studied whether auditors decision to qualify for going-concern problems could be predicted with publicly available information. She interviewed auditors to guide the selection of relevant financial statement cues, and ended up using the following ratios: Cash flow/total liabilities, current assets/current liabilities, net worth/total liabilities, long-term liabilities/total assets, total liabilities/total assets and net income before tax/net sales. In addition, she tested the effect of news, positive trends, and prior qualification on the linear model's ability to predict qualification from cases already identified as having going-concern problems. The best model included prior qualification, and predicted 90 % of the cases correctly. Using a

---

[1] See also Foster (1986) p. 559-560.

[2] An example of how behavioural theory can be used in selection of diagnostic cues in financial diagnosis tasks is by reviewing analyst reports to find signals and fundamentals reported relevant by these analysts. This approach has recently been used by Lev and Thiagarajan (1993), but their model was used to predict excess return.

subset of the Mutchler-data, Hansen et al. (1993) compared the predictive abilities of logistic regression to recursive partitioning on the same task. The prior qualification variable was left out, giving predictive accuracy of 82 % for the best model. Recursive partitioning did not perform better on this sample.

In the *bond rating* context, the dependent variable has a different scale than in the previous task contexts[1]. Horrigan (1966) developed a bond rating model using regression analysis, assuming that the dependent variable was interval scaled. Horrigan (1966) used primarily financial statement cues in the model, and was criticised by West (1970), who used cues applied from Fisher's (1959) model of market risk premiums. The methodological development of bond rating studies is very similar to how bankruptcy prediction models have developed. Pinches and Mingo (1973) applied factor analysis to select cues representing interesting dimensions, and used these cues in a 5-group linear discriminant model. Altman and Katz (1976) used quadratic discriminant analysis with 14 cues. The most significant where: Interest coverage, variability of interest coverage, variability of operating income, return on assets, total assets, market value of equity/book value of equity, and retained earnings/total assets. Kaplan and Urwitz (1979) used probit analysis including seven cues in their best model. The following cues proved most significant: Subordination, total assets, net income/total assets, and long term debt/net worth. Variability measures and beta did not improve predictions significantly. Neither did the probit analysis model, when compared to traditional linear regression. A table of the predictive accuracy of these models can be found in Altman et al., (1981), showing that classification accuracy varied from 56 to 77 % . These results should be interpreted with caution, since the number of classes of the dependent variable varied across studies. Nonparametric methods have also been applied to the bond rating task context. Reiter and Emery (1991) compared traditional OLS regression, probit analysis, and conjoint analysis. They did not find statistical support for preferring one model to another. One explanation may be that to apply conjoint analysis, the independent variables must be binary encoded, and a loss of information will occur[2]. Reiter and Emery (1991) also tested the effect of including several non-financial cues in the model, and found that some of these significantly improved accuracy. However, some of these variables (e.g. trouble with nuclear plant), seemed somewhat speculative, and were not selected on the basis of a theoretically derived hypothesis. Recursive partitioning has also been applied to bond rating (e.g. Buta, 1994).

Consequently, results from the bond rating studies generally show somewhat lower classification accuracy than bankruptcy prediction studies. However, it is difficult to establish if this difference is due to the differences in classification categories, or to a difference in the

[1] Variables measuring the rating of bonds are ordinal scaled.
[2] Binary transformations suggested by Edmister (1971) were used.

predictability of the tasks. Cues used to predict bonds have many similarities to successful cues used in bankruptcy prediction, but variability and coverage measures[1] are more frequently applied in bond rating studies.

The status of theory in bond rating models is very similar to its status in bankruptcy prediction. Emery and Reiter (1991) state that:

> *For the most part, the models used in past studies have been chosen on the basis of statements by the rating agencies about factors considered in rating a bond or other empirical techniques, such as factor analysis and stepwise regression. (Emery & Reiter, 1991, p. 149)*

The factors used by rating agencies are generally not official information. However, some agencies release information on distribution differences in particular cues for the different rating classes. Foster (1986) listed provisions on indenture agreements, protection afforded by existing assets and quality of management as factors, other than those traditionally measured by financial statement cues, of relevance to rating agencies

The nature of the bond rating task context makes theoretical development of a model less likely to occur. One exception is found in the study of Ziebart and Reiter (1992), relating financial variables to both bond rating and bond yields. They applied a causal modelling approach, and found that financial statement cues affected both bond ratings and bond yield directly. In addition, financial information indirectly affected bond yield through bond ratings. However, no specific conceptual model was developed for the direct effect of financial statement cues on bond ratings. The financial information used in this study was beta, total assets, interest coverage, net income/total assets and long term debt/market value of equity. In a stock return prediction task, Ziebart (1987) used four latent variables measured by 13 ratios. The latent variables were assumed to represent "liquidity", "leverage", "profitability" and "activity". No theoretical basis for the selection of latent variables was provided, and the "activity" concept was measured by turnover ratios.

In addition to bond rating, *bond rating changes* have been studied in predictive models. An early example was when Bhandari et al. (1979) attempted to predict rating changes of industrial bonds with a linear discriminant model.

Predictive studies of the *loan decision* have focused either on the credit granting decision or on the loan review process (Rosenberg & Gleit, 1994). Both task contexts involve financial

---

[1] Particularly, interest coverage measures.

diagnosis, but decision outcome depends more on financial diagnosis in the first. One reason is that loan history may be more relevant in loan review than in credit granting. Among the first developments was the loan review model of Orgler (1970), used to classify good and bad loans. This form of categorising the dependent variable is common within the loan decision context, and makes the results of these models comparable to bankruptcy prediction models. In fact, the failure concept is operationalised both as bankruptcy and as loan default[1]. Edmister (1971) developed a model to predict loan default in small businesses by using binary transformations of the variables equity/sales, net working capital/sales, current liabilities/equity, funds flow/current liabilities, inventory /sales trend, and quick ratio in addition to an industry related quick ratio trend. Doukas (1986) compared loan classifications with predictions from simple bankruptcy prediction models. Not surprisingly, it was found that the Altman (1968) model was a poor predictor of loan classifications as performed by Canadian bankers.

The methodological development of the loan decision studies follows a pattern similar to the bankruptcy prediction studies. Simple regression and discriminant analysis approaches have been replaced by methods with more realistic assumptions of the independent and dependent variables. Srinvasan and Kim (1987) were the first to compare parametric and nonparametric methods to the credit granting decision. They used customer files to classify customers into high risk and low risk categories. Current ratio, quick ratio, net worth/total debt, logarithm of total assets, net income/sales and net income/total assets were used to classify the customer into one of these two categories. Their models showed very high classification accuracy, with recursive partitioning performing best, classifying 92.5 % of the customers correctly in a cross validation procedure. These results were quite comparable to the best accuracy results reported in the bankruptcy prediction literature. A more traditional recursive partitioning application to the loan decision task was found in Shaw and Gentry (1990).

The position of theoretical models of the loan decision is very similar to the position of theory in the other task contexts reviewed above. A theoretical model of the process leading to loan default should prove very similar to bankruptcy process models, but few propositions have been made on such a theory. Some predictive studies of the loan review or credit granting process of loan officers are closely related to studies in the judgement modelling approach. Srinvasan and Kim (1987) used a behavioural model in their study, separating the assessment of default risk from other factors of relevance.

[1] For an overview of different operationalisations of "failure", see Karels and Prakash 1987, p. 576.

2.4 Other contributions of relevance to financial diagnosis

Most of the research questions listed in table 2.2, and reviewed within each approach to financial diagnosis have also been investigated in other disciplines. The literature on human information processing in accounting covers tasks such as internal control evaluations (Frederick, 1991), materiality judgements (Waller & Felix, 1989), analysis of performance reports (Shields, 1984), or intuitive forecasting (Danos & Imhoff, 1983), just to mention a few. Our review has been selective in only investigating studies which are of direct relevance to financial diagnosis as defined above. In this way, we have tried to avoid one of the main pitfalls pointed out by Gibbins and Jamal (1993) in *generalising findings from this literature across tasks*. However, in some of the studies of tasks not covered by our definition of financial diagnosis, particularly relevant research questions have been investigated.

In an auditing task, Biggs and Mock (1983) found two different overall problem solving approaches. For subjects following a *systemic strategy,* the data guided the problem solving behaviour in a sequential matter. In a *directed strategy* a subtask was defined to guide the search. Information necessary to perform the subtask was searched, and the subtask was completed before the next subtask was identified. Biggs and Mock (1983) found the amount of information attended to correlated with problem solving strategy. As expected, a systemic strategy resulted in subjects attending to more of the available information than with a directed strategy.

As opposed to studies of financial diagnosis task contexts, auditing studies with focus on *reasoning strategies* are more common. Two examples are Biggs et al. (1987) and Biggs and Mock (1983). In the Biggs et. al (1987) study, experienced auditors used two reasoning strategies in parallel. *Reasoning by assumption* meant that the auditor stated an assumption and let this assumption be an anchoring point for the reasoning to follow. *Reasoning by analogy* meant that the knowledge of previous and comparable situations guided the reasoning. Both these reasoning processes were heuristic. Reasoning by assumption may be considered a variation of the general anchoring strategy (Tversky & Kahneman, 1974), also found in Biggs and Mock (1983). Reasoning by analogy is closely related to pattern recognition.

In Biggs et al. (1987), the authors explicitly searched for evidence of *probabilistic reasoning* by the subjects. The lack of such findings were considered to support the hypothesis that some form of non-probabilistic reasoning about uncertainty, took place among the subjects.

Some of the studies outside the financial diagnosis task have focused on *representation* of knowledge of novice and expert subjects. The same representational forms as those assumed to guide judgement in financial diagnosis, have also been used within these studies. An example of a linear additive model of "knowledge representation" can be found in Peters (1990):

> *The model's knowledge base contains a series of cue weights, assigned by the author. The model selects relevant cues and combines their weights and assessed values using a linear combination rule (Peters, 1990, p. 89-90)*

The representation of knowledge in production rules is more common, and was found in several studies (e.g. Biggs et al., 1987; Meservy, Bailey and Johnson, 1986). The representation form of *production rules* was explicitly searched by Meservy et al. (1986), in an attempt to model the internal control evaluations of CPAs[1]. However, such rules were not explicitly mentioned by the CPAs, and had to be induced by task analysis. The same problems with identifying production rules were experienced in Biggs et al. (1987). The use of a schematic representational form was hypothesised by Meservy et al. (1986), but *schematic representational* forms were not used to represent knowledge in the model of the auditors.

Recently, several studies within the auditing discipline have focused on differences in representational forms among subjects. Frederick (1991) tested the assumption that experienced subjects had a more schematic organisation of memory, and found that auditors recalled more internal controls when these controls were organised by transaction flow. The student subjects showed no such difference. This finding was interpreted as support of Frederick's (1991) hypothesis of a schematic organisation of memory in experienced subjects. Another approach to investigate the schematic organisation of memory, is to search for *configural information processing*[2]. The lack of significant interaction effects found in judgement modelling and cognitive experimental studies of financial diagnosis have been used to suggest that configural information processing is not evident in financial diagnosis. Brown and Solomon (1991) studied whether configural information processing was used in tasks obviously requiring such information processing. They found support for subjects using configural information processing in situations where domain-specific knowledge implied that it was appropriate. The lack of similar findings in financial diagnosis may suggest that the financial diagnosis task does not require this kind of information processing. Another explanation may be that the configural information processing necessary in this task is not easily implemented in simple interaction terms of linear models.

---

[1] Certified Public Accountant

[2] Configural information processing is when the relevance of a cue is dependent on the context of other cues, and is often modelled by interaction terms in a linear model.

Sarah Bonner has studied the interaction between experience and attributes of the task, suggesting that experience differences will only be evident in tasks requiring task specific and not generic knowledge (Bonner, 1990, 1991). Using lens model measures of cue selection and cue usage, she found experience effects in an analytical risk assessment task both for cue selection and cue weighting. This line of research is reviewed in Bonner and Pennington (1991), summarising their conclusion in the following way:

> *In summary, there is a strong relation between the quality of performance in auditing tasks and the type of cognitive processes used in those tasks. Expert auditors perform better, on average, in tasks that require construction processes (information search, comprehension, hypothesis generation, and design) rather that reduction processes (hypothesis evaluation, estimation, choice). Performance level is also related to the presence of theory-based versus statistical reasoning and the quality of the knowledge available for processing. (Bonner & Pennington, 1991, p. 25-27)*

Relating these conclusions to financial diagnosis, the mixed results of experience effects may seem reasonable.

A final question related to financial diagnosis, but studied outside this task domain, is whether financial ratios traditionally used as cues in diagnosis are able to capture variations of significance in the underlying financial statement information. This question was investigated by Kinney (1987), in a study where he induced material accounting errors of different types in the underlying financial statement material. He studied whether these errors could be detected in financial ratios using simple detection rules. Two conclusions were made: First, even large material errors did not lead to change in ratios, when measured relative to their monthly natural variation. Second, by comparing the pattern of observed change in several ratios simultaneously, the particular type of error could very often be identified. These findings indicated that configural information processing may be relevant and necessary if financial ratios are to be used in financial diagnosis. However, the configural pattern was often found to be complex, leading Kinney (1987) to suggest that pattern search techniques could be useful.

2.5 Summary of findings on the financial diagnosis task

It is possible to unify some of the findings from the different approaches to the study of financial diagnosis reviewed above. However, conflicting results have also been identified in the studies referred to above. This section summarises some of the main findings that are supported by several approaches, and suggests that new explanations are necessary of some of the conflicting results found in the literature. The presentation is divided into five subsections. First, findings on the task itself are briefly summarised, followed by a summary of findings relating to search, reasoning and representation. Last, main findings on the response or outcome of the task are summarised.

Task

Research in the judgement modelling and predictive approaches are closely related. Both approaches assume simple models can be used to map financial statement information to judgements and economic events. The cognitive approach has provided us with a descriptive model of information processing in the financial diagnosis task. The processing is performed in two phases corresponding to perception and judgement, in which the first is dominated by information search, and the second is dominated by information integration and interpretation. This view is also supported by recent judgement modelling studies (Rodgers, 1991).

The findings referred to above make it reasonable to assume that the financial diagnosis task is largely similar across task contexts. This view is generally supported by other research in cognitive accounting (Bonner & Pennington, 1991, Gibbins & Jamal, 1993).

Methodologically, three findings relate to the task itself. First, experiments in financial diagnosis must be set up using realistic cue information, setting and response options (Holt & Carroll, 1980). Second, the applicability of traditional information processing theory and methodology have been questioned. One example is the question if verbal protocols accurately correspond to cognitive information processing (Anderson, 1985). A second example is the problems with identifying production rules (Biggs et al., 1987). A third example is the problems with the two-stage operationalisation often used to test propositions based on information processing theory (Choo & Trotman, 1991; Danos et al., 1984). Third, the use of judgement modelling based measures of cue usage has also made the results found in the experimental cognitive approach on factors affecting cue usage doubtful (e.g. Chewning & Harrell, 1990)

Search

We have seen that a strong theoretical basis for the selection of relevant cues does not exist. However, empirical investigation of variations in financial statement information has been used to guide cue selection (e.g. Gombola & Ketz, 1983). The number of different cues used in financial diagnosis studies is large, but to summarise the relevant cues, we have searched for the theoretical concepts assumed operationalised by the cues in a selection of studies. We have mainly concentrated on studies where the theoretical concepts are explicitly mentioned.

Table 2.7 summarises the use of cues to operationalise the 9 most widely mentioned theoretical concepts presumed relevant to financial diagnosis.

Concepts Studies	"Lever -age"	"Profit-ability"	"Liqui -dity"	"Debt cover-age"	"Asset balance"	"Cash posit-ion"	"Size"	"Inter-est cov-erage"	"Capit-al turn-over"	Task con-text[1]
Libby, 1975		X	X		X	X				BP
Casey, 1980	X	X	X		X	X				BP
Abdel-khalik & El-Sheshai, 1980			X	X						BP
Zimmer, 1980	X	X	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
Beaver, 1966	X	X	X	X						BP
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
Gilbert et al. 1990	X	X		X						BP
Laitinen, 1990	X	X	X						X	BP
Mutchler, 1985	X	X	X	X						G
Edmister, 1971	X		X	X						L
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

Table 2.7 Concepts operationalised by independent variables in selected studies of financial diagnosis.

As can be seen in table 2.7, most of the cues are used as operationalisations of theoretical concepts that have previously been identified in factor analysis studies of financial statement information (Gombola & Ketz, 1983; Pinches et al., 1973). This conclusion is obvious from the strong position of the "profitability", "leverage" and "liquidity" indicators. Somewhat surprising is the position of the "debt coverage" concept, a position attributable to Beaver's (1966) early findings. An opposite finding is the weak position of the "turnover" indicators. Inventory "turnover" was a significant factor in both the factor analysis studies referred to above. Only capital "turnover" indicators are among the most widely used indicators in

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

financial diagnosis[1]. Even though economic theory has been shown to be of little importance in the selection of cues, the position of cognitive theory is even weaker. No studies have suggested specific cues based on cognitive theory, such as cues that are particularly useful in economising on a limited information processing capacity. However, one of the main reasons for the popularity of financial ratios may be their ability to economise on information processing capacity. The ratios relate two or more financial statement items to provide new information, and they simplify comparisons with industry standards and evaluations of trends.

Cues measuring "trend" and "stability" or "variability" are not traditionally used, but in some studies including these, increased performance results have been found (e.g. Falbo, 1990). Cognitive descriptive research suggests that "trend" is an important concept, either on its own basis, or to define a standard for the ratios investigated.

Only small systematic differences can be found in the use of financial statement information between task contexts. Also, no large differences can be found in the assumptions of what is relevant financial information between judgement modelling and predictive studies. The use of external non-financial information in addition to the financial information is somewhat more common in the loan decision and bond rating task contexts.

Most judgement modelling studies of financial diagnosis, assume cues presented are cues used. However, several cognitive studies have indicated selective search, and use, of both income statement information and ratio information from more than one year. The number of cues used are assumed to be small, and ratios are the classical financial cue. Cue usage must be measured with reference to a model of the cognitive process, but is typically defined by standardised weights in a linear model in judgement modelling studies, or by the cue's presence in a rule in cognitive descriptive studies. This unfortunate connection between cue usage and the model of the judgement process is particularly obvious in studies comparing cue usage in different linear and nonlinear models. These studies have shown large differences in cue usage depending on what model is used, even though the models show quite comparable performance results (e.g. Schepanski, 1983; Selling & Schank, 1989).

The problem of simultaneous testing of cue usage and model has led to an assumption that subjects are outperformed by linear models because of inferior cue selection, not because of cue weighting. Another explanation of this findings may be that nonlinearity in the subject models are present. However, the inferior cue selection has also been found in cognitive descriptive studies of the task. This finding may also be reversed to guide model

[1] Capital "turnover" indicators are also used as operationalisation of "activity".

development. Thus, a property of the cognitive model should be that inferior selection of cues leads to good fit to the judgement variable, but worse fit to the predicted outcome.

Information overload has been investigated within several of the approaches to financial diagnosis, but no clear indication of an inverted U-shape performance with information load is found. However, performance is found to increase with increased data load up to a point of optimal data load. Agreement on what optimal data load is, has not been found, varying from 8 items in one study to more than 57 items in another.

A main, and widely supported assumption, is that subjects' information search is sequential (Bouwman, 198; Bouwman et al., 1987), but the search is guided by what is characterised by Bouwman (1983), as "checklists". The checklists are guiding structures used to detect interesting and relevant features in the data material, but they do not force a structure on the search. These "checklists" are assumed to be closely related to an organisation of knowledge flexible enough to allow a sequential search through the financial statement, and still can be used to detect relevant features and stimulus dimensions[1] during diagnosis[2].

Reasoning

Reasoning has mainly been investigated within the cognitive approaches to financial diagnosis, and the studies focusing this subject are few. In other approaches, integration of cues used is not investigated as a specific research problem, but assumed to be performed in a simple judgement model.

One assumption on reasoning in financial diagnosis is that it is qualitative (Bouwman, 1983). It has been questioned whether or not this assumption is an artifact of the use of verbal protocol methodology (Meservy et al., 1986), and judgement modelling studies using transformed quantitative information have not shown better modelling results (Schepanski, 1983). No signs of probabilistic reasoning have been found, but non-probabilistic forms of reasoning with uncertainty have been suggested.

Research shows that subjects reach high confidence in their preliminary decision early in the reasoning process, but that they still adjust their preliminary decision to confirming or disconfirming evidence. However, a bias has been found in the direction of the low misclassification cost decision. The early confidence may support the assumptions of hypothesis formation, but explicit formulations are mainly found among experienced

[1] The terms "features" and "stimulus dimensions" are introduced here to signal that both single cues and patterns of cues may relevant to the financial diagnosis. A common term used of these, is "features" if they are discrete and "stimulus dimensions" if they are continuous.

[2] This search strategy may be favourable if diagnostic features are spread out in the information material, and varies by firm.

subjects. Studies closely related to the financial diagnosis task have found reasoning by assumption to be frequently used, but the assumption was often not very clearly stated (Biggs et al., 1987).

Another reasoning strategy found in studies related to financial diagnosis, is reasoning by analogy (Biggs et al., 1987). Reasoning by analogy is a special form of reasoning with reference to a recognised or experienced pattern of financial information. Some authors have suggested reasoning to be replaced by pattern recognition in financial diagnosis (e.g. Bouwman et al., 1987, p. 22). However, the recognition suggested in these studies does not take the form of traditional *pattern recognition*, where the recognition of the pattern as a whole is related to conceptualisation (Dretske, 1981). Rather, it takes the form of recognising a pattern in already conceptualised cues of the financial information. Traditional theory of pattern recognition[1] does not assume such a conceptualisation to take place before the pattern is recognised. Thus, the "pattern recognition" suggested by, for example, Bouwman et al. (1987) is closer to what is usually termed *pattern classification* in cognitive science.

The form of pattern classification described above, and assumed in financial diagnosis, is related to configural information processing, in which patterns of cues of non-additive or nonlinear form is necessary for characterising or classifying the financial situation of the firm. Traditional tests have not found conclusive evidence that configural information processing is necessary and frequently used by financial diagnosticians, but studies in auditing (e.g. Brown & Solomon, 1991) have suggested that the lack of such findings may be explained by the robustness of linear models and the formulation of the financial diagnosis tasks.

The *importance* of pattern recognition is stressed by several authors (e.g. Bedard & Biggs, 1991; Biggs et al., 1993; Selfridge Biggs & Krupka, 1992). Bedard and Biggs (1991) found two important errors in the subjects' reasoning. One related to the recognition of relevant patterns and the other to the relation between the recognised patterns and a relevant hypothesis. However, this study is among the few explicitly *investigating* pattern recognition abilities of subjects. Biggs et. al (1993) characterise the situation in the following way:

> *....pattern recognition in financial analysis is largely untouched by audit judgement research. (Biggs et al., 1993, p. 97).*

Biggs et al. (1993) formulate two important research questions based upon this fact; "How do auditors recognise various trends in financial data ?", and "How do auditors link into patterns various measures of financial performance to identify problems ?" (Biggs et al., 1993, p. 97).

[1] Generally applied to visual pattern recognition (See Ashby, 1992).

In addition, research on patterns of financial data has shown that ratios have to be related in patterns if important aspects, such as materiality errors in the underlying accounts, are to be detected (Kinney, 1987). All these findings suggest that pattern recognition is of no less relevance to financial diagnosis than to auditing. However, the form of pattern recognition suggested in these studies should probably be termed pattern classification. The patterns of financial data suggested "recognised", are conceptualised in intermediate abstractions used to form a diagnosis.

Representation

Reasoning is performed over represented knowledge, pattern recognition is performed by the activation of stored patterns of knowledge, and pattern classification is performed using intermediate abstractions, all stressing the importance of representations.

The *representational content* in models of financial diagnosticians has only a limited theoretical foundation, and is derived by quantitative methods used in judgement modelling studies[1] and by protocol analysis in descriptive cognitive studies. A simple model has been assumed in some bankruptcy prediction studies with, for example, three processes leading to bankruptcy (Argenti, 1976, Laitinen, 1990). Domain knowledge is only seldom explicitly presumed in financial diagnosis tasks[2].

The knowledge represented in a judgement model of financial diagnosis consists of the weights in a linear weighting model. It further assumes that cue values are represented by some "unit" allowing weighting of the value and summation to take place. Other models of the same family assume non-additive or non-compensatory weighting of cues. Even though the integration is nonlinear in some models, the represented knowledge still lies in the weights. Such models does not rely on intermediate abstractions or internal representations in the traditional sense of the terms[3]. Further, the robustness of the models relative to response prediction has made it difficult to reject linear models and accept alternative nonlinear versions.

Knowledge in a cognitive descriptive orientation is traditionally represented in rules. Concepts representing the qualitative transformation of cue values initiate the reasoning over rules. Realisations of rule-based representations are traditionally done in production systems, and intermediate abstractions are also represented by concepts.

---

[1] Such as regression analysis, discriminant analysis or rule induction methods.

[2] One exception is in the descriptive study of Biggs et al. (1993), but their domain knowledge is very firm specific.

[3] Basically, in information processing theory sense of the terms.

Research shows only moderate self-insight in subjects own weighting processes (Mear & Firth, 1987a), and only moderate signs of the use of rules (Meservy, 1986). These findings have led to an assumption that the representational form is schematic. Concepts such as checklists, financial templates, schemata and frames of reference are used to describe the suggested schematic organisation of knowledge. Schematic representation of financial diagnostic knowledge has been suggested because this representational form has the flexibility necessary to explain some of the important behavioural findings on the task. One of these is the finding that search is performed sequentially even when stimulus material order is altered. The representational form required to explain this finding must be flexible. Another finding is that reasoning can be replaced by recognition during the integration phase of the task. Schemata have been proposed as a representational system suitable for this:

> *Schemata are recognition devices whose processing is aimed at the evaluation of their goodness-of-fit to the data being processed. (Rumelhart, Smolensky, et al., 1986, p. 36)*

Mixed results have been found on performance differences between experienced and inexperienced subjects in financial diagnosis (see Bonner & Pennington, 1991). Instead, knowledge organisation and cue utilisation differences have been suggested. One, generally accepted, hypothesis is that a schematic organisation of memory is more common among experienced subjects. Research has been interpreted as supportive to this view (Choo and Trotman, 1991), but tests of alternative organisations of memory are encumbered with large methodological problems. Typically, cognitive accounting studies on the organisation of memory rely on simultaneous tests of two hypotheses. The first relates to differences in knowledge organisation, and the second to measurable consequences of different organisations of knowledge to a specific task.

From the short summary of research on reasoning and representation related to the financial diagnosis task, it is difficult not to support the view of Biggs et al. (1993), referred to above, that a broad understanding of the relevance of these concepts to financial diagnosis does not exist.

Outcome

The standard assumptions on financial diagnosis performance results are still generally valid. These assume that environmental predictability by linear models is high, and new methods have shown even better performance. Subject accuracy is also generally high, but lower than the accuracy of models of the criterion variable. Individual subjects are outperformed by composite judges, but not significantly outperformed by their own models. Consistency and consensus are generally high, but self-insight is only moderate. Further, mixed results have been found regarding differences in experienced and less experienced subjects' performance, suggesting obvious performance differences do not exist. These findings justify the use of graduate business students as surrogates for the relevant subjects performing financial diagnosis in the industry.

Some research results have been found to weaken the position of the standard assumptions. The scale properties of dependent variables in different task contexts have made comparisons of environmental predictability across task contexts difficult. Some authors have even suggested that the scale of the dependent variables and the error measures used, may explain some of the standard assumptions (Mear & Firth, 1987b). Other authors have suggested that the standard assumptions are only valid for the initial parts of the task studied (Barnes & Huan, 1993). Among the strongest critiques of the standard assumptions, is the study of Hopwood et al. (1994), showing that predictability, subject accuracy and model accuracy were low when cost of misclassification, obviously non-stressed cases, and prior probabilities were considered. The mixed results on differences between experienced and less experienced subjects in financial diagnosis have been questioned by several authors, generally finding differences between the two categories of subjects to be related to other aspects of the information processing behaviour than task outcome. The relationship between task attributes and experience differences has been summarised by Bonner and Pennington (1991), leading us to conclude that the financial diagnosis task is not a task where large experience differences relating to outcome will be found.

Despite these and other critiques, the standard assumptions on financial diagnosis outcome still have a strong position due to the large amount of research supporting them in various tasks contexts.

Chapter 3. Cognitive theory

This chapter first argues that the financial diagnosis may be treated as a classification task. Being a new *cognitive* perspective on the financial diagnosis task, thorough introductions to classification theory in general, and to the most recent contributions to classification theory are warranted. Classification theory is introduced in section 3.1. Since connectionist theory of classification represents the most recent development, connectionist theory and connectionist models of classification are introduced and discussed in section 3.2. The principles laid out in this section, are later used to develop the connectionist model of financial diagnosis suggested in chapter 4.

Two approaches may be taken in the selection of cognitive theory relevant to financial diagnosis. In a *theoretical approach* , an investigation of theoretical paradigms of cognitive theory is performed. The purpose of such an investigation is primarily to identify the underlying assumptions of the theoretical perspectives. In revealing these assumptions, the choice of a suitable paradigm or theory can, in principle, be made. In a *task driven approach*, an investigation of the task, problems and the research contributions are performed. The purpose of the latter approach is primarily to describe the accumulated research on the problem and task, and to use this as a starting point in identifying lacking knowledge, and approaches to supply this knowledge. In this dissertation, the latter approach is taken. Chapter 2 contains the analysis of relevant research on the task under investigation.

The theoretical approach was followed in Pedersen (1988), in which we investigated the underlying assumptions of two paradigms in cognitive science. In principle, this knowledge could be used to select the theoretical perspective of greatest relevance to the financial diagnosis task. However, this choice is relative to the task, and therefore, selection of a relevant theoretical perspective depends on what aspects of the task the theory is best suited to explain. For a thorough investigation of the theoretical perspectives in cognitive science of relevance to financial diagnosis, we refer to Pedersen (1988). However, a short presentation of some findings from the analysis is given here.

In cognitive psychology, a division between non-analytic and analytic theories has been assumed to be perfectly correlated with a division of lower and higher order cognitive functions (e.g. Brooks, 1978; Estes, 1994, p. 5). The division has been put forward as the lower limit of cognitive theories (e.g. Dretske, 1981; Pylyshyn, 1984). What Pylyshyn defines as "cognitively penetrable" functions, have been proposed as this lower limit. Traditionally, the analytic tradition has been associated with a *computational theory of cognition* (e.g. Cummins, 1989, p. 108-113; Fodor, 1980; Haugeland, 1985; Pylyshyn, 1984).

In the computational theory of cognition, all cognitive systems are *symbol systems* (Newell, 1980). Cognition is the manipulation of these symbols. The cognitive system is best described by its representations and processes. All representation is performed using symbols. Symbols represent through designation. The manipulation of symbols is assumed to be performed without reference to semantics, but strictly according to syntax. However, the syntax is meaning preserving, making the system process information according to syntactic rules to reach a state of new meaningful information. In this way, the symbols are manipulated in meaningful ways relatively to the represented world.

In this thesis, we will generally define *information processing theory* as a cognitive theory based on this computational theory of cognition. The assumptions of the cognitive system as a symbol system underlie all cognitive theories applied to financial diagnosis within the cognitive approach reviewed in chapter 2.

Above, we have identified the two most important underlying assumptions of this theory, its representational and processing assumptions. The *representational assumption* of information processing theory states that all representation is symbolic. This means that there is a representational system containing symbols that designate. This designation is to objects in the represented world[1]. Thus, each meaningful object in the represented world is represented by one symbol in the representational system. This assumption is often referred to as the representational hypothesis of information processing theory (see Fodor, 1980).

Processing is the manipulation of the representations and thus, takes the form of symbol manipulation. The most important assumption is that this manipulation is independent of semantics, and as such, can be defined purely syntactic. However, the rules of symbol manipulation are defined so that syntactic manipulation of symbols is meaning preserving. Thus, results of syntactic manipulations of symbols is meaningful. This assumption is often referred to as the computational hypothesis of information processing theory (see Fodor, 1980)

These assumptions have been met with severe criticism. The *general criticism* has come from research in philosophy (e.g. Cummins, 1989; Dreyfus, 1972; Clark, 1989, 1993), linguistics (e.g. Lakoff, 1987; Searle, 1980), cognitive psychology (e.g. Brooks, 1978; Estes, 1994), general cognitive science (e.g. Smolensky, 1988), and in different areas of social science research (e.g. Nisbett & Ross, 1980; Nisbett & Wilson, 1977).

---

[1] Information processing theory is not restricted to a representational world outside the cognitive system (external world).

Recently, the information processing theory in *cognitive psychology*, and in particular, the rule-based versions of it, has been met with what has been described as an "anti-rule movement" (see Smith et al., 1992). The objections raised by the "anti-rule movement" stems from at least four research traditions. Instance memory research and instance based approaches to cognition are fundamental in classification and categorisation research (e.g. Medin & Schaffer, 1978; Nosofsky, 1984, 1986), but have gained increased attention in other areas of cognitive processing as well (e.g. Kvadsheim, 1992). This approach states that cognition is not the application of a set of abstract rules, but rather results from comparisons of stored and presented exemplars. A second line of research supporting the "anti-rule movement" comes from connectionist research, in which theories and models have been developed on cognitive phenomena previously assumed only to be explainable by rule-based accounts (see e.g. Rumelhart, Smolensky, et al., 1986; Seidenberg and McClelland, 1989; Sejnowski and Rosenberg, 1986). A third line of research arguing against the universal application of information processing theory to explanation of all cognitive functions, stems from evolutionary approaches to psychology (see e.g. Smith et al., 1992). Evolutionary approaches argue that "much of cognition may be attributable to specific mechanisms rather than to general purpose ones like applying abstract inferential rules" (Smith et al., 1992, p. 2). Thus, the principles of modularity (e.g. Fodor, 1983) argues against the universal application of information processing theory, and not necessarily in favour of it, as maintained by Fodor (1983). The fourth line of research supporting the "anti-rule movement" stems from what Smith et al. (1992) term the heuristic approach to cognition. Examples of this approach are the numerous studies by Kahneman and Tversky (Kahneman, Slovic & Tversky, 1982; Tversky & Kahneman, 1974), suggesting that "people often substitute judgements about similarity for normatively required rule-based reasoning" (Smith et al., 1992, p. 2). Even though the findings within the heuristic approach are interpretable from an information processing theory perspective[1], simpler and more obvious explanations can be given without its assumptions.

In addition to the general theoretical arguments and the empirical findings in different areas of cognitive psychology, *methodological arguments* have been raised against information processing theory. Several methodological principles follow from information processing theory. Two of these are the principles of operationalisation by model and the use of verbal protocol methodology. Operationalisation by model is not unique to information processing theory, but the representational and processing assumptions made by the theory necessitates explicit modelling of representations in symbolic form. However, these representations are based upon the truth of the representational hypothesis, but this hypothesis should be evaluated separately. The use of verbal protocol methodology has been criticised for being

[1] For example, based upon the principles of bounded rationality and limited information processing capacity (e.g. Simon, 1955).

"theory laden"[1] in a similar way, and the relationship between representational units and their manipulation, and verbal utterances is much debated (e.g. Ericsson & Simon, 1984; Nisbett & Ross, 1980; Nisbett & Wilson, 1977).

All this criticism strongly argues against the universal application of information processing theory to cognitive phenomena[2], but in our view, they are not sufficient to justify a position against information processing theory, per se. Surely, some cognitive functions may be best explained by information processing theory (see van Gelder, 1993). For applied science, such as cognitive accounting, a more fruitful position may be to investigate the results of research pursued with different orientations. Consequently, our position here is more pragmatic. It is based upon the achievements made in general cognitive science with recent approaches other than information processing theory, and an interest in investigating the potential of these approaches to financial diagnosis, more than a position against the assumptions in information processing theory, per se. Thus, the definition of financial diagnosis as a task only to be investigated in information processing theory terms is a position we argue against.

The second approach to the selection of relevant cognitive theory starts with an analysis of the task, problem or cognitive phenomena under investigation. Often this is referred to as a task analysis (Newell & Simon, 1972). From the task analysis of chapter 2, we can conclude that the judgement modelling and the predictive approaches to financial diagnosis treated financial diagnosis as a classification task. However, few of the contributions referred to classification theory in cognitive science. Instead, the cognitive approaches to financial diagnosis primarily applied theories from cognitive science developed for the explanation of choice, judgement and problem solving. Several important findings on how the financial diagnosis task is performed resulted from these applications. Still, there are several ways in which the financial diagnosis task differs from the typical choice, judgement and problem solving tasks.

Choice tasks have some basic characteristics. One is the importance of alternative actions and subjective preferences, and how these relate (e.g. March, 1978; Slovic, 1990). For financial diagnosis to be treated as a choice task, these characteristics should be of similar relevance. The financial diagnosis task may be considered as a choice between alternative conceptions of the financial health of the firm. However, explicit preference structures, and clear relationships between preferences and the chosen alternatives may be less obvious. Often, the diagnostic task ends with the chosen conception of the financial health of the firm, and the chosen alternative is not considered to have any further consequences to the

[1] See the review by Crutcher (1994).
[2] Universal application of information processing theory is suggested by several authors, and can be illustrated by the title of Newell's book: "Unified theories of cognition" (Newell, 1990).

diagnostician. Nor can the diagnostician be assumed to have a special preference for one conception before another.

A characteristic detail of *judgmental tasks* is the presence of uncertainty (Osherson, 1990). This uncertainty is often expressed by probabilities in tasks explained by a cognitive theory of judgement (e.g. Tversky & Kahneman, 1974). Furthermore, these probabilities describe the uncertainty of one, or a small number of dimensions or attributes of the stimulus. In the financial diagnosis task, probabilities related to stimulus dimensions are not readily available[1], and the number of relevant dimensions is large.

Two characteristics of *problem solving tasks* are the importance of goals and the role of subgoaling (see Holyoak, 1990). To give an example, the cryptarithmetic problems used in some studies of problem solving, have a goal situation in which a solution to the problem is found, but unlike the choice task, the solution is not found among predefined solutions. Thus, a design element is present in this task. In financial diagnosis, alternative responses are found among predefined classes or categories, and the design element is less important. Furthermore, evaluation of goal attainment is difficult because the "solution" can not easily be evaluated against a goal. Rather, the "solutions" are overlapping and not disjunctive. Another aspect of goals in problem solving, is that subgoaling occurs because the task is large or elements of design is involved. Typically, the solution is built by sequential goal attainment, [illegible]tle evidence pointing to the financial diagnosis task as involving such elements except for the familiarisation/analysis split. However, lack of subgoaling does not mean intermediate abstractions are irrelevant during financial diagnosis. If the development of a financial strategy to bring the firm from one diagnostic characterisation to another is considered, we have a task with important design elements, but this is an example of a financial task that goes beyond the purely diagnostic.

Despite the amount of knowledge generated from research applying the conceptions of the financial diagnosis task referred to above, knowledge of several aspects of the task was reported lacking in chapter 2. Examples of such knowledge were the role of template matching (e.g. Bouwman et al., 1987), pattern recognition and pattern classification (e.g. Bedard & Biggs, 1991), schematic organisation of memory (e.g. Choo & Trotman, 1991), and analogical reasoning (e.g. Biggs et al., 1987), just to mention a few. These concepts have received only limited attention within traditional information processing theory. However, they are important parts of research on cognitive *categorisation* and *classification* (e.g. Smith & Medin, 1981), and on *induction* (e.g. Holland, Holyoak, Nisbett & Thagard, 1986). Thus closer investigation of financial diagnosis as a classification task seems warranted. In section

[1] As explained in chapter 2, the most relevant stimulus dimensions in financial diagnosis consist of financial cue values or patterns of financial cue values.

1.1, we introduced this approach to financial diagnosis, and showed that other diagnostic tasks have been investigated as classification tasks.

Simulated *medical* diagnoses have a long tradition in classification research. This task context has been used extensively in studies of correlated stimulus dimensions (e.g. Medin, Altom, Edelson & Frecko, 1982; Shanks, 1991) and base rate effects (e.g. Estes, Campbell, Hatsopoulos & Hurwitz, 1989; Gluck & Bower, 1988a, 1988b; Shanks, 1992). Applications of categorisation and classification theory to more realistic, practical diagnosis tasks now begin to appear in medical diagnosis (e.g. Brooks et al., 1991). To illustrate the position taken by these applications, consider the opening statement of Brooks et al. (1991):

> *Medical diagnosis is primarily a categorization task. (Brooks et al., 1991, p.278)*

A similar position is taken in other diagnostic areas, such as in psychiatry:

> *Clinical diagnosis is a classification task which uses features or dimensions of relevance across individuals (e.g. the patients mood) to categorize aspects of the patient's functioning into one or more of a finite set of diagnostic disorders, such as those in DSM-III-R.*[1] *(Mumma, 1993, p. 283)*

When comparing these definitions of diagnosis to the description of financial diagnosis given by Methlie (1994, see section 1.1), several similarities can be recognised. Stimuli consist of relevant information, typically derived from the financial statement and represented by financial cues. The response takes the form of a characterisation, often by selecting a predefined class. The financial diagnosis task is the mapping of these stimuli to the predefined classes. This definition of financial diagnosis was introduced in chapter 2, and variations in the response were shown to be dependent upon the task context of the diagnosis.

Several aspects of financial diagnosis are different from medical diagnosis. Since not all firms suffer from relevant "diseases", the task of the diagnostician is more like the general medical practitioner when asked for a medical certificate. This situation requires a more general approach to the task, and makes initial formulations of hypotheses less relevant. Next, the producers of the financial statement showing the "symptoms" have the best knowledge of the firm's "diseases", and often do their best to hide relevant "symptoms" in their report. Furthermore, the analyst performing the diagnosis is rarely interested in how identified "diseases" should be "cured". However, these differences do not make the application of classification research less relevant to financial diagnosis than to medical diagnosis.

---

[1] Diagnostic and Statistical Manual of Mental Disorders (our note).

3.1 Classification theory

A classification task can be described by assuming a vector of stimulus dimension values. This vector may consist of continuos or binary values, traditionally termed features (Garner, 1978). In the case of financial diagnosis, the stimulus dimension vector consists of the relevant real valued financial statement cues I_{pi} in which i indicates the relevant stimulus dimension where i=1,2,...,n, and p indicates the pattern presented, where p=1,2,...,N. The classification is the assignment of the stimulus to one of M prespecified classes C_j, where j=1,2,...,M (Lippmann, 1993).

Categorisation is distinguished from classification by Estes (1994), in classification being the pure assignment of the stimulus to one of M prespecified classes, whereas categorisation "carries the further implication that knowledge of the category to which an object belongs tells us something about its properties" (Estes, 1994, p. 4). The further implications carried by knowledge of category membership are particularly useful in diagnosis. Categorisation implies missing information of properties can be inferred (Smith & Medin, 1981, p.1), and that similarity evaluations of stimulus and previous experiences can be performed. In this thesis, our term classification is used synonymously to the categorisation term of Estes (1994), because the assignment of an object to a class in itself is of less interest unless inferences can be drawn from the classification. This interchangeable use of the terms classification and categorisation is common within cognitive psychology (see Nosofsky, 1984; Smith & Medin, 1981)

Three main theories of classification are frequently mentioned in standard textbooks on cognitive psychology (e.g. Ellis & Hunt, 1993). All theories assume that a stimulus is represented by an array of feature-, or stimulus dimension values (Estes, 1994). *Attribute theory* or *definitional theory* assumes that specific features are necessary, and collectively sufficient, to define a stimulus as member of a class. Class membership is evaluated by identifying presence or absence of these features. Consequently, a class is defined by a list of necessary and collectively sufficient features. Models implementing definitional theory includes traditional symbolic systems, typically relying on production system representations. The definitional theory has been heavily criticised since the 1970's, leading to alternative theories of classification.

Prototype theory (Rosch, 1978; Rosch & Mervis, 1975) assumes that the representation of a category or class is done by averaging the feature values of cases within a class, and by representing the average feature values as the representation of a prototypical or best member of a class or category. The evaluation of class membership is performed by evaluating similarity in feature values of the presented stimulus to the prototype(s) in each class. Implementations of prototype theory have been done in, for example, feature frequency models (see e.g. Hurwitz, 1994) and in connectionist models (e.g. Gluck & Bower, 1988a, 1988b).

Exemplar theory (Estes, 1986; Medin & Schaffer, 1978; Nosofsky, 1984, 1986) does not assume that averaging of feature values occurs across members of a class. Instead, exemplar theory assumes that each exemplar is represented in memory. This is done by a memory trace representing the relevant feature values of each exemplar. Representations of exemplars are activations in psychological space. Evaluation of class membership is a consequence of the similarity evaluations performed when new stimuli are presented. Implementations of exemplar theory have been done in the formulations of general context theory (Nosofsky, 1984), and in more recent connectionist models (Estes, 1993, 1994; Kruschke, 1992).

As shown by Smith and Medin (1981), both the prototype and exemplar theories can solve the problems raised by the critics of definitional theory, but depending on their operationalisation, new problems arise.

These classification theories have recently been implemented in connectionist models (e.g. Estes, 1994; Hurwitz, 1990; Kruschke, 1992), whereas other classification models built independently from a *connectionist* approach to cognition, now also have appeared in classification research. Three examples of the latter models are the simple adaptive network model of Gluck and Bower (1988a, 1988b), the configural adaptive network model by Gluck (1991), and the attentional connectionist model of Schanks (1992). These connectionist models all have *implicit* prototype representations, and have been classified as prototype models (McClelland & Rumelhart, 1986, p. 173; Robins, 1992 p. 46; Shanks, 1991 p. 433). Thus, connectionist models of classification have been developed both within, and independently of, classification research. This is a situation very different from other areas in cognitive science, were the positions taken by researchers in the information processing and connectionist paradigms seem rather irreconcilable (e.g. Fodor & Pylyshyn, 1988; Smolensky, 1988).

In this section, we introduce the traditional classification theories and their relationship to traditional information processing theory. We use these theories to show that many of the weaknesses of information processing theory have led to the development of alternative theories of classification. Next, we introduce some recent connectionist theories of classification and some recent connectionist implementations of traditional classification theories.

3.1.1 Definitional theory

Definitional theory is a theory of categorisation and classification that assumes class membership is determined by a set of *defining* features or stimulus dimension values. Other terms frequently used are attribute theory (Ellis & Hunt, 1993), classical theory (Ashby, 1992; Lakoff, 1987; Smith & Medin, 1981) and set theory (Lakoff, 1987; Rumelhart & Norman, 1985).

Definitional theory makes three assumptions (Smith & Medin, 1981, p. 23-24). The representation of a class is a *summary representation* of the entire class. Consequently, the representation is the result of an abstraction process. The abstracted representation consists of a list of defining features. By defining features, we mean features *necessary* and *collectively sufficient* for classification. Last, classes are represented in a hierarchy, in which defining features of a superset are *nested* in subsets. By nested features, we mean that defining features of a superset are found in the subset. However, the subset also contains some features not shared by the superset.

The assumption of abstraction is shared among all theories of classification, including exemplar theory. The assumption of defining features states that every exemplar of a class must have the features, and if they are found, the exemplar is a member of that class. This assumption excludes disjunctions of features as defining (Smith & Medin, 1981, p. 24). In definitional theory, the assumption of nesting gives categorisation its inferential power. By nesting of features, the classification of an object makes it possible to infer how features of supersets are true of the subsets. For example, if high profitability is a defining feature of a success firm, and a growth firm is a subset of a success firm, it can be inferred that the growth firm has high profitability.

Definitional theory is a representational theory. The simplest processing model based upon the representational assumptions in definitional theory, is a model equating problem solving as hypothesis testing and classification (e.g. Bruner, Goodnow & Austin, 1956). The model simply states that one can start with the hypothesis of the stimulus as member of a class, and check whether the defining features are present. If not, proceed to the next hypothesised class. Estes (1994) refers to the differences between analytic and non-analytic approaches to cognition, and views "categorisation-as-hypothesis-testing" as the analytic theory of classification:

In the hypothesis-testing approach, concept formation is treated, in effect, as a form of problem solving. (Estes, 1994, p.6)

The problem of the learner in such studies is viewed as one of formulating hypotheses about the critical features and testing the hypothesis against observations of a sequence of category exemplars until an adequate hypothesis is discovered (Estes, 1993, p. 16)

There is a close connection between the conception of classification as hypothesis testing and definitional theory (Estes, 1993). If classification is based upon definitional theory, rules of logic can be used to deductively reason from features to class in classification (Lakoff, 1987, p. 7). Thus, definitional theory and rule-based approaches to classification go hand in hand with the hypothesis-testing view of classification. This is a relationship not only found in classification research. In the words of Estes (1994), the relationship can be described in the following way:

...there has traditionally been a sharp opposition between approaches to categorization and induction centred on the discovery and use of rules and approaches based on processes of learning and memory. Over several decades since the early work on concept formation, these approaches have diverged, as witness the debates between proponents of connectionist and rule-based treatments of language acquisition and processing and reasoning, and the efforts to choose between interpretations of concept formation and categorization based on instance memory and those based on rules and hypothesis testing. (Estes, 1994, p. 244 (original citations left out)).

We can illustrate this relationship by showing how definitional theory can be related to rules. The theory assumes a two-stage decision process when classification is performed. Since definitional theory assumes that features of the stimulus can be detected from relevant stimulus dimensions, this *detection* is the first stage of the classification. This implies a characterisation of relevant stimulus dimension values as binary features, and is an example of the transformation from analog to digital form proposed by Dretske (1981). This stage is traditionally performed by implementing a decision bound on the stimulus dimension, and can be implemented in a simple production rule (e.g. Newell, 1990; Newell & Simon, 1972). Consequently, the decision bounds are always orthogonal to the stimulus dimensions (Ashby, 1992). The next stage is to assign objects to a category based upon the combination of identified features. This assignment is also done by the application of rules. If the rule contains several conjunctions, a system of limited cognitive capacity (Simon, 1955) might introduce intermediate abstractions to represent the temporary conclusion on an evaluation of a stimulus dimension. In this way the classification can be hierarchical, in which the intermediate abstraction plays an important role (Chandrasekaran & Goel, 1988).

When applied to financial diagnosis, the use of rules and the formation of decision rules based upon definitional theory can be illustrated as in figure 3.1.

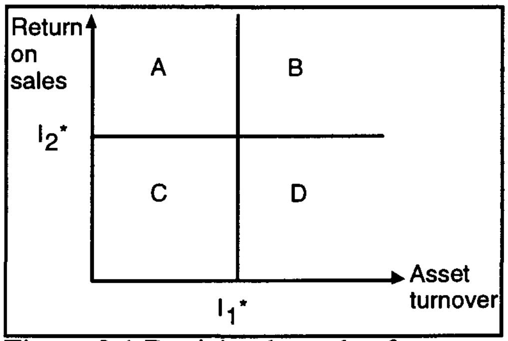

Figure 3.1 Decision bounds of definitional theory

In figure 3.1, the class B of highly profitable firms - success firms, can be identified with the application of the rule; if $I_1 > I_1^*$ and $I_2 > I_2^*$, then the object is in class B. To use the rule, we must first convert the values of the stimulus dimensions into digital form (Dretske, 1981). This is similar to what Methlie (1993, p.147) has termed "qualification". Next, the rule can be applied. Similarly, in hypothesis-testing terms; if the initial hypothesis is to test if the object is in category B, the features that must be identified are the definitional features of the rules listed above.

The representational assumptions of definitional theory have been criticised on both general and empirical grounds. Smith and Medin (1981, p. 32) list four general criticisms; the exclusion of functional features[1], the existence of disjunctive concepts, the existence of unclear cases, and the failure to specify defining features. Only the last three criticisms are considered relevant, since definitional theory does not exclude functional features per se, even though the use of functional features has not been common in experimental studies using definitional theory (Smith & Medin, 1981, p. 27). In addition to these general criticisms, empirical criticisms have been raised against the ability of definitional theory to explain typicality findings (Rosch, 1978), the use of non-necessary features in classification (Smith & Medin, 1981, p. 43-45), and empirical findings related to the nesting of concepts (Smith & Medin, 1981, p. 47-49). In the following, we summarise some of these criticisms.

With the assumption of defining features of a class, definitional theory rejects *disjunctive classes or categories* (Smith & Medin, 1981, p. 28). Smith and Medin (1981, p. 29) question the prevalence of disjunctive concepts, but to illustrate the necessity of disjunctions in classification, consider the financial diagnosis situation of figure 3.1. If we want to assign objects in decision regions B and C to the same class, disjunctions of and-productions must be formed. The introduction of disjunctions makes it possible to use rule-based classification even if the requirements in definitional theory of necessary and collectively sufficient features are not satisfied.

[1] The term functional features are used by Smith and Medin (1981) of *abstracted* features functional to the classification task.

The nesting of features assumed in definitional theory does not allow judgements of a class as a subset of another to be *unclear*, or the membership of a class to be *graded* (Zadeh, 1965). However, unclear subset classification is found in human classification studies (Smith and Medin, 1981, p. 29), and graded membership is typical of many classes (Lakoff, 1987). As a general argument against definitional theory, this argument is weak, but if empirical findings of unclear classification exist, the argument is relevant.

The last general criticism of definitional theory stems from the well known argument of Wittgenstein (1953) that some categories lack defining features. However, Smith and Medin (1981) suggest that this argument is weak as a general argument, because it may be that we have not searched for the right kind of defining features. Consequently, the general arguments against definitional theory do not make it clear that definitional theory should be rejected, but by adding the empirical criticism raised against definitional theory, Smith and Medin (1981) find the arguments against definitional theory to be strong.

The first of these arguments comes from a series of studies on typicality effects in classification (e.g. Rips, Shoben & Smith, 1973; Rosch, 1978; Rosch & Mervis, 1975). Subjects are able to judge how typical an exemplar is of a class, and these judgements are highly reliable across raters. The interesting findings are that typicality is negatively correlated with response time and positively correlated with accuracy in classification tasks. Furthermore, typical members are the first learned by children, and they are the most likely to be mentioned by subjects when asked for an example of a class. Since definitional theory assumes defining features are sufficient to classification, all exemplars should be equal members of a class. To explain these findings, rather speculative processing assumptions must be introduced for the representations of definitional theory to be sustained (Smith & Medin, 1981, p. 36-37).

The next empirical argument against definitional theory comes from findings that subjects often use non-necessary features in classification tasks (Rips et al., 1973). The use of non-necessary features is closely related to a probabilistic view of classification. Non-necessary, but easily identifiable features may be reliable indicators of class membership in most cases, even though there may be exceptions. Not to take advantage of such easily identifiable features would be disadvantageous, and a theory of classification should incorporate their usage.

If subclasses are nested by defining features of a superclass, definitional theory predicts that a subclass "should always be judged more similar to an immediate superordinate than to a distant one" (Smith & Medin, 1981, p. 47). This is often the case in experimental studies, but exceptions have been found indicating further weaknesses in definitional theory (Smith & Medin, 1981, p. 48).

A processing assumption often criticised in implementations of definitional theory is the assumption of independent decisions (Ashby, 1992). This assumption states that the decision bounds of one dimension is independent of the value of another stimulus dimension. This assumption greatly simplifies classification because stimulus dimension evaluations can be performed individually. However, this assumption can be overcome in rule-based representations by assuming that $I_1{*}$ is different for different values of I_2. This means the rules are sensitive to the context of the other stimulus dimensions, and this is often illustrated in a decision tree. In the predictive literature, several rule induction algorithms with the property of inducing such context sensitive rules have been developed (e.g. Quinlan, 1986). This assumption is closely related to the use of disjunctions in rules. A class may have no defining features, so the exemplars must be classified using a set of disjunctions. Then, stimulus dimensions can not be evaluated independently. Consider firms with four features; good situation (GS), bad situation (BS), good trend (GT) and bad trend (BT). The disjunction used for classification of a success firm may be: (GS and BT) or (GS and GT) or (BS and GT). The only conjunction left is defining of a distressed firm: (BS and BT), but neither of the features were defining for success firms, and the BT and BS features were only defining for distressed firms in the context of each other. Consequently, independent evaluation of features and stimulus dimensions is closely related to disjunctive categories, and to retain independent evaluations, the theory must reject disjunctions.

Rule-based accounts of classification are closely related to definitional theory, and if definitional theory is correct, rule-based accounts of the simplest kind can perform the necessary classification. In this section we have shown how certain criticisms of the definitional theory could be overcome while still using a rule-based account of classification. Most empirical criticism stemming from the findings of unclear cases have been incorporated in rule-based accounts[1], but when probabilistic classification is to be accounted for, prototype theory is more common.

3.1.2 Prototype theory

Prototype theory has partly been developed to account for the general and empirical findings used as arguments against definitional theory. Major researchers on prototypes have stressed that the prototype concept should be used to illustrate certain typicality effects in classification, and not as a proposal for a general theory of classification (Rosch, 1978).

[1] Using, for example, the principles of Zadeh (1965).

Several definitions of the prototype concept have been given, but they all refer to some *abstract representation of commonalties among stimuli.* Different formulations of a theory of classification based upon the idea of a prototype are found (Smith & Medin, 1981).

Prototype theory states that when subjects "is presented a set of stimuli for purposes of learning, they abstract the commonalties among the stimulus set and the abstracted representation is stored in memory" (Ellis & Hunt, 1993, p. 217). Ellis and Hunt (1993), further equate prototype and schema representations to illustrate the abstract character of the prototype.

The first ideas of a prototype are often attributed to Wittgenstein's (1953) criticism of definitional theory. He questioned the assumptions of clear boundaries and common properties of categories, and used the category of games as an example (Lakoff, 1987). According to Wittgenstein, the category "game" included plays that had no necessary and collectively sufficient common properties that characterised them as a "game". More formal tests of the assumptions of definitional theory were performed by Rosch (e.g. Rosch & Mervis, 1975) in a series of experiments investigating the proposition of definitional theory that no exemplars should be considered more typical than another.

The theoretical status of prototype theory is somewhat unclear. Lakoff (1987 p. 44) cites Rosch in stating that prototypes, in her opinion, did not constitute any particular model of processes, but should rather be used as a convenient grammatical fiction. Consequently, in Rosch's view (Rosch, 1978), prototypes do not constitute any particular theory of classification and categorisation. Lakoff (1987) characterises prototype *theory* as a proposition of Rosch to be a misunderstanding of her intentions. Thus, the term prototype is used as a *framework* for models including the idea of an abstract representation of commonalties, or as stated by Estes (1994):

> *Formal theory in the categorization area has a curious aspect in that prototype theory is by far the most visible variety in the literature......although it can be credited with none of the close quantitative accounts of categorization data that have appeared during the last decade, the majority of which have been achieved by exemplar-similarity models....... The popularity of prototype theory appears to be attributable to a combination of factors, among them its intuitive appeal, its long history, and some results of experiments employing categories of objects produced by means of variations in experimenter-defined prototypes. (Estes, 1994, p. 51-52 (references in the text excluded))*

Despite the vague notion of the prototype concept, operationalisations of prototypes and developments of classification theories inspired by the idea of a prototype can been found. Common to these operationalisations is the assumption that the representation of a category is a summary description of the entire class. In contrast with definitional theory, the representation of the class does not contain a set of definitional features, but rather a central tendency (Smith & Medin, 1981, p. 61[1]). Three different formulations of prototype theory were mentioned by Smith and Medin (1981).

A featural approach presumes feature representation of the stimulus. The features are salient and have a "substantial probability of occurring in instances" (Smith & Medin, 1981, p. 62) of the class. However, the features may not be defining. The representation of a class consists of a feature list and associated weights of each feature. These feature weights can represent both salience and conditional probabilities (Smith & Medin, 1981, p. 62). In some models (see Hurwitz, 1994), feature frequency is assumed represented in the weights. In other models, the weights have a more complex interpretation (Gluck & Bower, 1988a, 1988b). Classification is traditionally performed in the simplest featural approach by adding weights on features present until a threshold for the class in question is reached. Other versions of the featural approach assume more complex featural evidence accumulation. Two examples are the feature comparison model of Smith, Shoben and Rips (1974), and the well known contrast model of Tversky (1977[2]). Another formulation following the principles of a featural approach, is the spreading activation model of Collins and Lofthus (1975).

Featural formulations of prototype theory have recently been implemented in connectionist models. Connectionist models of classification have been developed from advances in connectionist research in other areas of cognitive science (e.g. Rumelhart & McClelland, 1986). These operationalisations have partly been developed independently from traditional classification research, but follow the assumption of an abstract representation of the prototype. In these connectionist models of classification, a prototype is represented "implicitly" in the model by weight patterns. The simplest connectionist model of classification is the model now often characterised as the "standard connectionist model[3]" (Estes et al., 1989; Gluck & Bower, 1988a, 1988b). This model is illustrated in figure 3.2.

---

[1] Smith and Medin (1981) use the terms "probabilistic view" of classification theories based upon the idea of a prototype.
[2] See Smith, 1990.
[3] We will later use the term "adaptive network model" of this model, but to illustrate that this model is basic to connectionist models of classification, we use the term "standard connectionist model" of Shanks (1991) here.

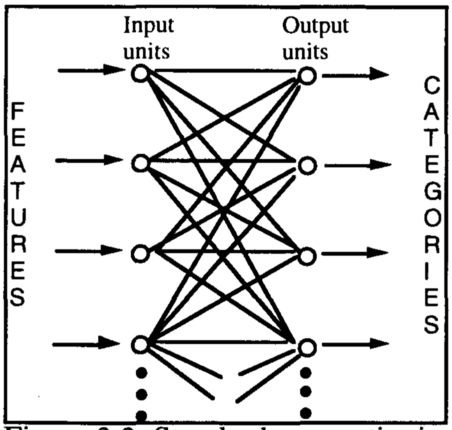

Figure 3.2. Standard connectionist model (From Shanks, 1991)

A set of features are represented by input units, whereas class or category membership is represented by output units. The connections between input and output units each have an associated weight, indicating the salience and probability of the feature, given the category. The outputs of the class or category membership units are the weighted sum of the activation of the input units. By transforming the weighted sum in a sigmoid function or a choice rule (e.g. Luce, 1963), class membership probabilities can be modelled. This simple model has been extended and refined by several researchers (e.g. Gluck, 1991). A more thorough understanding of these models requires general knowledge of the connectionist paradigm and connectionist theory. A broader introduction to connectionist theory is given in section 3.2. After central terms in connectionist theory have been explained in section 3.2, we will return to these models, and to their relevance and application to the financial diagnosis task.

The second formulation of prototype theory does not assume a featural representation of stimuli, but uses a stimulus dimension representation. In this approach, a stimulus is represented by the real values of each stimulus dimension (Smith and Medin, 1981). By using stimulus dimension values, the interpretation of prototypes and exemplars as represented by points in multidimensional psychological space comes natural. Thus, this operationalisation assumes that stimulus dimensions are salient, and that a classes are represented by either an average stimulus dimension value[1] (Reed, 1972; Smith & Medin, 1981, p. 102), or by some other "ideal point" in psychological space. Usually, this ideal point is represented by some "focal member" of the class (Estes, 1994, p. 54).

With the interpretation of prototype representations as points in psychological space, similarity between an object and a prototype can be interpreted as a function of distance in psychological space. Similarity is inversely related to distance in psychological space, and can be used to perform a classification. Shepard (1958, 1987) assumed that similarity of two items in psychological space is defined by:

$$s(p,a) = e^{-cd_{pa}}, \tag{3.1}$$

[1] Or some other measure of central tendency.

where $s(p,a)$ is the similarity of item p to a, c is a constant and d_{pa} is the distance between item p and a. When applied to prototype theory, a indicates the prototype of a class j. However, other similarity measures are often used in prototype theory. A simple distance model (Smith & Medin, 1981, p. 107) assumes that a stimulus object is classified as belonging to a class if the distance between the object and a prototype of the class is below some threshold. The comparative distance model (Smith & Medin, 1981, p. 110; Reed, 1972) assumes that the object is assigned to the class with the closest prototype, measured by some comparative distance measure. However, models incorporating similarity measures of the kind proposed in equation (3.1) have been developed using both the central tendency (Ashby, 1992, p. 474-475) and the "ideal point" (Massaro & Friedman, 1990[1]) representations of the prototype.

When applied to financial diagnosis, the use of prototypes and similarity measures in psychological space can be illustrated as in figure 3.3.

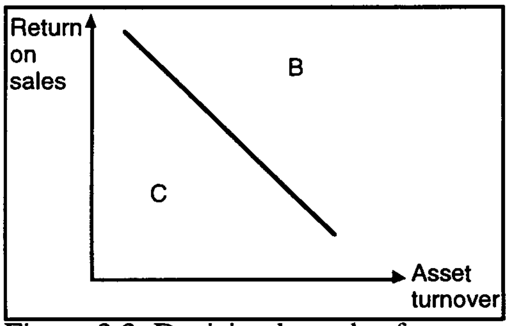

Figure 3.3. Decision bounds of prototype theory

When referring to figure 3.3, a firm can be classified by measuring the comparative distance of the object to the two prototypes B and C, and classifying the firm in the class with the shortest distance to the prototype. In figure 3.3, the axis corresponds to the perceived dimensions "return on sales" and "asset turnover". If these perceptual dimensions correspond to measured dimensions, classification with prototype theory equals the use of a linear classifier (Ashby, 1992), such as linear discriminant analysis.

The third operationalisation of prototype theory also mentioned by Smith and Medin (1981), assumes that the prototype is represented by a *template*. Several operationalisations of the template concept exist, but their common assumptions are that the template is "isomorphic to the object it represents, unanalysable, and inherently relational" (Smith & Medin, 1981, p. 131). The assumption of isomorphism implies that abstraction in a template model is different from our previous prototype models. The template is considered to be more perceptually similar to the objects and to be more holistic. Thus, parts of the template can not be analysed separately.

[1] See Nosofsky, 1992, p. 35-36.

Studies using the template operationalisation of prototypes often study objects with perceptual features. The traditional application presumes heavy pre-processing of perceptual objects to the comparable size and orientation. Next, a grid-like template is matched to the perceived and pre-processed representation (Smith & Medin, 1981, p. 132-134). The approach holds a strong position in machine perception and pattern recognition, but as a psychological model, the formulation has received less attention.

Both the featural and the stimulus dimensional prototype theories can explain the empirical findings incompatible with definitional theory (Smith & Medin, 1981, p. 163). Here, we focus on the explanations given by the featural approach. The featural approach allows *disjunctive* categories. The critical weighted sum needed to place an object in a class can be achieved by various combinations of features (Smith and Medin, 1981, p. 65). Furthermore, *unclear* cases are allowed when an object has features shared by many classes, and the weighted sum of the features is below a critical value for all classes. The featural approach emphasises features with high probability of occurrence within a class, but does not require these features to be *defining*. The same assumption also explains subjects' use of *non-necessary* features. If we assume that typicality is related to the weighted sum of shared features, *typicality* effects can easily be explained. For example, short classification time for typical objects can be explained because the sufficient weighted sum is achieved more quickly when these objects are evaluated. The finding that some objects were rated more similar to their distant superordinates than to their nearest superordinates was incompatible with the assumption of *nested features* in definitional theory. In the featural approach, weighted sums for an object may be higher for a distant superordinate than for an immediate one, even though this will be the exception rather than the rule (Smith & Medin, 1981, p. 71).

Even though prototype theory is able to explain the findings used as arguments against definitional theory, prototype theory has generated some new problems. Some of these problems are related to the particular prototype theory operationalisation, while other are more general. Here, we focus on the general problems. Definitional theory was found to be too constrained in its definition of categories or classes. The opposite may be a problem for prototype theory. One example of lack of constraints, is the ease with which disjunctive classes are accepted in prototype theory. Even though most categories are not disjunctive, there is nothing in prototype theory that "favours low degrees of disjunctiveness over high ones" (Smith & Medin, 1981, p. 88). Another example may be that relationships between classes may be a part of their representation, indicating that a more complex prototypical representation than found in most prototype theory operationalisations, is necessary[1] .

---

[1] Such as a frame (Minsky, 1975).

Another problem following from the simplicity with which a class is represented, is the difficulty of prototype theory in dealing with correlated features. Studies have shown that subjects are sensitive to correlated features (Ashby, 1992, p. 451). Introduction of conjunctive features has been suggested as a way to represent correlated features in prototype theory, but this will result in an explosion in the number of features in the prototype representation[1]. A similar problem is the treatment of context effects in prototype models. Context effects are relevant when features are correlated, and must be treated similarly.

A problematic finding for prototype theory to explain, is the finding that other individual members of a class than the prototype, may have an effect on classification performance (Brooks, 1978; Medin & Schaffer, 1978). This has led researchers to formulate theories in which several exemplars of a class are given significance in the class representation.

3.1.3 Exemplar theory

Another line of criticism of the definitional theory and the rule-based accounts of classification has come from researchers stressing the importance of instances or exemplars in cognitive processing (e.g. Estes, 1986; Medin & Schaffer, 1978; Nosofsky, 1984). Their work in classification research is paralleled by similar instance approaches to other cognitive phenomena, such as memory (Brooks, 1978), perception (Whittlesea & Brooks, 1988), judgement (Tversky & Kahneman, 1974), reasoning (see Smith et al., 1992) and choice (Kvadsheim, 1992).

Exemplar theory has been developed as an alternative to the assumption in prototype theory, that the only class "exemplar" relevant to classification is the prototype, and it has been shown that category exemplars, other than the prototype, can have "pronounced effects on categorization performance" (Ashby, 1992, p. 451).

In its extreme form, such as in the proximity model of Reed (1972)[2], exemplar theory states that each class is represented by all the instances the subject has encountered of that class. Thus, no abstraction occurs, but this seems an implausible model. At the other extreme, consider a model where each class is represented by its best exemplar or a "focal exemplar" (Rosch & Mervis, 1975). However, this model falls into the category of prototype theory. Consequently, most exemplar theory models state that classification is performed by comparing the stimulus presented to a set of exemplars in each class. The responded class is the class with the highest evaluated similarity to the presented stimulus.

---

[1] The same idea has been implemented in the standard connectionist model to overcome the same problems (e.g. Gluck, 1991).

[2] See Smith and Medin, 1981, p. 146.

Several formulations of exemplar theory exist. Among the most widely known are context theory (Medin & Schaffer, 1978), later generalised to general context theory by Nosofsky (1984, 1986). Other models developed from exemplar theory are the MINERVA model of Hintzman (1986) and the array model of Estes (1986). Here, we focus on the widely known context theory and its generalisation to continuous stimulus dimensions in general context theory (Nosofsky, 1984, 1986).

In context theory, it is suggested that subjects learn to attend selectively to stimulus dimensions. Consequently, exemplars are only represented to the extent that they differ on stimulus dimensions, and abstraction in the form of selective attention takes place (Medin & Florian, 1992). A consequence of selective attention is that the subjects can attend to properties that occur frequently, and thus develop a "detailed representation of typical exemplars but only an incomplete or collapsed representation of atypical exemplars" (Smith & Medin, 1981, p. 153).

The core of context theory lies in its similarity processing assumptions. With a featural representation of the exemplars, classification is performed by assigning the object to the class with the highest conditional probability computed as:

$$P(C_j|p) = \frac{\sum_{a \in C_j} s(p,a)}{\sum_{a \in C_j} s(p,a) + \sum_{a \notin C_j} s(p,a)}, \tag{3.2}$$

where $P(C_j|p)$ is the probability of stimulus p being classified in category j, $s(p,a)$ is the similarity of stimulus p to a stored exemplar representation a. In context theory, $s(p,a)$ is computed by a *multiplicative* similarity rule (Estes, 1994; Medin & Florian, 1992):

$$s(p,a) = \prod_{i=1}^{n} s_i, \tag{3.3}$$

where s_i is the similarity between p and a on feature i. For a match between two features, $s_i = 1$, and for a mismatch $0 \le s_i \le 1$. The use of a multiplicative similarity rule implies similarity is measured sensitively to both correlated features and feature frequency (Medin & Florian, 1992). The diagnosticity of a feature is relative to the context of other features, thus the term *context* theory is applied. Context and feature frequency insensitivity were important limitations of the prototype theory formulations presented in section 3.1.2 above.

While the context theory presumes a featural stimulus representation, *generalised context theory* (Nosofsky, 1994, 1986) generalises the theory to continuous stimulus dimensions. When continuous stimulus dimensions are introduced, the interpretation of representations of exemplars as points in multidimensional psychological space is useful. Nosofsky (1984) has shown that the similarity measure in equation (3.3) is equivalent to the similarity measure proposed by Shepard (1958, 1987), and referred in equation (3.1) when:

$$d_{pa} = \sum_{i=1}^{n} | p_i - a_i |. \tag{3.4}$$

Thus, similarity between a presented object and a represented exemplar is an exponential decay function of psychological distance when psychological distance is computed using a city-block metric. The city block metric is computed as the summed absolute valued distance summed over all stimulus dimensions.

Furthermore, Nosofsky (1984) has shown the close relationship between equation (3.2) of context theory and the Luce (1963) choice rule for stimulus identification, so that the similarity measure of equation (3.1) can be used with the classification model in equation (3.2) even when the stimulus dimensions are continuous.

Since general context theory operates in a psychological space, its stimulus dimensions can be determined using multidimensional scaling (Shepard, 1962). A multidimensional scaling solution (MDS) is obtained, and exemplars are shown as points in n -dimensional space, where n is the number of stimulus dimensions in psychological space. Solutions in this space should be in accordance with equation (3.1). Classification of new exemplars is done by computing the similarity of the new exemplar to all represented exemplars of each category. The probability of being classified in a particular category is the summed similarity of the new exemplar to all exemplars in the category divided by the summed similarity of the new exemplar to all exemplars as shown in equation (3.2). Similarities used for summation must be in accordance with equation (3.1).

When applied to financial diagnosis, the use of general context theory can be illustrated as in figure 3.4.

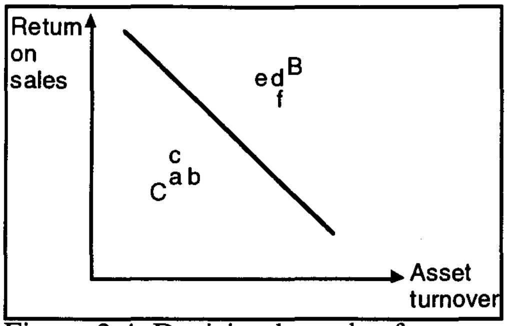

Figure 3.4. Decision bounds of exemplar theory

In figure 3.4, the stimulus dimension values of six firms are shown as points in multidimensional psychological space. The dimensions shown are perceived "return on sales" and "asset turnover". Similarity in the psychological space is computed using equations (3.1) and (3.4)[1]. Three exemplars of each class are shown. The decision bound in psychological space is the contour for which the probability of being classified in class C and B are equal. In this simple case, this bound is linear, but as the distributions of exemplars in each class change, the decision bound will be defined by nonlinear contours.

A change in exemplars of a class is illustrated in figure 3.5. The exemplars g, h, i , j, k, l and m are new exemplars placed on the category decision bound. When exemplars in class B are changed, the decision bound becomes nonlinear. This illustrates the sensitivity of the theory to differences in exemplar distributions of classes. Consequently, several conditions may change the classification of a firm to the "success firm" class. One example is the selective attention to dimensions, which changes the metric properties of the multidimensional space shown. Another example is a change in the distribution of represented exemplars within a class.

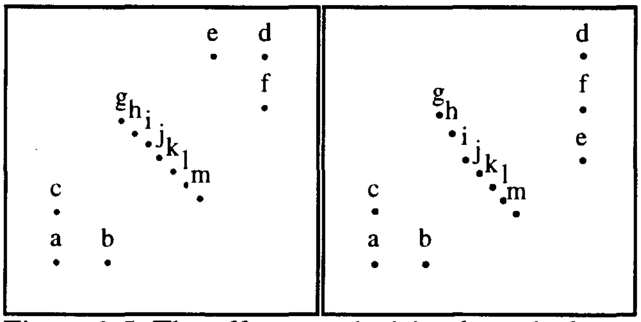

Figure 3.5. The effects on decision bound of changing exemplar distribution in a category.

General context theory has been implemented in a connectionist model by Kruschke (Kruschke, 1992, 1993a; Nosofsky & Kruschke, 1992). The implementation, ALCOVE[2] , is illustrated in figure 3.6. The illustration shows the simple case of two stimulus dimensions derived from MDS as dimensions in psychological space. Values on these stimulus dimensions are represented by the two input units in the bottom part of figure 3.6.

[1] Remember that the illustrated space is Euclidean, so that direct interpretation of distances and similarities is difficult.
[2] Attention learning covering map.

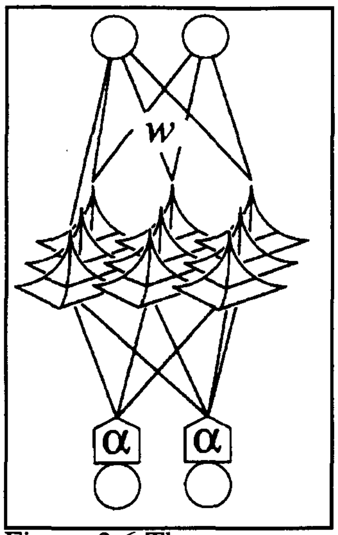

Figure 3.6 The ALCOVE model (From Kruschke, 1992)

The parameter α allows *selective attention* to the psychological dimensions. The middle layer of units in ALCOVE consists of exemplar nodes positioned at the stimulus dimension values of each exemplar in psychological space. Their values are computed as an exponential decay function of the sum of the city block distance between their position and the "position" of the stimulus object in psychological space summed over all stimulus dimensions. Consequently, the units in the middle layer have a value corresponding to the similarity of the stimulus object to each represented exemplar. The output units of ALCOVE represent category or class "activations". These "activations" are used in a traditional choice model (e.g. Luce, 1963), just as class probabilities were computed in the generalised context model. Class nodes and exemplar nodes are connected by *associative weights* that indicate the association between an exemplar and the different classes. The main difference between general context theory and the ALCOVE implementation is that the attentional parameters and associative weights are estimated by error driven learning. However, a more thorough understanding of the model requires general knowledge of the connectionist paradigm and connectionist theory. A broader introduction of connectionist theory is given in section 3.2. After central terms in connectionist theory have been explained, we will return to these models, and to their relevance and application to the financial diagnosis task.

The empirical findings raised as arguments against definitional theory can be explained by exemplar theory (Smith & Medin, 1981, p. 150-151). Assuming a featural approach, such as in the context theory, *disjunctive* classes are an implicit part of an exemplar representation. *Unclear* cases are explained by, for example, failure to retrieve a sufficient number of exemplars from the relevant category. Features may not necessarily be shared by all exemplars in a category, thus features are not *definitional.* However, some features may be present in most exemplars and thus, they explain the importance of *non-necessary* features. *Typicality* effects are explained in exemplar theory by the presumption that some exemplars share more of their features with other exemplars in a category. These exemplars are judged more typical, and stimuli with features similar to these exemplars should retrieve a sufficient number of exemplars in that class more quickly (Smith & Medin, 1981, p. 150). Because of abstraction, not all exemplars of a superclass are represented as members of all their subclasses. Thus, a less typical member of a class may be represented as an exemplar of a superset, but not explicitly as a member of its subset. This can explain the finding that some exemplars are judged more similar to their distant superset than to their immediate subset, used to argue against the assumption of *nested* categories in definitional theory.

Some of the problems caused by the lack of constraints on prototype theory are equally relevant to exemplar theory. However, we have shown above how some formulations of exemplar theory show sensitivity to both correlated stimulus dimensions and to the distribution of exemplars within a class. Thus, exemplar theory is often considered to be even less constrained than prototype theory. Some of the more specific criticism of exemplar theory comes from empirical research on base rate effects in human classification[1] (e.g. Estes et al., Gluck & Bower, 1988a, 1988b; 1989; Medin & Edelson, 1988; Shanks, 1992). Most of these effects have been modelled by connectionist models of classification, leading some researchers to suggest that error driven learning should be a necessary property of a classification theory (Kruschke, 1993a; Medin & Florian, 1992, p. 230).

We now turn to the explanations of the general principles of connectionism in order to provide the necessary and sufficient terms and concepts to apply these models to a classification task like financial diagnosis.

3.2 Connectionist theory

Even though connectionism only recently has gained massive attention in cognitive science, many of the main ideas were formulated some years ago. Several versions of the history of connectionism exist; some entertaining (e.g. Papert, 1988), some popularised (The economist, 1987) and some more formal (e.g. Cowan & Sharp, 1988). By using the term "new connectionism" (Quinlan, 1991), some researchers pay attention to the history of this theoretical perspective. Examples of early connectionist research are parts of the work by Hebb (Hebb, 1949) on learning, but the main supplier of early connectionist ideas was Frank Rosenblatt (Rosenblatt, 1958, 1962). His influential work on pattern recognition and perception was unfortunately almost forgotten after Minsky and Papert showed some main limitations of Rosenblatt's simple perceptron models (Minsky & Papert, 1969). Funding of connectionist research in the years that followed was scarce, but a few research communities continued their work on connectionist models. Two examples are the research performed by James A. Anderson (e.g. Anderson, 1977), and Stephen Grossberg (e.g. Grossberg, 1982). Even though some research was done on artificial neural networks, the renewed interest in connectionist models did not burst until the release of the two-volume documentation from the "PDP[2]" research project at the University of California, San Diego (McClelland & Rumelhart, 1986, Rumelhart & McClelland, 1986).

[1] Other effects causing problems for exemplar theory are mentioned by Estes (1994, p. 252), and Estes et al. (1989, p. 557), and most of them are summarised in Medin and Florian (1992, p. 214-229).

[2] Parallel distributed processing

With the PDP research project, the first applications of connectionist models in cognitive science became widely known. Examples of cognitive phenomena covered by this research project were aspects of language (McClelland & Elman, 1986, McClelland & Kawamoto, 1986), distributed memory (Rumelhart, Hinton & Williams, 1986) and schema representations (Rumelhart, Smolensky, et al., 1986), just to mention a few. However, several connectionist models had been developed on memory and other cognitive functions prior to this project (see Hinton & Anderson, 1981)

In both cognitive science and cognitive psychology, the application of connectionist theory is now widespread, and includes research on perception (e.g. McClelland & Elman, 1986), recognition (e.g. Zipser, 1990), classification and categorisation (e.g. Gluck & Bower, 1988a, 1988b; Kruschke, 1992), judgement (e.g. Grossberg & Gutowski, 1987), choice and decision making (e.g. Usher & Zakay, 1993), reasoning (see Levine & Aparicio, 1994), problem solving (e.g. Hampson, 1990) and several aspects of language (e.g. Sejnowski & Rosenberg, 1986). As opposed to other theoretical approaches to many of these phenomena, parameter estimation in a majority of the models is done by learning. Thus, learning and knowledge acquisition are integral parts of the theory.

At least three terms are used[1] in the literature introducing connectionist theory. In cognitive science and cognitive psychology introductory literature on the subject (e.g. Clark, 1989; Quinlan, 1991), the term connectionism is used to present the "paradigm" under which connectionist theory is used. In parts of this area, specific processing and representational assumptions of connectionist models are made, and the term "parallel distributed processing" (e.g. Rumelhart & McClelland, 1986; McClelland & Rumelhart, 1986) is often used of these models. However, the most frequently used term in the area is "neural networks" or "artificial neural networks" (e.g. Fausett, 1994; Gallant, 1993; Hertz, Krogh & Palmer, 1991; Wasserman, 1989, 1993).

In this thesis, we limit the use of the term connectionist theory to models of cognition, based upon a set of connectionist principles to be introduced below. This view is consistent with traditional use in cognitive science, cognitive psychology and philosophy (Bechtel, 1993; Estes, 1994; Rumelhart & Todd, 1993). The term "parallel distributed processing" is used primarily by connectionist researchers advocating the importance of distributed representations in connectionist models, among who the most radical position has been taken by Smolensky (1988) (see e.g. Touretzky & Pomerlau, 1994).

---

[1] Sometimes interchangeably.

We will use the term "connectionism" to cover a set of models and theoretical contributions with the following characteristics in common:

- A *useful* source of knowledge to understanding cognition is knowledge of how the brain works.
- It will be possible to define an aggregate level of description for cognition, but an understanding of this level of description will be constrained by an understanding of the underlying subcognitive processes.
- Cognition is not the manipulation of passive symbols, but the result of large collectives of active processing units connected together.
- A cognitive system is made up of large collectives of units computing simple operations in parallel.
- Representation is not restricted to symbolic entities, but may be implicitly defined in the connections of the units of a system.
- The state of the cognitive system is a result of a process internal to the system. This process is the result of the output adapting to the external stimuli.
- The connections of the system are modified during the system's interaction with its operating world. This is how the system learns.

Connectionism as a concept covers more than a single theory. The term should be used similarly to the term "information processing theory" of a perspective or paradigm (Smolensky, 1988; Shanon, 1992).

Models using the principles stated below on functional components and operating principles in connectionist models *without reference to models of cognition*, we will term "neural networks" or "artificial neural networks". Consequently, these terms are applied to a wider and larger area of research spanning neural network research within, for example, statistics (e.g. Ripley, 1993; Cheng & Titterington, 1994), engineering (e.g. Dagli, Burke & Shin, 1992), medicine (e.g. Bassøe, 1995), finance (see Refenes, 1995), or economics (e.g. White, 1992), just to mention a few active applied areas.

In this chapter we give a brief description of connectionist theory, its functional components and its operating principles. We show some of the most widely known models in connectionism, and discuss connectionist models' relationship to artificial neural networks. We further discuss some of the most important aspects of cognitive modelling with connectionist theory, and present important connectionist models of classification.

3.2.1 General connectionist theory

Following the principles of Newell and Simon (1972) in their introduction of information processing theory, we present connectionist theory by first stating the *functional components* of a connectionist system. Next, we explain its *operating principles*. The presentation of functional components roughly follows Rumelhart and McClelland (1986). Other ways of characterising the models have been proposed, such as by the adaptive filter formalism (Carpenter, 1989), by learning principle (Hinton, 1989), by complexity (Zeidenberg, 1990) or by history (Cowan & Sharp, 1988). In addition to presenting the functional components and operating principles of connectionist models, we summarise these properties for some of the most widely used connectionist models.

A connectionist model is defined by combining specific functional components and operating principles in a particular proposed model. The scientific study of a system with a specific combination of these properties often constitutes a research project in connectionism.

The application of a connectionist model to a cognitive phenomena requires special attention to the environmental constraints provided by prior knowledge of the cognitive phenomena under investigation. When consideration for these constraints is implemented in the connectionist model, a connectionist theory of the cognitive phenomena is provided. Responses made by this model can be evaluated against established theory on the cognitive phenomena to provide a sufficiency test of the model.

3.2.1.1 Functional components

A connectionist system has the following eight functional components (Rumelhart & McClelland, 1986, p. 46):

* A set of active processing units.
* A state of activation.
* An output function for every unit.
* A pattern of connectivity defining the topology of the system.
* A propagating principle/rule for combining signals in the network that determine input to units.
* An activation principle/rule to produce the current state of activation.
* A learning principle/rule that changes the system's response based on experience.
* An operating environment that supplies the system with input, and provides a world for the system's response to take place.

The *processing units or processing elements* [1] of a connectionist system only loosely resemble properties of real neurons (Hebb, 1949). In a way they can be considered abstract neurons (Feldman & Ballard, 1982, p. 211; Tank & Hopfield, 1987, p. 62). The level of abstraction of these artificial neurons goes as far as to the functional level. This implies that units in a connectionist system remain units at a psychological level of explanation without necessarily being implemented in neural structure in a corresponding way (Smolensky, 1988, p. 9). The units are active because they can perform simple operations, and because they are not operated upon by some mechanism external to the unit.

In a connectionist system, units can be classified as *input* units, *hidden* units or *output* units. Input units receive signals from the system's environment and output units give the system's response. The organisation of input, hidden and output units defines the system's topology as explained below.

Each unit has a *state of activation*, and the overall state of activation for all units defines the system's state of activation. This is indicated by the *activation vector.* A unit's state of activation is normally computed as a function of its input. The legal states of activation of a unit define the type of unit as either *discrete* or *continuous*. The simplest discrete activation vector is a binary vector.

The *output function* determines the output of the unit for a given state of activation. The simplest output function is the direct function, in which output equates the state of activation. This output function limits the operations of the system, and more complex output functions are normally used. Threshold functions may be defined in various ways. The *binary* threshold function used in the original "perceptron" (Rosenblatt, 1962), gives zero as output for activation states below the threshold, and one otherwise. The *bipolar* threshold function of the "adaline" (Widrow & Hoff, 1960) gives minus one for activations below the threshold, and plus one otherwise. Other nonlinear output functions such as the *sigmoid* function of the traditional "backpropagation network" (Rumelhart, Hinton & Williams, 1986), are also used. In, for example, the "Boltzman machine", a *stochastic* output function is used (Hinton & Sejnowski, 1986).

The *topology* of the connectionist system is determined by the connections between units. One way of organising these connections is by arranging units in layers, and by having rules that determine how layers can be connected. The simplest systems[2] have only two layers; an input layer and an output layer. Systems of this kind are severely limited in the number of

---

[1] Also referred to as neurons, nodes, or processing elements. In the following we will generally use the term units.
[2] In terms of topology.

functions they can implement (see Minsky & Papert, 1969). Between the input and output layers, a number of hidden layers may be implemented. The hidden units of these layers can take care of interaction effects in inputs, and form intermediate abstractions representing such effects. In simple *feedforward* networks (e.g. Rosenblatt, 1962), the connections are allowed to pass signals from a layer below to the layer in question. In *recurrent* networks (e.g. Elman, 1990), the output of a unit is folded back to units in layers below the layer in question. *Competitive* networks typically allow connections between units in the same layer (e.g. Rumelhart & Zibser, 1986). Combinations of different connection types allow complex topologies to be modelled.

The connections can be either *inhibitory* by having a negative weight, or *excitatory* by having a positive weight. Thus, competition between units in a layer can be introduced with inhibitory connections (e.g. Rumelhart & Zipser, 1986). The connection weights are used for representation in the connectionist system. The set of weights are often shown in matrix form and the weight matrix then defines the representational properties of the network[1].

Output of units is propagated through the network as input to other units. The *propagation rule* is generally very simple, such as when the net input to a unit is the weighted sum of the outputs of units connected to the unit in question. However, more complex rules in which net input is a product function of outputs and weights (Peng & Reggia, 1989), or rules where excitatory and inhibitory weights are treated differently, also exist.

The *activation rule* controls the computation of the current activation state of a unit. In some models, activation is just a function of the net input coming in to the unit at time t. In other models, the activation of a unit at time t is a function both of the net input at time t and the activation of the unit at time t *-1*.

The *learning rule* of a connectionist network determines the way topology and connection weights change as a function of the system's operation in its environment. Topological change can be done by changing a weight value to zero; often termed "pruning", or from zero to a particular value. Thus, topological change is a special case of general learning. *Supervised* learning is when the system has an explicit "teacher" to tell if the output of the system is correct. The generalised delta rule or backpropagation rule (Rumelhart, Hinton & Williams, 1986) is a learning rule for supervised learning. If there is no "teacher", the system can still learn how to structure its representation to structure in its environment. Such learning is called *unsupervised*. The learning principle of the ART theory of Carpenter and Grossberg (1987, 1990) is an example of unsupervised learning.

[1] Even though these representations may not be fully interpretable without reference to processing.

The connectionist system interacts with its *environment* by receiving signals at its input units, and by giving response at its output units. Signals from the system's environment are received by letting the input units have their state of activation "clamped" by the environment. In this way the environment is represented by patterns of signals with presumed stable probability distributions, at least during the period of information processing. Similarly, activation of output units is interpreted as system response. However, the environment constrains the connectionist models by more than input and output. The environmental constraints are given by the task environment and application area of the system, and theoretical and empirical knowledge of the cognitive phenomena under investigation must be incorporated into the model. This important aspect of the environment is more thoroughly treated in section 3.2.3.

3.2.1.2 Operating principles

The *dynamics* of a connectionist system defines its operation principles. These principles are the second part of a functional description of a connectionist system. The principles are explained in two phases, the *recall phase* and the *learning phase*. We first explain the dynamics of the recall phase, then turn to the learning phase.

Depending on the relationship between input and output, a connectionist system can be either *heteroassociative* or *autoassociative*[1]. A heteroassociative system traditionally has the simplest operating principles. In such a system the typical input is not a part of the system's output, but differs from it. This means the system is instantiating a "cognitive function" (Cummins, 1989) in which the typical stimulus-response schema can be applied. Such instantiations take the form of a mapping function. In an autoassociative system, input and output traditionally are of the same pattern. Such a system may be used for several purposes even though the stimulus-response schema can not be applied. Some examples are the completion of a pattern, or the creation of a compressed representation of the stimulus (Chalmers, 1990).

A heteroassociative system is normally implemented in a *feedforward* network. In this network, a stimulus is presented at the input units by clamping the activation vector at the input layer. Clamping can be performed with continuous or discrete activation values. This leads to an output at each input unit defined by the output function of each unit. Output now

[1] The terms used here are traditionally used for different connectionist memory types. An autoassociative memory recalls by folding the input back upon itself. This is traditionally done by allowing feedback in the system until the system stabilises. In a heteroassociative memory what is recalled is different from the input and traditionally called a response or inference. This is possible without letting the network feed signals backwards through the recall phase. As a consequence of these differences of operation principles the terms are now also being used to classify the networks themselves (Zeidenberg, 1990).

propagates through the connections of the system, normally from the input layer to succeeding layers. In a simple system, this means direct propagation to the output layer, and in more complex systems this means propagating outputs to the (or the first of several) hidden layer(s). The topology of the system, determined by the weights with a value different from zero, gates the propagated output. At the receiving units, the propagation rule calculates the net input at each unit, and the activation rule determines what state of activation this net input gives the units of the layer in question. At this layer, the corresponding output function is initiated and outputs are propagated further until the output layer is reached.

In an autoassociative system, the recall phase is traditionally more complex. This complexity is partly caused by the fact that input and output no longer are separated in a clear cut way. Input may be part of the output, or input may be folded back onto itself. In most cases, the response of an autoassociative network is a completed pattern. However, pattern completion may also be achieved by a feed forward network where input and output patterns are equivalent. What further traditionally distinguishes an autoassociator from a heteroassociative system, is that the system often incorporates a feedback mechanism for the propagated signals, a property typical for the *recurrent* networks traditionally used to implement autoassociation. This means the signals are propagated in both directions through layers, but it also means propagation between units within the same layer is allowed. Since such propagation is allowed, the *time* dimension is particularly relevant in these networks. The number of times propagation between layers and between units in the same layer is allowed, must be controlled during recall and learning. Traditionally this control is performed by some constraint satisfaction principle, stopping recurrent propagation when a criterion is satisfied. Such systems also often have a complex propagation rule, even though the principle itself is quite simple. During recall, the system searches through its representational space to find a state that satisfies some specified criterion. The traditional criterion is some measure of the match or mismatch between the current representational state and the state imposed upon the system by its clamped input units. This measure is often termed the system's energy (Hinton & Sejnowski, 1986; Tank & Hopfield, 1987) or with a sign reversal, its harmony (Smolensky, 1986). Thus, recall in recurrent networks is often interpreted as constraint satisfaction. An interesting property of some connectionist systems is that the search for a global optimal match can be performed by letting the units' states of activation be determined completely locally.

Of the two learning principles, *supervised* learning gives the simplest dynamics of the system during the learning phase. In supervised learning, the system will have its connections changed as some function of the difference between the response of the system and the correct response. Several learning rules can be applied to decide how this change should be done. Easy learning is when the system has no hidden units. Many simple learning rules can be applied to such learning, but when the system has one or more layers of hidden units, the learning becomes hard. The reason for using the term "hard" is the so called "credit assignment" problem. The problem is how to find out which connections to change as a function of the system error. One way of solving this problem in supervised learning is by propagating the error backwards through the system and make "stepwise" corrections of the connections between each layer. This is what is done in the backpropagation algorithm (Rumelhart, Hinton & Williams, 1986). A special case of supervised learning is when the teacher has a repertoire limited to a nominal value, the simplest being the binary "right" or "wrong" values. The limited "teaching" input in such reinforcement learning makes correlational approaches to the "credit assignment" problem possible when there are few connections, but as the number of connections increases, the efficiency of the algorithm decreases dramatically (Hinton, 1989).

In *unsupervised* learning there is no teaching signal. This means that the learning principle must develop some organisation based upon its inputs only. The traditional Hebbian learning rule (Hebb, 1949) used without a teacher, is often applied to such situations. In principle, it states that the weights of the connections should be changed as a function of the connecting unit's pre- and postsynaptic activities. Simply stated, this means that a connection should be strengthened if both its presynaptic and postsynaptic units are strongly active. Variations of the Hebbian learning rule are among the most widely studied unsupervised learning rules (see Hertz et al., 1991, p. 197-215).

3.2.1.3 General connectionist and backpropagation models

The combination of different functional components and operation principles presented above, can give a very large number of specific connectionist models, but only a few of these models have been widely studied. An important criterion for the combination of functional components is that it must give the model some easily analysable properties. As an example, one wants the output function to be differentiable if the learning function is supposed to minimise some error in the system's response. Threshold functions must be excluded as output functions when this criterion is used.

As an example of how functional components and operating principles are combined in connectionist models, the properties of some early connectionist models are summarised in table 3.1.

Principle	Perceptron	Adaline	Madaline	Linear associator (BSB)
Reference	Rosenblatt, 1962	Widrow & Hoff, 1960	Widrow, Winter & Baxter, 1987	Anderson, 1977
Units	Several input and output units in one layer	Several input and one output unit in one layer	Several input and output units in one layer	Several input and output units in one layer
Activation	Binary	Linear	Linear	Linear
Output function	Linear threshold	Bipolar threshold	Bipolar threshold	Limited continuos
Topology	Fully/randomly connected hierarchically	Fully connected hierarchically	Fully connected hierarchically	Fully/random bidirectional inhibition
Propagation	Weighted sum	Weighted sum	Weighted sum	Weighted sum
Activation rule	Direct input dependent	Direct input dependent	Direct input dependent	Cycle dependent
Learning rule	Perceptron convergence rule	Widrow-Hoff rule	Widrow-Hoff rule	Hebbian learning or Widrow-Hoff rule
Recall operation	Feedforward heteroassociative	Feedforward heteroassociative	Feedforward heteroassociative	Feedback autoassociative
Learning operation	Supervised	Supervised	Supervised	Supervised

Principle	Selforganizing map	Hopfield network	Boltzman machine	Harmony theory
Reference	Kohonen, 1977	Hopfield, 1982	Ackley, Hinton & Sejnowski, 1985	Smolensky, 1986
Units	Several input and output units in one layer	Several input and output units in one layer	Visual and hidden units in two or several layers	Visual and hidden in one layer
Activation	Linear	Linear	Linear	Linear
Output function	Competitive normalised	Binary threshold	Bipolar and binary stochastic	Bipolar and binary stochastic
Topology	Fully with lateral connections	Fully bidirectional	Fully bidirectional symmetric	Fully bidirectional
Propagation	Normalised weighted sum	Weighted sum	Weighted sum	Weighted sum
Activation rule	Cycle dependent	Cycle dependent	Direct input dependent	Direct input dependent
Learning rule	Kohonen rule	Hopfield rule	Two phase Boltzman	Trace learning
Recall operation	Feedback autoassociative	Feedback autoassociative	Feedback autoassociative	Feedback autoassociative
Learning operation	Supervised	Supervised	Supervised	Supervised

Principle	Backpropagation network	Radial basis function network	Competitive learning	Adaptive resonance theory
Reference	Rumelhart, Hinton & Williams, 1986	Moody & Darken, 1989	Rumelhart & Zipser, 1986	Carpenter & Grossberg, 1987
Units	Input, hidden and output in two or several layers	Input, hidden and output units in three or several layers	Input units in one layer, and one hidden layers of units	Complex two system organisation
Activation	Linear	Distance and linear	Linear	Linear
Output function	Continuos sigmoid	Gaussian and sigmoid	Binary threshold with competition	Binary threshold and linear layer dependent
Topology	Fully hierarchical	Fully hierarchical	Fully connected with inhibiting clusters	Fully connected with special gain control and vigilance units
Propagation	Weighted sum	Distance and weighted sum	Weighted sum	Weighted sum
Activation rule	Direct input dependent	Direct input dependent	Direct input dependent	Cycle dependent
Learning rule	Backpropagation	Combined rules	Competitive learning	ART learning rule
Recall operation	Feedforward heteroassociative	Feedforward heteroassociative	Feedforward heteroassociative	Feedback autoassociative
Learning operation	Supervised	Un- and supervised	Unsupervised	Unsupervised

Table 3.1. Selected early connectionist models

The selection of models is restricted to the few models most widely studied. Among other models of considerable interest, but which are specifically connected to a particular researcher or research group are, for example, the counter propagation network (Hecht-Nielsen, 1987), the bi-directional associative memory (Kosko, 1987), the cognitron and the neocognitron

(Fukushima, 1975), and the ART2 and ART3 models (Carpenter & Grossberg, 1987, 1990), just to mention a few. The reinforcement learning based networks (Barto & Anandan, 1985) are worth mentioning as well as models which have been developed in areas peripheral to the traditional connectionist field, such as the Darwin models of Reeke and Edelman (1988), and work on genetic algorithms (Goldberg, 1989).

Some of the models, such as the counterpropagation network and the radial basis function networks, are hybrid models which combine different functional components and operating principle for different parts of the system. In, for example, radial basis function networks, the first hidden layer traditionally learns with unsupervised learning, while the rest of the hidden layers and the output layer learn by supervised learning. Several such hybrid models have been developed, and new suggestions are introduced in artificial neural network and connectionist research in every new issue of the research journals in the fields.

A connectionist model is traditionally illustrated in a simplified network structure. The most widely applied connectionist model, the multilayer perceptron often termed the "backpropagation network" is illustrated in figure 3.7.

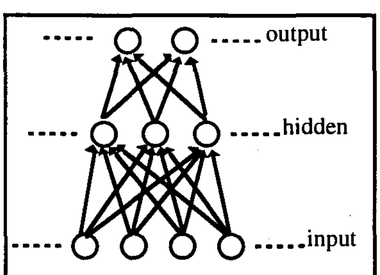

Figure 3.7 Backpropagation network with one hidden layer

The functional components and operating principles are indicated in table 3.1. Input is presented at the input layer and is propagated to the hidden layer[1], where it is summed and transformed by the output function. The outputs of the hidden units are propagated further to the output layer at which a similar summation and transformation is performed to produce the output of the system. An error is computed by comparing the system output to a target. This error is used by the backpropagation learning rule to change the weights of the connections so that error is minimised during further processing.

Since the backpropagation network is fundamental to our models of classification, the formal properties of its functional components and operating principles need further elaboration.
A backpropagation model is a feedforward network which applies the backpropagation learning rule (Le Cun, 1985; Parker, 1985; Rumelhart, Hinton & Williams, 1986; Werbos, 1974). For simplicity purposes this network is termed a backpropagation model. The network and algorithms are well presented in several textbooks (e.g. Gallant, 1993; Rumelhart & McClelland, 1986; Wasserman, 1989). In the following, we primarily use symbols from

[1] Several hidden layers may be used.

Rumelhart, Hinton and Williams (1986). The signal coming into a unit j of a backpropagation model when pattern p is presented is computed as:

$$net_{pj} = \sum_{\forall i} o_{pi} w_{ji}. \tag{3.5}$$

Here, p indicates the pattern presented and i indicates units in a layer passing outputs to units indicated by j. The term w_{ji} indicates weights of connections between units j and i, and o_{pi} the output of unit i when pattern p is presented. For the first hidden layer, o_{pi} is the input I_{pi} of unit i when pattern p is presented. A bias[1] is introduced by setting o_{p0} to 1.

The signal is transformed in an output function. Backpropagation requires that this function is continuously differentiable, and the function is asymmetric if input is scaled to [0,1] or symmetric if input is scaled to [-1,1]. The most widely used output function is the standard asymmetric sigmoid function on the form:

$$o_{pj} = \frac{1}{1+e^{-net_{pj}}}. \tag{3.6}$$

The feedforward pass follows the principles given above. For a model with one hidden layer, the transformed output of the hidden units is weighted, summed and transformed again to form the output of the network. From the output of the network, the error, E, is computed as:

$$E_p = \tfrac{1}{2}\sum_{\forall j} (t_{pj} - o_{pj})^2, \tag{3.7}$$

where t_{pj} is the target value of unit j when pattern p is presented. For all patterns presented, the total error can be computed as:

$$E = \sum_{\forall p} E_p. \tag{3.8}$$

Learning is done in the network by adjusting the weights according to the backpropagation learning algorithm. This algorithm performs gradient decent on the error surface. To perform this gradient decent, we must compute the partial derivative of the error with respect to the weights. This is found using the chain rule. In general we have:

$$\frac{\partial E_p}{\partial w_{ji}} = \frac{\partial E_p}{\partial o_{pj}} \frac{\partial o_{pj}}{\partial net_{pj}} \frac{\partial net_{pj}}{\partial w_{ji}}. \tag{3.9}$$

---

[1] Also termed "threshold".

The first term in equation (3.9) is different for output and hidden units. For output units the first term is:

$$\frac{\partial E_p}{\partial o_{pj}} = -(o_{pj} - t_{pj}), \tag{3.10}$$

and the second term is

$$\frac{\partial o_{pj}}{\partial net_{pj}} = o_{pj}(1 - o_{pj}). \tag{3.11}$$

The product of the two terms of equations (3.10) and (3.11) is often termed δ, thus:

$$\delta_{pj} = \frac{\delta E_p}{\delta net_{ji}} = o_{pj}(1 - o_{pj})(t_{pj} - o_{pj}). \tag{3.12}$$

The third term of equation (3.9) is:

$$\frac{\partial net_{pj}}{\partial w_{ji}} = o_{pi}. \tag{3.13}$$

The change in the weights Δw_{ji} is generally:

$$\Delta w_{ji} = \eta \delta_{pj} o_{pi}. \tag{3.14}$$

Here, η is termed the learning coefficient or learning rate. The new weights are now computed as:

$$w_{ji}(t+1) = w_{ji}(t) + \Delta w_{ji}(t+1). \tag{3.15}$$

To smooth the weight changes, the previous weight change can be included in the calculation of the new weights:

$$\Delta w_{ji}(t+1) = \eta \delta_{pj} o_{pi} + \alpha \Delta w_{ji}(t), \tag{3.16}$$

where α is a smoothing parameter called the momentum term.

All equations above, except (3.10) and (3.12), also apply to weights of a hidden layer. However, units in a hidden layer do not have a traditional target. The terms in the chain rule for determining the partial derivative of the error with respect to the weights given in equation (3.9) are the same, but the formula for δ of a hidden layer is different. It can be shown (Rumelhart, Hinton & Williams, 1986; Smith, 1993) that:

$$\frac{\partial E_p}{\partial o_{pj}} = \sum_{\forall k} \delta_{pk} w_{kj}. \tag{3.17}$$

In equation (3.17), the subscript k indicates that δ_{pk} is a parameter of unit u_k where u_k is a unit in a layer above the layer of unit u_j. In this way, the term δ can be thought of as being propagated back from units in a layer above the layer in question. For a network with one hidden layer, the δ_{pj} calculated in the output layer using equation (3.12) is propagated as δ_{pk} in equation (3.17). Consequently, δ_{pj} of a hidden layer is calculated as:

$$\delta_{pj} = o_{pj}(1 - o_{pj}) \sum_{\forall k} \delta_{pk} w_{kj}. \tag{3.18}$$

Weight change in the hidden layer is performed using equations (3.14) to (3.16) above. For networks with more than one hidden layer, the principles of equation (3.18) is followed correspondingly.

Since backpropagation should perform gradient decent in E and:

$$\frac{\partial E}{\partial w_{ji}} = \sum_{\forall p} \frac{\partial E_p}{\partial w_{ji}}, \tag{3.19}$$

we should not adjust the weights on each iteration. However, it can be shown that small η gives only a small departure from true gradient decent in E when example based learning is used[1] (Rumelhart, Hinton & Williams, 1986). Weight adjustment after all patterns have been presented, is termed epoch learning.

Several modifications have been suggested of the functional components and operating principles of the backpropagation model described in this section, such as modifications of the learning rule by using second order methods (e.g. Johansson, Dowla & Goodman, 1992), or the introduction of pruning during learning (e.g. Weigend, Rumelhart & Huberman, 1991), just to mention a few.

[1] Example based learning means weight adjustments are performed after each example has been presented to the model.

3.2.2 Connectionism and artificial neural networks

In the introduction to this section, we reserved the terms "neural networks" or "artificial neural networks" to research based upon the functional components and operational principles of connectionism *without* reference to cognitive models. There are mainly four reasons for this interest in non-cognitive neural networks.

First, much research on neural networks is related to the simulation of biological neural networks (see e.g. Gluck & Rumelhart, 1990). Many functions performed by biological neural networks are surely not cognitive functions, but may, nonetheless, be of significance to the understanding of human and animal behaviour. In this case, the models are not psychological, but models at the implementational level. Much original research within the neural network field results from this perspective. Our formulation of connectionism relates to models at the molar level. Thus, we make no claim of the biological relevance of our models. However, we clearly make claims that our models are cognitive models. We will not pursue, or refer to, the research on biological neural networks other than when findings are easily transferable to the molar level.

The other three reasons why large amounts of research on neural networks, without reference to cognitive models, are found, lies in the mathematical and statistical properties of artificial neural networks as modelling frameworks. First, artificial neural networks provide a general framework for describing and representing statistical and mathematical models (see Cheng & Titterington, 1994). Second, neural networks provide a rich set of functional forms, and third, neural networks provide a similarly rich set of estimation methods for parameters of these functions. We will elaborate somewhat on these three elements of artificial neural networks.

To illustrate how neural networks can provide a general framework for describing mathematical and, in particular, statistical models, consider a simple single layered[1] perceptron with several input units and one output unit. This model is illustrated in figure 3.8.

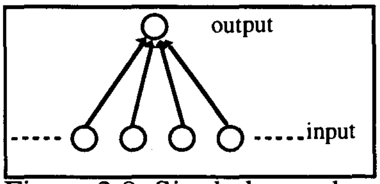

Figure 3.8. Single layered perceptron

If the single layered perceptron has continuous inputs, a direct output function at the input and output layers, and the activation at the output layer is computed using equation (3.5), then the model has a very familiar structure. The model then performs a functional mapping similar to the simple linear regression model. Traditionally, the parameters of this model are estimated by

[1] Here, the single layer refers to the single layer of connections.

analytic least squares minimisation (OLS), but the parameters could also be estimated by gradient descent. The general framework of neural networks can also be used to illustrate and formalise several other statistical models (see Cheng & Titterington, 1994; Kuan & White, 1994; Ripley, 1993; White, 1989). Some examples are neural network implementations of linear and quadratic discriminant analysis, logistic regression, and principal components analysis (Cheng & Titterington, 1994, p. 5).

The simple linear regression function implemented in the network of figure 3.8 may also indicate that neural networks incorporating hidden layers, modularity, recurrency and so on, can provide a whole new set of "statistical" models. Only some of the standard neural network configurations have been investigated as mathematical or statistical models, but the backpropagation model, some self-organising networks (Kohonen, 1977, see Oja, 1989), and some radial basis function networks (Moody & Darken, 1989; Poggio & Girosi, 1990) have received considerable attention among statisticians. It has been shown that some of these networks are neural network implementations of known statistical models, while others provide new and unexplored statistical models.

It has also been shown that some multilayer perceptrons, such as the backpropagation model and some radial basis function neural networks under certain, not very restrictive assumptions, belong to a larger family of universal approximators[1] (Cybenko, 1989; Hornik, Stinchcombe & White, 1989; Park & Sandberg, 1991; Poggio & Girosi, 1990). Belonging to the same family are several other mathematical and statistical models, such as projection pursuit (Friedman & Stutze, 1981) and multivariate adaptive regression splines (MARS) (Friedman, 1991)[2]. To illustrate, the output of unit k in a backpropagation model with one layer of hidden units is:

$$o_k = F(I_{pi}, W) = f(\sum_{\forall j} g(\sum_{\forall i} I_{pi} w_{ji}) w_{kj}), \tag{3.20}$$

where, F illustrates that output is a function of input and weights of the network only, and f and g are some, traditionally similar, nonlinear functions, such as the sigmoid function. The main differences between, for example, the family of functions in projection pursuit and F of equation (3.20), are that in projection pursuit, f is linear and g is unknown. Situations in which f is a linear function have been thoroughly investigated by White (see Kuan & White,

---

[1] Backpropagation networks with one hidden layer can approximate any *continuous* function, while backpropagation networks with two hidden layers have been show to be generally universal. An open question in artificial neural network research is, however, the degree of approximation for the different models. One question of particular interest is if there are situations in which a network with more than one hidden layer has equally good approximation properties with fewer parameters (weights) as a network with only one hidden layer (see Kuan & White, 1994, p. 10-11)

[2] See Ripley (1993), p. 107-108, or Geman et al. (1992), p. 6, for other examples of universal approximators.

1994), demonstrating the similarities between the two families of approximators. Consequently, artificial neural networks provide new functional forms, and some of these functions have the capacity of universal approximation.

The impressive approximation abilities of certain neural networks are, however, of limited value if a method to estimate the parameters can not be found. Traditional analytical least squares minimisation and log likelihood maximisation methods are insufficient in these models (Kuan & White, 1994). Fortunately, for some neural networks, methods for estimating the unknown parameters of models with universal approximation properties have been developed. For certain feed forward neural networks with one layer of hidden units, a set also containing the backpropagation network, White (1990) has shown how gradient descent on mean squared error can be used to estimate the unknown parameters. Several methods of gradient descent on mean squared error exist. The traditional formulation of Rumelhart, Hinton and Williams (1986) is in principle sufficient[1], but several new methods and modifications have been suggested (Fahlman, 1989; Jacobs, 1988; Johansson et al., 1992). However, these methods share the principles of error minimisation by gradient descent, suggesting that the third reason for the interest in artificial neural network research lies in its methods for parameter estimation in complex nonlinear models.

Kuan and White (1994) summarise the last two reasons for the growing interest in artificial neural network research in the following way:

> *Thus, in addition to introducing us to an interesting new class of flexible function forms, the artificial neural network field has drawn our attention to a remarkably simple estimation procedure for complex models, of interest in its own rights. (Kuan & White, 1994, p. 19)*

In addition to the property as an illustration and formalisation framework, the properties of artificial neural networks referred to above have led to a considerable number of applications, as well as a growing interest for artificial neural networks in the mathematics and statistics communities.

The favourable properties of neural networks as modelling frameworks do, however, not come without a cost. Of these costs we will briefly mention two problems of significance to connectionist modelling, but which have been paid little interest in cognitive science applications of the models. The first of these problems is how to determine the generalisation properties of the model. The second is closely related to generalisation, and results from the

---

[1] By letting the learning parameter decrease over learning time. See Kuan and White (1994, p. 18-19).

fact that we, so far, only have considered networks of a fixed complexity. This problem is the determination of the complexity of the mapping function provided by the model.

All models, but in particular flexible functional forms, are in danger of *overfitting* the sample data. Thus, the estimates of prediction or classification error based upon the sample is under-estimated (Cheng & Titterington, 1994, p. 20; MacKay, 1992, p. 451). Most often, we assume that the size and distribution properties of the sample prevent this problem from occurring. The property of universal approximation is an asymptotic property, based on sufficiently large samples. How large a sufficiently large sample is, however, is highly context dependent (White, 1989, p. 110). Some suggestions say the number of weights times 100 is sufficient (White, 1989), while others argue that the number of weights times 10 is enough (Baum & Haussler, 1989). In contrast to these rather large sample size requirements stand several practical applications showing surprisingly good generalisation properties with smaller samples (see Wasserman, 1993, p. 229)

The second, and closely related problem, is that what constitutes a sufficiently large sample does not only depend on the nature of the sample and the number of inputs, but also on the *complexity* of the artificial neural network. For a multilayer perceptron, the traditional way of regulating complexity is by adjusting the number of hidden units and hidden layers. We concentrate here on the number of hidden units. It is obvious that the dimensionality of W in equation (3.20) and thus, the number of free parameters, is partly determined by the number of hidden units. By increasing the number of hidden units, the flexibility of the function increases and the danger of overfitting increases (Cheng & Titterington, 1994, p. 20; Geman, Bienenstock & Doursat, 1992; Ripley, 1993; Smith, 1993). Furthermore, the need for large samples is greater the more hidden units the network contains, if one is to get "good" approximation and prevent overfitting.

To find the optimal complexity of the model, two principal solutions can be used. One is to start with less complexity and increase it, while the other solution starts with high complexity and reduces it. The first solution is used in *constructive* methods (e.g. Fahlman & Lebiere, 1990). In these methods, one typically starts with a small number of hidden units, and introduces new units when necessary. The second solution refers to reduction in both hidden units and in individual weights, and is the use of *pruning* methods (e.g. Karnin, 1990; Weigend et al., 1991). These methods traditionally introduce a complexity penalty on the error measure, and the gradient descent method performs error minimisation on the combined error measure. Since both solutions traditionally rely on sample data in their determination of optimal complexity, they neglect the interaction between the sample size and complexity determination problems mentioned above.

Fortunately, resampling methods derived from modern nonparametric statistics (e.g. Efron & Tibshirani, 1993), can be used to overcome the two problems simultaneously. Of particular interest are methods that separate estimation sample[1] and test sample. Simple "intuitive" methods (White, 1990), do this separation once, but N-fold cross validation can provide better estimates of prediction and classification error (Efron & Gong, 1983, p. 37; Efron & Tibshirani, 1993, p. 237-255; Stone, 1974). The idea is that parameter estimation is based upon the N "leave-one-out" samples, and tests of generalisation properties (generalised prediction or classification errors) are based upon the N "left-out" observations. The measure of generalisation error in this procedure is traditionally termed "cross validation error". Several authors recommend this measure for assessing the generalisation abilities of artificial neural networks (Cheng & Titterington, 1994, p. 20; Moody, 1993; White, 1990, p. 539). The measure has been refined and several versions of the cross validation principle exist (see Moody, 1993; Moody & Utans, 1995). These procedures are computation intensive, and if combined with constructive or pruning methods, the demands on computational power increase even more. A method based upon complexity determination in the learning sample with constructive methods, and assessment of generalisation ability with cross validation, has been developed by Moody (Moody & Utans, 1995). In this thesis we independently develop a method in which both complexity and measures of generalisation properties are determined with the use of cross validation. This is done by finding the connectionist model with the best generalisation properties while the models "grow" in complexity.

To summarise, artificial neural network research has much to offer cognitive science research using connectionist models. The comparisons of neural network models with more traditional statistical methods have greatly increased the understanding of properties of artificial neural networks, and their advantages and limitations. Some of these limitations are of particular significance to connectionist modelling, such as the problems of sample size and complexity determination. The methods traditionally used to overcome these problems have largely been unattended in connectionist modelling; a rather unfortunate situation[2].

Still, it is important to bear in mind that research in artificial neural networks is not constrained by biological or cognitive plausibility, and the uncritical application of models and methods from the artificial neural network community to cognitive modelling is not recommended. This point has also been expressed by artificial neural network researchers, such as White (1989):

---

[1] In artificial neural network and connectionist terminology this sample is termed "training sample" or "learning sample".

[2] One example is the problem of overfit treated above. In Estes et al. (1989) the authors state: "The exemplar model does better, but its account of the test data does not come close to the accuracy with which either model can account for learning data" (p. 569). This may well be due to overfit, a problem not considered by the authors.

To the extent that biological or cognitive processes or constraints suggest useful approaches to learning, we are free to adopt them. To the extent that such processes or constraints get in the way of using an artificial network to encode empirical knowledge, we are free to dispense with them. (White, 1989, p. 91)

When this principle is followed by artificial neural network researchers, cognitive scientists are recommended to carefully evaluate the vast amount of models and algorithms provided within the field. Successful adoptions and refinements of models developed within the neural network community have also been done by cognitive scientists (e.g. Kruschke, 1992). Even though the impression may have been created here, that new models mainly are developed within the artificial neural network community only, many of the main developments within the field have come from cognitive scientists.

3.2.3 Connectionist modelling and environmental constraints

In section 3.2.1, we laid out the eight functional components of a connectionist system. Seven of these are components of the particular connectionist model chosen, while the environmental component is determined by the *task environment and application area* of the model. The environmental component provides the constraints on the formulations of the connectionist model, and the basis for testing its validity. Theory and empirical research results on the particular cognitive phenomena under investigation represent constraints on the connectionist model. Above, we saw how neural network research was not constrained by these environmental constraints, but connectionist modelling definitely is.

The constraints are highly relevant when cognitive functions are modelled. Operationalisation is not done directly from theory to empirical measures, but by models when connectionist theories are tested. Thus, environmental constraints should be built into the modelling simulations. In addition, the empirical findings and the established theory of the investigated cognitive phenomena can provide a basis for additional evaluation of model adequacy. The environmental constraints are traditionally related to stimulus and input representation, to response and output representation, to internal representational constraints, and to constraints on the overall processing behaviour of the model. We will briefly explain the relevance of these constraints.

Two major constraints on the modelling of cognitive phenomena are the selection of relevant features or stimulus dimensions, and the selection of a proper representational form for these dimensions. In principle, the relevant stimulus dimensions are hopefully identified by task analysis, analysis of previous research and the established practices of the research tradition.

In, for example, financial diagnosis, the relevant stimulus dimensions are cues of the financial statement. However, task analysis may reveal other relevant stimulus dimensions previously believed to be irrelevant, such as when new aspects of the task context are taken into consideration. Thus, the principles of stimulus dimension selection are no different from traditional independent variable selection in other domains.

However, connectionist models typically differ from these domains in their requirement that a relevant *representation* of stimulus dimensions must be selected. Traditionally, two opposing representational designs are found. The situation is illustrated in figure 3.9.

When representations are *local*, there is a one to one correspondence between the number of stimulus dimensions and the number of units used to represent them (Hanson & Burr, 1990; Hinton, 1989; Sharkey, 1991). When *distributed* representations are used, the stimulus dimensions are distributed over a (traditionally larger) set of units. For input representations this is often termed place coding (see Kruschke, 1993b, p. 28), or coarse coding (see Hinton, McClelland & Rumelhart, 1986, p. 92)[1]. A local representation of the two stimulus dimensions in figure 3.9 consists of two units representing the value of each of the stimulus dimensions one and two, respectively. A distributed representation is illustrated with the "receptive fields" of four units. These units have an activation illustrated by the "isoactivation contours". Activation along the contours depends on the combined value of the two stimulus dimensions of the object represented. Distributed representations are thus less sensitive to distortion in the stimulus, and are often assumed to have implicit generalisation properties (Clark, 1989, 1993; Hinton et al., 1986).

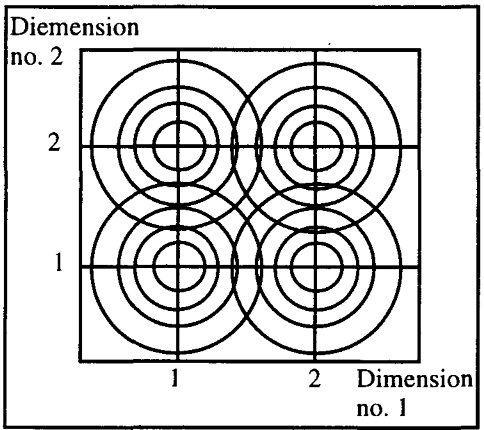

Figure 3.9 A distributed (coarse coded) representation of two stimulus dimensions.

For output representations there are similar task and theory constraints as for the input representation, but the question if distributed representations should be used is not of similar relevance[2]. One obvious constraint is that the response should be made by a cognitive system, such as a human being. In addition, the response should be given in forms interpretable as relevant response to the task. This presumption limits the relevance of artificial response

[1] Other terms, such as "superpositional storage" have also been used on memory systems using distributed representations (see Clark, 1993).

[2] Distributed output representations may be used, but these must always be translated to local responses. This translation may well be performed by an additional output layer above the distributed representation.

forms not ordinarily used within the task context. In classification tasks, output representation can be interpreted as a class choice, or as class posterior probability. With the last interpretation, the model approximates a posterior probability classifier (see Lippmann, 1993). When classification is measured by a continuous response value, output representation can be interpreted as expected continuous response given the stimulus presentation, and the model approximates an expected value as, for example, a regression model does (Cheng & Titterington, 1994; Smith, 1993).

An interpretation of model response must be made with reference to relevant responses identified by task analysis, and then, the traditional elements of validity, reliability and measurement error applies to these responses as for any other dependent variable.

A major difference between connectionist models and linear, or other models within a stimulus-response paradigm, is the development of an internal representation in connectionist models[1]. As shown in equation (3.20), the internal representation consists of a weight matrix. However, these weights are interpretable only with reference to unit activation and outputs because representational units in connectionist models are not necessarily conceptual[2]. Traditionally, they are interpretable as subconceptual units (Rumelhart & McClelland, 1986; Smolensky, 1988), and such interpretation must take place during processing. In principle, internal representations may be conceptual, and correspond to local representations as explained above. However, distributed internal representations are more typically developed in connectionist models because of their representational capacity and implicit generalisation properties (see Sharkey, 1991). Connectionist models are indirectly interpretable in conceptual terms by analysis of weights and unit outputs during processing. Several traditional methods are applicable to this interpretation, such as principle components analysis and cluster analysis. In addition, a whole set of illustration principles, such as Hinton diagrams (Hinton, 1989)[3], and analysis methods have been developed to facilitate transformation of connectionist representations to allow conceptual interpretation (Clark, 1993, p. 41-67; Hanson & Burr, 1990; Gorman & Sejnowski , 1988; Sanger, 1989; Sejnowski & Rosenberg, 1986). Special attention has been paid to the interpretation of internal representations in connectionist models as rules of traditional information processing models (see Gallant, 1993, p. 315-328; Towell & Shavlik; 1993). As is evident, connectionist models can show rule-following behaviour, but the interpretation of single units and weights as representing these rules may be impossible (Smolensky, 1988).

---

[1] Kohonen (1995) has even proposed this as an exclusive property of connectionist models: "Only neural networks are able to create new information processing functions, such as specific feature detectors and ordered internal representations for structured signals, in response to frequently occurring signal patterns. Also, only neural networks can create higher abstractions (symbolisms) from raw data completely automatically" (Kohonen, 1995, p. 57).

[2] See the distributed input representation shown in figure 3.9.

[3] For an application, see Bremner, Gotts and Denham (1994).

Connectionist models' internal representations are often uninterpretable without reference to processing. This illustrates the important relationship between processing and representation in such models. A similar relationship between representation and processing is not unknown in previous cognitive theories (e.g. Kosslyn, 1980), but is at odds with traditional assumptions of information processing theory (Newell, 1990; Pylyshyn, 1984). To say that processing and representation are integral parts of connectionist model behaviour, is not equivalent to saying that general processing assumptions independent of representation does not exist . The difference between information processing theory and connectionist theory is that the processing in information processing models is independent of the conceptual interpretation of internal representations[1]. In connectionist models it is assumed that processing is independent of the subconceptual representations (Pedersen, 1988; Smolensky, 1988)[2]. Thus, connectionist models give predictions that only allow simultaneous evaluation of processing and representational assumptions, but constraints on these assumptions help us evaluate their realism. These representational assumptions can be tested by applying analysis methods of connectionist models' representations (Hanson & Burr, 1990; Sanger, 1989; Sharkey, 1991), and by comparing the developed representations to competence theory or other established knowledge in the task domain.

Processing assumptions must be made, but can be evaluated by comparing the model's generalisation ability and its error production to other models and findings in the task domain. An example of how environmental constraints can be used to evaluate these aspects, is by comparing model predictions after changes in stimuli have been made to actual responses under similar conditions. In, for example, classification tasks, biases and errors found from observation of behaviour (e.g. Tversky & Kahneman, 1974) can be used to evaluate if similar errors and biases are produced by the model under similar environmental conditions. This strategy has been extensively used in categorisation and classification research by connectionist modellers (see e.g. Gluck & Bower, 1988; Shanks, 1991, 1992). These evaluations are performed within the framework of operationalisation by model. The framework helps us to show several important aspects of theory testing in connectionist research, and is illustrated in figure 3.10

[1] They are syntactic.

[2] A similar view on recent research on cognitive phenomena in general has been presented by Estes (1993).

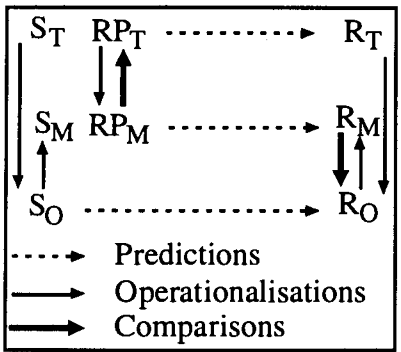

Figure 3.10
Operationalisation by model

The model in figure 3.10 has three levels, the conceptual or theory level (T), the model level (M) and the observational level (O). As illustrated, stimulus (S) and responses (R) are operationalised at the observational and at the model level, but representation and processing (RP) are only theoretically assumed and operationalised at the model level. Predictions can be made at all three levels. Responses are evaluated by comparing model predictions to observed behaviour. Representation and processing assumptions are evaluated by extracting representation and processing principles from the model during processing, and then by comparing these to theory on the task not explicitly assumed in our original operationalisations[1]. In addition, the model can be evaluated against observed behaviour with variations in stimulus conditions when theory predicts a specific behaviour. If the model generalises under such conditions, the model is strongly supported.

Two other aspects of connectionist models are often used to evaluate their validity. Connectionist researchers often assume that the parameter estimation procedure used, also models the learning aspects of the task (Kruschke, 1992, 1993a, 1993b). Thus, if learning data are collected, they can be compared to model predictions during learning to evaluate the model's course of learning. If a proposition is only made of the estimated model's behaviour at the molar level, the course of learning is irrelevant as a test of the generalisation abilities of the model. The second aspect used to evaluate connectionist models[2] is how it scales to larger and more realistic size (Clark, 1989). Some models make assumptions that do not scale well. One example is the configural cue model of Gluck (1991), which assumes that a representational unit exists for any conjunction of cues. For multidimensional stimuli, this assumption does not scale well.

These important principles of connectionist modelling of cognitive functions have been summarised by Seidenberg (1993):

> *Rather, one starts with a set of principles concerning learning and the representation of knowledge. If the principles are identified correctly, modelling should merely involve incorporating domain-specific variables such as different types of stimulus inputs, motoric responses, and learning experiences. The relevant generalizations about the*

---

[1] Notice how this evaluation is made impossible in information processing theory operationalisation by model, because representations must be explicitly formulated in the model. Thus, connectionist models provide a two way validation not possible in information processing models.

[2] Typically against information processing models.

domain in question should then fall out of the model. That is, it will develop the correct sorts of representations, obviating the need to build them in by hand. (Seidenberg, 1993, p. 231)

The principles laid out by Seidenberg (1993), and referred to above, stress the importance of environmental, and thus, domain specific constraints on connectionist modelling. The suggestion is that "explanatory theories can be derived from general connectionist principles in conjunction with domain-specific boundary conditions" (Seidenberg, 1993, p. 231).

The environmental constraints on connectionist modelling of cognitive phenomena apply correspondingly to models of financial diagnosis. Constraints on input representations are given by available and diagnostic cues of the financial statement. As shown in chapter 2, some agreement exists on such cues. These cues are most conveniently represented locally, but distributed representations may be relevant. Output is constrained by the task contexts of financial diagnosis, but can also be represented by continuous responses or linguistic terms. Within the context of classification, distinct classes must be identifiable in the response material. Internal representation and processing assumptions are determined by the applied connectionist model. Model applicability is evaluated by testing the generalisation properties of the model to unseen cases. Other aspects of the model can be evaluated against general findings on behaviour in the financial diagnosis task. Internal representations and processing principles can be evaluated by comparing the representations developed in the model to established findings on knowledge presumed relevant to the financial diagnosis task. The representations are not readily available in conceptual terms, but must be interpreted using analysis methods developed for connectionist models (e.g. Hanson & Burr, 1990; Sanger, 1989). Competence theory derived from predictive studies and knowledge derived from cognitive processing studies of financial diagnosis may be a valuable reference point for such evaluations.

With this understanding of connectionist principles, functional components, models, their relationship to artificial neural networks and their role in modelling cognitive phenomena, we turn to connectionist models of categorisation and classification, and finally to their application to financial diagnosis.

3.2.4 Connectionist models of categorisation and classification

In section 3.1, we presented the three traditional theoretical approaches to categorisation and classification in cognitive psychology. Connectionist models have been developed of both prototype- and exemplar theory. In addition, connectionist models have been developed by transfer of models from other areas of connectionist and neural network research. In this section we review and discuss some of these contributions.

One of the first, and definitely one of the most influential, connectionist models of classification was the simple *adaptive network model* [1] of Gluck and Bower (e.g. Gluck & Bower, 1988a, 1988b). In a series of experiments, they tested the predictions of this single layered perceptron model of classification on a simulated medical diagnosis task. The basic architecture of this model is shown in figure 3.2. However, the model of Gluck and Bower has four input units and one or two output units varying across simulations. The model learns using the delta rule of Widrow and Hoff (1960), shown by Gluck and Bower to be equivalent to the Rescorla-Wagner learning rule for associative learning (Rescorla & Wagner, 1972). The output is converted to probabilities by using a sigmoid output function at the output units when learning is finished. In the experimental setup, subjects learned probabilistic classification of a rare and a common diagnosis. The succeeding simulations showed the capacity of the simple adaptive network to model two interesting findings in subjects' classifications, not easily explained by other models of classification. First, subjects showed base rate neglect when presented with an ambiguous symptom. Second, subjects judged the diagnosticity of a cue relative to other cues present in a situation. From these results, the modelling of base rate neglect has received considerable attention (Myers, Lohmeier & Well, 1994; Schanks, 1990). Replications and extensions of the Gluck and Bower (1988a, 1988b) study have been performed, showing superiority of the model to exemplar models (e.g. Estes et. al,. 1989), but also several of its weaknesses, as well as limitations in the conclusions drawn by Gluck and Bower (Estes et al., 1989; Myers et al., 1994; Shanks, 1990).

The simple adaptive network model of Gluck and Bower (1988a) has several limitations. Two of these also reported by Gluck and Bower (1988b, p. 180), are the adaptive network's inability to model subjects' sensitivity to correlated cues in classification, and classification based on nonlinear combination of cues. A third weakness is the sensitivity of the model to "noise" in the data (Gluck, 1992), which makes the model predictions generalise poorly. Even though the simple adaptive network model was able to account for "base rate neglect", its inability to reproduce the "inverse base rate" effect[2] found by Medin and Edelson (1988) has been one of its major empirical weaknesses (see Shanks, 1992, p. 10)[3].

---

[1] Also termed the "component cue model" (Gluck 1991), and the "standard connectionist model" (Shanks, 1992).

[2] The effect occurs when subjects are presented with a novel feature combination and the features have previously been seen in the context of another feature in which the original features had high diagnosticity. When presented in the novel combination, one may assume that the most common class is selected, but it was shown by Medin and Edelson (1988) that subjects select the rare category.

[3] It shares this weakness with even the most sophisticated exemplar theory models.

A minor modification of the adaptive network model was suggested by Shanks (1992) to make it account for inverse base rate effects. He suggested that the extent to which a stimulus is unexpected, may have special relevance both during learning and in the processing of a cue[1]. By using a modification of the delta rule (Widrow & Hoff, 1960) suggested by Wagner (1978), Shanks formulated his *attentional connectionist model.* In this model, learning is modulated by the "unexpectedness" of a stimulus (Shanks, 1992, p. 12). By incorporating this modification, the asymptotic weights of the model show selective attention to cues. Except for the modification of the learning rule, the model is similar to the adaptive network model. Shanks (1992) shows how the model consistently predicts the inverse base rate effect of a magnitude similar to the one found in his experimental data.

To overcome the other weaknesses of the simple adaptive network model, Gluck and Bower (1988b) suggested two refinements; the introduction of hidden layers in their model, and the configural cue model.

The configural cue model introduced in Gluck and Bower (1988b) was similar to the simple adaptive network model except that the input representation consisted of elementary features and all conjunctions of elementary features. The implausibility of this input representation led Gluck to formulate the *configural network model* (Gluck, 1991) in which only elementary features and *pairwise* conjunctions of features are represented at the input layer. Except from the input representation and the use of two output units, the configural network model is similar to the simple adaptive network model. Gluck (1991) tested the abilities of this model to replicate the findings of Medin and Schwanenflugel (1981) that subjects found a particular nonlinearly separable categorisation task easier to learn than a highly similar linearly separable categorisation task. As proposed, the configural cue model replicated the findings of Medin and Schwanenflugel (1981) rather well. Besides this finding, the configural network model has received little attention. The main reason is probably due to the representation of conjunctions, which, even though only pairwise conjunctions are considered, is a somewhat unrealistic assumption for the representation of complex objects (see Estes, 1994, p.75).

Another solution to overcome the original limitations of the adaptive network model was suggested by Gluck (1992), in a model replacing the original elementary feature representation with a distributed input representation. Except from this change in input representation, the *distributed stimulus sampling* model is similar to the adaptive network model (Gluck & Bower, 1988a). The distributed network model uses a representation of inputs derived from stimulus sampling theory of Estes (1950). It was shown to reproduce the "inverse base rate effect"[2] found by Medin and Edelson (1988) in addition to being far more

---

[1] This proposition is consistent with information theory assumptions (see Shanon, 1992).

[2] Gluck (1992) uses the term "relative novelty effect".

robust to "noise". Thus, it shows better generalisation properties than the original adaptive network model. A major problem, however, is how to select the distributed representation. Using stimulus sampling theory is just one of a large number of ways to convert the local representation into a distributed one. As shown in section 3.2.3, it is of great importance that the particular distributed representation is selected so that important relational properties of the stimulus are preserved[1].

It seemed that a distributed representation of inputs could overcome many weaknesses of the adaptive network model of Gluck and Bower (1988a). A major question, however, was how this distribution should be created. The largest degree of distribution of relevance for N different stimulus patterns with n stimulus dimensions would be to transform the n-dimensional local input vector into a N-dimensional distributed input vector of binary values. This transformation is finite for feature dimensions, but more problematic for stimulus dimensions. Having this problem in mind, Estes (1993, 1994) developed a *similarity network model* which introduces a layer of pattern units between the input and output layer. The number of pattern units equals the number of different patterns (N), and their activation is computed using similarity measures from exemplar theory. The activations of the pattern units are somewhat similar to the activations of the distributed units shown in figure 3.9[2]. New pattern units are introduced as long as new exemplars are presented at the input layer. Known patterns are used to modify pattern-to-output connections by error based learning, using a variant of the well known delta rule (Widrow & Hoff, 1960). Compared to previous exemplar models of Estes (1994), this model uses abstracted exemplar representations, since only one unit is found for each different stimulus pattern. As usual, the output units correspond to classes, and their activation is a weighted sum of the pattern unit activations. Output is produced using the traditional sigmoid output function, but learning is based on the pre-transformed activation values (Estes, 1994). The model has been tested on a variety of experimental data, and has generally showed good fit, and the capacity to overcome the weaknesses of the adaptive network model of Gluck and Bower (1988a).

However, two weaknesses of the similarity network model of Estes (1994) are serious. First, the estimation of the similarity parameters in the model is done separately, and is not subject to estimation during learning. This means selective attention is not driven by learning in the model. Second, the similarity network model is only formulated for feature dimensions, making transfer to continuous stimulus dimensions difficult in its present formulation.

---

[1] One example of such sensitivity is that the "receptive fields" of the distributed representations are placed in areas of the psychological space where stimulus objects are found. For example, radial basis function networks (Moody & Darken, 1989) use cluster analysis techniques to find these places.

[2] Even though the "isoactivation" contours vary with the similarity measure applied.

The problems with selective attention and continuous stimulus dimensions have been treated in further developments of the similarity network model. Both Hurwitz (1990) and Kruschke (1992) have suggested models that reduce its limitations. Because far more research results have been published on the model formulated by Kruschke (1992), and Kruschke acknowledges that similar developments have been made by Hurwitz (1990), we have chosen to present the ALCOVE model of Kruschke (1992) here.

Kruschke (1993) refers two lines of research as influential in his developments. From the neural network community, the research on radial basis functions (Moody & Darken, 1989; Poggio & Girosi, 1990) has suggested models with Gaussian hidden unit output functions placed at points in multidimensional input space. These suggestions are closely related to similarity functions placed at points in multidimensional space as suggested in some exemplar theory models. From the research on general context theory (Nosofsky, 1984, 1986), the idea of similarity as an exponential function of city-block distances is very similar to the ideas of radial basis functions researchers. However, radial basis function models traditionally use unsupervised learning to set the parameters of the hidden layer, and general context theory traditionally uses log likelihood estimation procedures. The idea of Kruschke was to implement general context theory in a radial basis function like model, and to use supervised learning for parameter estimation (Kruschke, 1993).

The result of this research, the ALCOVE model (Kruschke, 1992), is illustrated in figure 3.6. As briefly explained in section 3.1, the feed forward computations in ALCOVE are an implementation of general context theory. The inputs to ALCOVE are stimulus dimension values derived from multidimensional scaling. Thus, stimulus dimension values are values in a psychological space. The selective attention given to each of these dimensions is adjusted with a parameter not traditionally included in general context theory, the parameter α. This selective attention parameter is used to "stretch" and "shrink" the dimensions in psychological space to better discriminate exemplars in different classes, and to concentrate exemplars in the same class. The centre of the hidden units is placed at the stimulus dimension values of the exemplars[1]. The activation of the hidden units are given by:

$$o_{pj} = e^{-c\sum_{\forall i} \alpha_i |h_{ji} - I_{pi}|} \quad . \tag{3.21}$$

Here, o_{pj} is the activation of hidden unit j when pattern p is presented, c is a parameter determining the width of the receptive fields, h_{ji} is the centre of the receptive field of hidden

[1] In Nosofsky and Kruschke (1992), the formulation with hidden units placed at the exemplar positions is termed ALEX, and it is referred to the original formulation of ALCOVE as having its hidden units scattered randomly out in psychological space. However, as far as we can see, the formulation of ALCOVE given in Kruschke (1992), places its hidden units at the exemplars in multidimensional psychological input space.

unit j, and I_{pi} is the input on stimulus dimension i of pattern p. To illustrate how the receptive field is formed in ALCOVE, the isoactivation contours of four hidden units in two-dimensional psychological space are shown in figure 3.11[1].

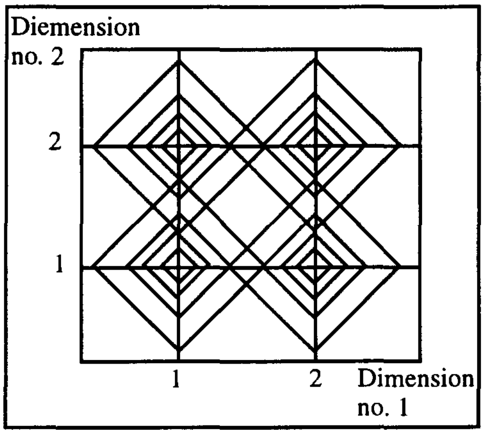

Figure 3.11 The isoactivation contours of ALCOVE

Since the city block distance metric is used, the receptive fields get diamond shaped isoactivation contours. In traditional radial basis function networks, the hidden units are Gaussian. Thus, the receptive fields of radial basis function networks are similar to those shown in figure 3.9[2]. The consequence of the diamond shaped receptive fields is that the hidden units of ALCOVE are most sensitive to changes in stimulus dimension values along the diagonal of the receptive fields.

Classification is performed by the associative weights connecting hidden units and output units corresponding to categories or classes. The activation of the hidden units is computed as the weighted sum of associative weights and hidden unit activation values. The traditional choice rule of Luce (1963) is used for classification, and is applied to the activation of the output units. This choice rule has previously been shown to be equivalent to the use of a sigmoid output function at the output units (Nosofsky, 1992).

Parameter estimation in ALCOVE is done by supervised, error based, learning. The principles of the backpropagation learning rule are applied (Rumelhart, Hinton & Williams, 1986). For the output layer this is done by applying the delta rule (Widrow & Hoff, 1960) in the same manner as in equation (3.14). However, in ALCOVE, as in most other connectionist classification models, the error is computed before the output is transformed in the Luce (1963) choice rule.

Supervised learning in the hidden layer is more difficult to implement in ALCOVE. Traditional radial basis function networks (Moody & Darken, 1989; Poggio & Girosi, 1990) have two unknown parameters in their hidden layer. The position of the receptive fields h_{ji} is traditionally determined using some kind of unsupervised learning procedure, such as Kohonen learning (Kohonen, 1977), which actually performs k-means clustering[3] and places

[1] Actually, the stimulus dimensions shown here are post attenuated stimulus dimensions adjusted for the α parameter.
[2] The choice of city block and exponential similarity gradients in ALCOVE is founded in research by Shepard (1987), and makes the receptive fields of ALCOVE different. See e.g. Kruschke (1992, p. 23).
[3] See Hertz, et al., 1991.

one hidden unit in the centre of each cluster[1]. The other free parameter of the radial basis function networks is the width of the receptive fields. This parameter is traditionally determined using some k-nearest neighbour heuristic (Wasserman, 1993, p. 154). In ALCOVE, the width of the receptive field is constant across hidden units and is determined by the parameter c. From the published material on ALCOVE (Kruschke, 1992, 1993a, 1993b; Nosofsky and Kruschke, 1992), it is not clear to us how c is determined. The receptive fields are placed at the exemplar positions in psychological input space. These parameters are not subject to learning, but are fixed in ALCOVE. Instead of letting these parameters be learned, ALCOVE introduces supervised learning of the selective attention parameter. The learning of this parameter is performed using backpropagation of error derivatives, similar to Rumelhart, Hinton and Williams' (1986) formulation.

ALCOVE has been shown to be able to model several aspects of human classification. Kruschke (1992) showed how the model could learn to attend to relevant dimensions, learn to attend to correlated dimensions, reproduce the base rate neglect of the models by Gluck and Bower (1988a), and learn nonlinearly separable classification tasks. As such, it represents a major improvement on the simpler connectionist models reported above. The inverse base rate effect of Medin and Edelson (1988) can not be modelled by ALCOVE in its present form, but Nosofsky and Kruschke (1992) have shown that incorporating exemplar specific selective attention, the model will replicate this effect also. ALCOVE must be considered one of the most impressive connectionist models of classification and categorisation presently available (Estes, 1994; Robins, 1992). However, some limitations have been found, and as will be shown, some of these are of particular importance to the modelling of financial diagnosis.

One limitation of ALCOVE mentioned by Nosofsky and Kruschke (1992), is that using gradient descent learning, it will be unable to reproduce the abrupt shifts in attention found by human subjects in some tasks (Nosofsky & Kruschke, 1992, p. 244). This type of limitation is, in our opinion, of less relevance to a model of classification at the molar level, even though it certainly is of relevance to a model of classification *learning*. Of much greater relevance to classification tasks depending upon abstract stimulus dimensions or change in stimulus dimensions during learning, are the limitations discussed in Kruschke (1992, pp. 40-41). These limitations are simultaneously responsible for many of the favourable properties of ALCOVE, and are stated by Kruschke (1992) as:

[1] Notice that this implies that the number of hidden units is smaller than the number of exemplars. Thus, abstraction in the form of location of units at prototype centres takes place, actually implementing a multiple prototypes theory.

A third property of learning in ALCOVE is that attention learning can only adjust to the relative importance of the dimensions as given. ALCOVE cannot construct new dimensions to attend to. (Kruschke, 1992, p. 25)

Since the stimulus dimensions in ALCOVE are derived by multidimensional scaling, their psychological relevance seems well founded. However, several tasks require the combination of stimulus dimensions into *abstracted* stimulus dimensions for classification to be effective. As an example, classification based upon financial data may require the use of abstracted stimulus dimensions, such as "trend" or "stability" (Falbo, 1991). Such stimulus dimensions may be learned during classification, but can not be formed by ALCOVE. A methodological limitation of ALCOVE is that it requires both similarity judgements and classification data as input to the model, making the data collection more comprehensive.

A model formulated particularly to *build* abstracted stimulus dimensions during learning is the multilayer perceptron model of Rumelhart, Hinton and Williams (1986), termed the *backpropagation* network. This model was criticised by Kruschke (1993b) for not being able to reproduce subjects' behaviour on a filtration task (see Kruschke, 1993b), to suffer from catastrophic forgetting (McCloskey & Cohen, 1989) and for learning nonlinearly separable categories too slowly. Taraban and Palacios (1993) have shown that the criticisms of Kruschke (1993b) against backpropagation networks were wrong on the first and third argument. Progress has also been made on preventing catastrophic forgetting (see Clark, 1993, p. 145-147) in such models.

The backpropagation model is illustrated in figure 3.7, and the equations controlling response production and learning in the model are given in equations (3.5) to (3.19).

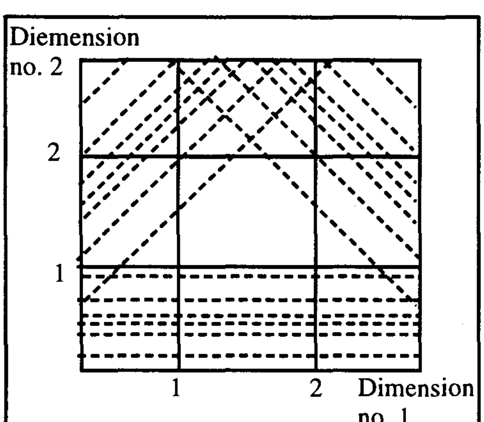

Figure 3.12 Isoactivation contours of backpropagation

The backpropagation model can use distributed or local input representations, and is not restricted to distributed encoding. Hidden units can be introduced in one or several hidden layers. The "receptive fields" of the hidden units in backpropagation models are illustrated in figure 3.12.

Traditionally, outputs of hidden units are sigmoid functions of the weighted sum of inputs, but other output functions can also be used. In figure 3.12, sigmoid output functions are assumed, and activation and output are equated. As is illustrated, the direction of the receptive fields is arbitrary and is determined during learning. Model output is similarly a sigmoid function of the weighted sum of hidden unit activations, and consequently, a *function* of the "receptive fields" shown in figure 3.12. Thus, the functional form of the model is as given in equation (3.20), proven to have universal approximation properties (Hornik et al., 1989).

Learning is performed by gradient descent on error as shown in equations (3.7) to (3.19), and consists of weight adjustments determined by the backward propagated output error derivatives. The model has been applied to model several cognitive phenomena with considerable success (e.g. Rumelhart & McClelland, 1986; Seidenberg & McClelland, 1989; Sejnowski & Rosenberg, 1986; Taraban, McDonald & MacWhinney, 1989). Recently, Taraban and Palacios (1993) successfully applied the model to several classification phenomena, and showed how the model could produce many of the empirical findings on human classification. They also showed how the model could be modified to overcome previous criticism (Kruschke, 1992, 1993a, 1993b). Taraban and Palacios (1993), stressed the importance of *feature abstraction* in the backpropagation model as a property not shared by ALCOVE, and this ability is of great significance to many classification tasks of higher complexity (Chandrasekaran & Goel, 1988). Since the ability to build internal representations and to form abstracted stimulus dimensions of relevance to classification is necessary to a classification model of financial diagnosis, and this capacity currently is restricted to the backpropagation model, we concentrate on this model in chapter 4.

Chapter 4. Connectionist models of financial diagnosis

We concluded from chapter 2 that both the judgement modelling and the predictive approaches had treated financial diagnosis as a classification task, and that the cognitive approach basically had overlooked classification theory as a relevant cognitive theory of financial diagnosis. In chapter 3, we introduced cognitive classification theory, and showed how recent approaches shared an interest in applying connectionist ideas of cognition to classification problems. We ended chapter 3 with the suggestion that certain connectionist models of classification could provide the means to successfully model complex classification tasks like financial diagnosis.

As shown in chapter 3, a close relationship exists between connectionist theories of cognitive classification and artificial neural network models developed for classification from a predictive perspective. We showed how some artificial neural networks had universal approximation properties, and thus, that they could approximate a posteriori classifiers. This was suggested as one of the reasons for the success of artificial neural networks as classification and prediction devices. This success has not gone unattended in the community interested in financial diagnosis from a predictive perspective. Several applications of artificial neural network models to financial diagnosis are found. Whether any of these models have cognitive relevance, is, however, an open question.

Before introducing our model of financial diagnosis built by applying a connectionist model of cognitive classification, we review some of the artificial neural network applications to the financial diagnosis task. We wish to establish if any of these applications have cognitive relevance, and if so, build this relevance into our own model.

The review of artificial neural network applications in financial diagnosis is found in section 4.1. In section 4.2, an attempt is made to unify the theoretical perspectives and empirical findings referred to in chapters 2 and 3, into a connectionist classification model of financial diagnosis. Furthermore, some propositions that can be derived from the model are presented and elaborated.

4.1 Neural networks in financial diagnosis

A survey of connectionist and artificial neural network applications to business administration tasks was performed. Leading connectionist, artificial neural network and cognitive accounting journals and conference proceedings were searched. This resulted in a relatively large number of studies, illustrated by the summary provided in appendix A. The studies cover applications to problems in economics and finance, such as currency exchange rate prediction (e.g. Mehta, 1995; Refenes, 1993), stock price prediction (e.g. Schöneburg, 1990; White, 1988), futures price prediction (e.g. Grudnitski & Osburn, 1993), derivative securities prediction (e.g. Hutchinson, Lo & Poggio, 1994) and pricing of material goods (e.g. Chakaborty, Mehrotra, Mohan & Ranka, 1992)[1]. In addition, applications in management (e.g. Jung & Burns, 1993) and in marketing (e.g. Brown, 1992; Wray, Palmer & Bejou, 1994) are found. In accounting, several of the studies are applications of relevance to financial diagnosis, such as e.g. bankruptcy prediction and loan evaluation studies. A summary of selected applications is shown in table 4.1.

Reference	Task[2]	Model	Benchmark	RA[3]	Cog[4]
Altman, Marco & Varetto, 1994	Bankruptcy prediction[5]	Backpropagation	Discriminant analysis	N	N
Erxleben, Baetge, Feidicker, Koch, Krause & Mertens, 1992	Bankruptcy prediction	Backpropagation	Discriminant analysis	N	N
Martin-del-Brio & Serrano-Cinca, 1993	Bankruptcy prediction[6]	Selforganizing map	None	Y	N
Odom & Sharda, 1990	Bankruptcy prediction	Backpropagation	Discriminant analysis	N	N
Poddig, 1995	Bankruptcy prediction	Backpropagation and LVQ	Discriminant analysis	N	N
Raghupathi, Schkade & Raju, 1991	Bankruptcy prediction	Backpropagation		N	n[7]
Rahimian, Singh, Thammachote & Virmani, 1993	Bankruptcy prediction	Backpropagation, Athena and simple perceptron	None	N	N
Salchenberger, Cinar & Lash, 1992	Bankruptcy prediction[8]	Backpropagation	Logit model	N	N
Tam & Kiang, 1992	Bankruptcy prediction[9]	Backpropagation	Discriminant analysis, logistic regression, KNN and ID3	N	N
Tam, 1991	Bankruptcy prediction	Backpropagation	See Tam and Kiang, 1992	N	N
Udo, 1993	Bankruptcy prediction	Backpropagation	Regression analysis	N	N
Wilson & Sharda 1994	Bankruptcy prediction	Backpropagation	Discriminant analysis	N	N
Dutta & Shekhar, 1988	Bond rating	Backpropagation	Regression analysis	N	N

cont...

[1] Two collections of applications have also recently been edited by Trippi and Turban (1993) and Refenes (1995).

[2] The task context terms of section 2.1 is used. The term "stock prediction" is used for tasks including stock price or return predictions, and stock classifications.

[3] Representational analysis (RA) is marked Y if the study contains an analysis of how the representations of the neural network performs the vector mappings, and N elsewhere. Lowercase letters are used to indicate doubt about the classification.

[4] Cognitive (Cog) is marked Y if the study refers to the neural network as a cognitive model or compares it to a cognitive model of the task performance, and N elsewhere. Lowercase letters are used to indicate doubt about the classification.

[5] The authors use the term "corporate distress diagnosis".

[6] Martin-del-Brio and Serrano-Cinca (1993) study *bank* classification and bankruptcy prediction.

[7] However, Raghupati et al. (1991) state: "Various financial ratios may be giving some intermediate features such as immediate financial health of the company, long-term financial health, recent revenue generating trends, and others. Based on these higher-level features, the network may be arriving at a categorizing decision".

[8] Salchenberger et al. (1992) study failure of *thrift institutions*.

[9] Tam (1991), and Tam and Kiang (1992) study *bank* bankruptcy prediction.

cont...

Kim, Weistroffer & Redmond, 1993	Bond rating	Backpropagation	Discriminant, regression and logistic analysis, ID3	N	N
Moody & Utans, 1995	Bond rating	Backpropagation	Regression analysis	n[1]	N
Singleton & Surkan, 1995	Bond rating	Backpropagation	Discriminant analysis	N	y[2]
Surkan & Ying, 1991	Bond rating	Backpropagation	None	y[3]	N
Surkan & Singleton, 1990	Bond rating	Backpropagation	Discriminant analysis	N	y[4]
Utans & Moody, 1991	Bond rating	Backpropagation	Regression analysis	N	N
Coats & Fant, 1993	Going-concern judgement	Cascade correlation	Discriminant analysis	N	N
Barker, 1990	Loan decision	N/A	None	N	N
Deng, 1993	Loan decision	Backpropagation	None	N	N
Nottola, Condamin & Naim, 1992	Loan decision	Backpropagation	ID3	y[5]	N
Piramuthu, Shaw & Gentry, 1994	Loan decision	Backpropagation and 2. order modification	Probit analysis and ID3	N	N
Romaniuk & Hall, 1992	Loan decision	Feed forward network with cell recruitment learning	None	n[6]	N
Srivastava, 1992	Loan decision	Backpropagation	None	N	n[7]
Kryzanowski, Galler & Wright, 1993	Stock prediction	Boltzmann Machine	None	N	N
Refenes, Zapranis & Francis, 1995	Stock prediction	Backpropagation with variations	Regression analysis	y[8]	N
Wong, Wang, Goh & Quek, 1992	Stock prediction	Backpropagation	None[9]	N	N
Yoon, Swales & Margavio, 1993	Stock prediction	Backpropagation	Discriminant analysis	y[10]	N
Yoon, Guimaraes & Swales, 1994	Stock prediction	Backpropagation	Discriminant analysis[11]	N	N
Berry & Trigueiros, 1993	Ratio analysis[12]	Backpropagation	Discriminant analysis	Y	n[13]

Table 4.1 Applications of connectionist and artificial neural network models to financial diagnosis.

[1] Moody and Utans (1995) use sensitivity analysis to determine the importance of input units, and thus, study the input-output mappings of corresponding variables.

[2] Singleton and Surkan (1995) state that "Neural network success........ suggests that neural networks may have captured some of the judgement exercised by these analysts".

[3] Surkan and Ying (1991) test the sensitivity of the network to exclusion of inputs. This is used to simplify the model so that it is represented in only one sigmoid function.

[4] Surkan and Singleton (1990) state: "There is a hope that some of the intermediate representations may be identified with concepts used by humans to analyze this bond classification problem".

[5] Nottola et. al (1992) use ID3 to extract rules from the input to hidden unit output mapping.

[6] Romaniuk and Hall (1992) give examples of rules extracted from the neural network by "traversing" the network. The exact method of this "traversing" is not explained.

[7] However, Srivastava (1993) states about the model: "It simulates human judgement and integrates it with mathematical analytical tools".

[8] Refenes et al. (1995) perform sensitivity analysis to investigate the input-output mappings.

[9] Wong et al. (1992) focus on integrating the artificial neural network with an expert system.

[10] Yoon et al. (1993) investigate the effect of the different inputs on the classification, not the representation as such.

[11] Focus is on integrating the artificial neural network with a rule-based expert system.

[12] Berry and Trigueiros (1993) use ratio analysis to predict industry classifications.

[13] In their original paper presented at INNC, 1990, they state: "The emerging organization reproduces the way an expert in ratio analysis chooses variables...Experts put together several points of view around a few significant variables. And extended ratios seem to be trying the same sort of procedure" (Trigueiros & Berry, 1990, p. 12).

In table 4.1, the studies are ordered by task context as in the review tables of chapter 2. Furthermore, the applied connectionist or artificial neural network model is indicated. Most studies compare the connectionist or artificial neural network model with more traditional models. These models are termed "benchmarks", and are shown for each study. The last two columns of table 4.1 indicate if the study reports an analysis of the representations developed by the model after learning, and if the author claims that the model has cognitive relevance[1]. Our review of these studies is organised by task context, starting with the applications to the bankruptcy prediction context.

One of the first applications of artificial neural networks to *bankruptcy prediction* was Odom and Sharda's (1990) study, comparing a backpropagation model and discriminant analysis on a sample of 129 firms. They used the cues of Altman (1968), and found that the artificial neural network outperformed discriminant analysis on training and test samples. However, no test of significance or cross validation results were reported. In a further analysis of the results, Wilson and Sharda (1994) reported *significantly* better performance for the backpropagation models. A similar study, using the same cues and data, was reported by Rahimian et al. (1993), and showed results very similar to Odom and Sharda's (1990) study.

Another small sample study[2], concentrating on cue selection and network sizing was reported by Poddig (1995). His performance results were comparable to Odom and Sharda's (1990), showing superior performance for the artificial neural network. In the study by Udo (1993), a linear regression model was used as benchmark, making comparisons with previous studies difficult.

Two large sample studies[3] have been performed on the bankruptcy prediction task. Altman et al. (1994) tested several artificial neural network models for bankruptcy prediction against a two-stage discriminant analysis based system developed for Italian business conditions. A simple backpropagation model, a model with time-series organisation of inputs, and a model with conceptual organisation[4] were developed. In general, the network performance on different test samples was comparable to the results obtained with discriminant analysis, but no systematic cross validation results were reported. Only limited analysis of network representations was performed, and no claims of cognitive relevance were made, even though the authors compared model behaviour under special conditions to the judgements made by

---

[1] In our terms; if the model is connectionist.

[2] The total sample consisted of 150 firms.

[3] The Altman et al. (1994) study used a sample size of 1108 firms, while the Erxleben et al. (1992) study used 3539 data sets on average covering 3 years of cues, so that the sample size is approximately 1180 firms.

[4] The model was set up to perform sub classifications of the firms in eight conceptual areas, and to combine these sub classifications in a final diagnosis.

analysts under similar conditions. The "illogical types of behavior" (Altman et al., 1994, p. 526) of the artificial neural network under these conditions were considered unacceptable[1]. Variables used as inputs were not reported. A similar large sample study performed under German business conditions, was performed by Erxleben et al. (1992) showing very similar results when the model was compared to discriminant analysis.

The studies reported so far on the bankruptcy prediction task used traditional manufacturing and retailing firms as classification objects. Three studies concentrating on *bank failure* predictions were the studies of Tam (1991), Tam and Kiang (1992)[2] and Martin-del-Brio and Serrano-Cinca (1993). Tam (1991) used 19 cues collected from two consecutive years of 118 banks in a comparison of single and multilayer perceptrons to linear discriminant analysis, logistic regression analysis, nearest neighbour algorithms, and a recursive partitioning algorithm (ID3). The study used cross validation results, adjusted for misclassification costs and base rates, and concluded that the backpropagation model outperformed the other models, but significance was not reported. Similar institutions were studied by Salchenberger et al. (1993). They used backpropagation models to predict the failure of savings and loan institutions. When compared to logistic regression, the artificial neural networks predicted significantly better. Initial experiments were used to select cues representing the theoretical concepts "capital", "assets", "management", "earnings" and "liquidity".

Common to all these bankruptcy prediction studies are that they do not contain representational analysis, and for obvious reasons, no cognitive claims are made[3]. A study concentrating on representational analysis was performed by Martin-del-Brio and Serrano-Cinca (1993), but the artificial neural network used was a Selforganizing map (see Kohonen, 1995). This network was used to cluster 66 Spanish banks based upon 9 well known financial ratios. An advantage of these networks is their illustrative power, but no traditional tests of performance was reported[4]. Representational analysis was performed with weight maps[5] to illustrate how regions of solvent and insolvent banks were formed by unsupervised learning. The analysis further revealed that solvent banks were placed in different regions depending on their value on operationalised theoretical concepts such as "profitability" and "liquidity". However, no cognitive claims were made for the models, or for the cognitive relevance of theoretical concepts presumed operationalised in these models.

---

[1] With reference to the lens model of Brunswik (1952) shown in section 2.1, this conclusion is made of a model of the left hand side of the lens model by comparing it to an "intuitive" understanding of the rights hand side, and is thus, unwarranted.

[2] The Tam (1991) and Tam and Kiang (1992) articles report the same study.

[3] These models are purely on the left hand side of the lens model (see section 2.1).

[4] The network is autoassociative.

[5] Weight maps are similar to Hinton diagrams when the units are fully interconnected and not organised in layers.

Raghupathi et al. (1991) used a backpropagation model in a context very similar to Odom and Sharda's (1990), but concentrated more on the course of learning. A suggestion that the artificial neural network model was able to "extract higher-level features" (Raghupathi et al., 1991, p. 156) of relevance to the classification, is interesting. Even though the authors suggested these higher level features were complex, no direct claims of their cognitive relevance were made.

A closely related study of auditors *going concern judgements*, is the cascade correlation application of Coats and Fant (1993). The cascade correlation algorithm (Fahlman & Lebiere, 1990) is a constructive algorithm building multilayer perceptrons by sequentially adding hidden units to increase performance. With the going concern judgement of auditors, a *behavioural* response variable was used, but in the Coats and Fant (1993) study, this variable was suggested as an indicator of bankruptcy. Cues of the Altman (1968) model were used as inputs to predict the going concern judgement of 282 firms. The resulting network performed significantly better than discriminant analysis in predicting qualifications for going concern problems, but was outperformed by the discriminant analysis on the "healthy" firms. Unfortunately, no analysis of representations was performed, but such analysis may have had cognitive relevance due to the behavioural response variable used.

One of the first applications of artificial neural networks in accounting and finance was the *bond rating* study of Dutta and Shekar (1988). They tested the ability of several backpropagation models to predict the bond ratings of 47 companies, using 6 or 10 cues of the financial statement and other sources. In their first report (Dutta & Shekar, 1988) they compared a backpropagation model with linear regression models, and used a transformed response variable consisting of two classes. In this case, the backpropagation model outperformed the regression model, but only simple validation was performed, and no significance tests were reported. In an extension reported in Dutta, Shekar and Wong (1994), different benchmarks[1] were used, and different transformations of the rating variable were reported. The same relationship was found between the artificial neural networks' and the benchmarks' performance as in the original report. In addition, Dutta et al. (1994) reported the prediction errors of the artificial neural networks to be within a single rating class distance, while the benchmarks often missed by more than one class.

Very similar, small sample applications, have been reported in a series of studies by Singleton, Surkan and Ying (e.g. Singleton & Surkan, 1995; Surkan & Singleton, 1990; Surkan & Ying, 1991). In Surkan and Singleton (1990), they tested backpropagation models against discriminant analysis in a two valued transformation of bond ratings of a collection of

---

[1] Variations of logistic regression and logit analysis models were used.

18 firms[1]. As above, the backpropagation models outperformed the discriminant analyses. In Singleton and Surkan (1995), a similar result was obtained for prediction of bond rating *changes* of the same firms. In Surkan and Ying (1991), the backpropagation model was analysed further, and attempts were made to simplify the mapping performed by the network in order to implement it in one simple sigmoid function. To arrive at this form, some analysis of network representations was performed. However, more interesting were the propositions made in Surkan and Singleton (1990), that the intermediate representations created in the backpropagation models may "be identified with concepts used by humans to analyse this bond classification problem" (Surkan & Singleton, 1990, p. 286). However, no further attempt to pursue this claim was made.

An application more closely resembling traditional research on bond ratings is the application of Kim et al. (1993), using a backpropagation model and input variables derived from Belkaoui (1980) to predict the six class debt ratings of 168 firms. Benchmarks in the study were linear regression, discriminant analysis, logistic regression and a recursive partitioning algorithm (ID3). Unfortunately, no cross validation was performed, but results indicated significantly better performance of the backpropagation model than of any other benchmark, except logistic regression. Even though not all performance differences were significant, the differences were always in favour of the artificial neural network.

Of considerably higher quality than the studies reviewed so far, are the studies by Moody and Utans (Moody & Utans, 1995; Utans & Moody, 1991), applying backpropagation models to predict the bond ratings of 196 industrial firms. They started with 10 cues of the financial statement assumed to represent the theoretical concepts "leverage", "coverage" and "profitability". The number of input variables and the size of the network were determined by constructive algorithms and posterior pruning. The reported results were cross validated, and showed significantly better performance of the artificial neural network than of the regression model benchmark, even when the response variable remained scaled to 16 classes. This work concentrated on constructive and pruning algorithms, and on the use of cross validation. Consequently, it contributed mainly to the artificial neural network community (see Moody, 1993), but the application to bond rating showed that the use of constructive algorithms and cross validation was practically applicable[2]. Parts of the sensitivity analysis used for network pruning involved analysis of representations, but interpretation of the representations was not considered in these studies. Furthermore, no cognitive claims were made.

---

[1] Data were collected for 7 years, resulted in 126 "patterns".
[2] Despite their computational requirements.

One of the first applications of artificial neural networks to the *loan decision* was a study by Barker (1990), applying a backpropagation model as part of an expert system. However, training cases, network structure, and performance results were not reported. A similar study applying artificial neural networks to refine the induced rules of the knowledge base in a credit granting system, was performed by Deng (1993). Even though the author stressed the artificial neural network's ability to extract knowledge from past granting decisions, the material used for learning was severely limited[1], no performance results and no benchmark comparisons were reported. Srivastava (1992) reported using the data of Abdel-Khalik and El-Sheshai (1980), but on closer inspection it seems that they used the recursive partitioning model of Messier and Hansen (1988) developed *from* the Abdel-Khalik and El-Sheshai (1980) data, to construct a "health" index as part of a model predicting loan denial. No traditional report was made of network topology and performance results, and input data collected from other sources than the financial statement were used in the final model. Piramuthu et al. (1994) also reported a test using the Abdel-Khalik and El-Sheshai (1980) data, but of more interest was a loan evaluation study reported in the same paper. In this study, the authors tested the ability of a backpropagation model to predict the loan classification of 100 firms using financial statement cues and loan information as input variables. A comparison was made with probit analysis and a recursive partitioning algorithm (ID3), showing superior performance of the backpropagation model used. Modifications in the original backpropagation learning rule (Rumelhart, Hinton & Williams, 1986) were suggested to improve learning speed[2], but the performance results remained the same. Furthermore, no reports of significance and cross validation results were made. A multilayer perceptron with a "cell recruitment" constructive algorithm was used by Romaniuk and Hall (1992) in a study of the creditworthiness decisions on 50 firms. However, the model was not validated, and the study concentrated on how rules could be generated from the model to make the model "explain itself" (Romaniuk & Hall, 1992, p. 20).

In a study by Nottola et al. (1991), a larger training set of 600 *company evaluations* of "practical experts" (Nottola et al., 1991, p. 510) was used to develop a backpropagation model. However, no traditional report was given on performance. Instead, Nottola et al. (1991) focused on how recursive partitioning (ID3) could be used to induce a set of rules explaining the behaviour of each hidden unit. This suggestion is new and interesting, but as shown in chapter 3, representations of hidden units in a backpropagation model are typically distributed, and *separate* analysis of the behaviour of each hidden unit may be of little value.

Common to all of the applications to the loan decision context was a focus on the models as loan evaluation and granting *systems*. Comprehensive evaluation of the systems against

[1] Deng (1993) used a sample size of 14.
[2] Piramuthu et al. (1994) applied second order methods to improve learning speed.

traditional benchmarks was not performed. Despite this systems orientation, surprisingly little analysis of knowledge representation was done, and no claims of cognitive relevance were made.

Stock prediction studies are treated as a class containing several contributions which have some relevance to financial diagnosis. They apply a fundamental analysis perspective on the pricing of stocks, and try to model the relationship between fundamentals and prices by using artificial neural networks. However, the fundamentals used, and the measure of stock price, stock evaluation or stock returns used as dependent variables, differ considerably across the studies. Kryzanowski et al. (1993) used a Boltzmann machine (Ackley et al., 1985) to classify stocks into three classes expected to show increased, decreased or stable stock returns. Inputs were *binary* trends of 14 financial ratios representing the theoretical concepts "profitability", "debt" and "liquidity and activity". In addition, seven similarly coded macroeconomic variables were used as inputs. Performance was not tested against traditional benchmarks, but was shown to be better than chance. In the study of Yoon et al. (1993), backpropagation models were used to classify stocks into two classes of "good" and "bad" stocks, as evaluated by official sources. Inputs were four financial cues[1], and the model was used to classify the stocks of 151 firms. The model was compared to discriminant analysis, and showed superior performance even though significance was not reported. In a later report, Yoon et al. (1994) reported integrating the model with a rule-based system, but the "integration" consisted of supplying the artificial neural network with a rule-based explanation facility. A similar "integration" was reported by Wang et al. (1992), but in this study, the inputs to the artificial neural network were linguistic values derived from a rule-based fuzzy information processor. It was also shown how this information processor could be implemented in an artificial neural network. Again, no traditional evaluation of prediction results was performed.

In a more formal study, Refenes et al. (1995) started with the principles of arbitrage pricing theory (Ross, 1976), and investigated the relationship between three assumed factors and excess return of 143 stocks. Despite the high quality of the study, the derivation of the three factors was not explained, and thus, the relevance of these factors to financial diagnosis is unclear. The arbitrage pricing model traditionally assumes all factors are determined by external sources. Other studies applying artificial neural networks to aspects of stock valuation and pricing by using macroeconomic and time series data, have been performed with varying success (e.g. Schöneburg, 1990; White, 1988). Continued research is pursued from this perspective, but its relevance to financial diagnosis is limited. ,

---

[1] The cues were current ratio, return on equity, price/earnings ratio and price/sales ratio.

A final study, reported by Trigueiros and Berry (1990, Berry & Trigueiros, 1993), applied backpropagation models to classification of firms into classes assumed to be equally affected by environmental factors. This dependent variable was not a financial diagnosis variable, but the study is of some interest because of its representational analysis and cognitive claims. The authors discussed many of the limitations with traditional methods in prediction and classification due to the distribution properties of financial cues[1]. The backpropagation models were tested against discriminant analysis, and were found to perform significantly better. Trigueiros and Berry (1990) compared the behaviour of the backpropagation models to the presumed behaviour of experts when input data were varied, and reported several similarities. Even though the cognitive relevance of these comparisons is limited, the authors point in a direction of research of considerable interest.

Several of the other studies shown in appendix A are related to financial diagnosis, but are applications to tasks not covered by our definition of financial diagnosis. Two examples of artificial neural network applications using financial statement data for analysis purposes are the studies by Liang, Chandler, Han and Roan (1992), predicting LIFO/FIFO inventory evaluation, and by Sen, Oliver and Sen (1995), predicting corporate mergers[2]. These studies were excluded from this review because the classification studied was not unambiguously related to financial health. Some studies were excluded despite their use of response variables corresponding to aspects of financial health, because they primarily used input variables external to the financial statement (e.g. Yoon & Swales, 1991). Finally, some studies were of financial diagnosis tasks, but investigated other classification objects than firms, such as individual loan applicants (Jensen, 1992) or individual homeowners (Collins, Gosh & Scofield, 1988; Yamamoto & Zenios, 1993).

From this review of connectionist and artificial neural network applications to financial diagnosis, we must conclude that no *connectionist* applications are found. All the studies used neural network principles from a predictive perspective, and without explicit reference to a cognitive model of the task. In some tasks contexts, such as bankruptcy prediction this is unproblematic because the dependent variable used, is not behavioural. In other task contexts, such as bond rating, in principle, a behavioural dependent variable is used, and cognitive relevance is of importance. Furthermore, few of the studies performed extensive analysis of network representations in order to present and understand why the artificial neural networks' performance was superior. As an example, no search for the "higher level features" proposed by Raghupathi et al. (1991), or "concepts used by humans" proposed by Surkan and Singleton

---

[1] Examples are the multivariate normality assumptions of discriminant analysis and the assumption of uncorrelated independent variables in regression analysis.
[2] These studies have dependent variables of similar interest as the variable in the Trigueiros and Berry (1990) study, but because of their purely predictive perspective, the studies are not reviewed.

(1990), was performed. However, the propositions of these intermediate abstractions are interesting and deserve further investigation.

4.2 A connectionist model of financial diagnosis

By using the findings reported on the financial diagnosis task, the theory proposed on cognitive classification, and the models provided by connectionist theory, we can formulate a connectionist model of financial diagnosis. The introduction of the model is divided into three parts. In the *presumptions* part, the presumptions and premises of a connectionist model of financial diagnosis are summarised using relevant research findings reviewed in chapters 2 and 3. In the *presentation* part, the model is introduced and illustrated graphically. In the *propositions* part, we elaborate on some propositions presumably supported by an empirical investigation and simulation of the model.

The *presumptions* of the connectionist model of financial diagnosis can be organised in three groups. The first group of premises is based upon the assumption that financial diagnosis is a classification task. This assumption does not seem unwarranted when considering the studies reviewed in chapter 2. From a cognitive point of view, however, it means treating the task at a "lower level" of cognition (see Osherson & Smith, 1990; van Gelder, 1993) than is traditional in cognitive processing studies of financial diagnosis. However, diagnostic tasks in other areas[1] are studied from this perspective (Brooks et al., 1991; Mumma, 1993). The term classificatory diagnosis (Chandrasekaran & Goel, 1988) has been used of complex classification tasks in which classification is presumed to depend upon a classification hierarchy and the use of intermediate abstractions by the cognitive system. Thus, a model of financial diagnosis as classificatory diagnosis should be able to implement classification hierarchies and form intermediate abstractions.

The second group of premises stems from the application of relevant theory of cognitive classification. First, the model of financial diagnosis should be based upon established theory of cognitive classification. Second, models derived from the theory must have been shown to explain most of the empirical findings on human classification, such as base rate effects (Medin and Edelson, 1988) and classification learning effects (e.g. Shepard, Hovland & Jenkins, 1961). Third, model operationalisations of the theory should have well known formal properties[2].

---

[1] Such as medical and psychiatric diagnoses.

[2] Such as a well understood mapping function, formal properties as a posterior classifier, or known limitations with relevance to cognitive classification.

The third group of premises stems from the constraints that the task put on the model. These constraints are provided by research on financial diagnosis and relate to stimulus dimensions, relevant responses, and the processing and representational assumptions of the model. In financial diagnosis, stimulus dimensions are real valued, correlated and sometimes configural. Thus, the model must show sensitivity to these properties of stimuli. Furthermore, selection of relevant stimuli must have a basis in theory of financial diagnosis and empirical findings on cue usage. Relevant responses of financial diagnosis are clinical narratives or linguistic terms transformed to express class memberships, or direct assessments of class memberships. If financial diagnosis is similar across task contexts, these linguistic terms or membership assessments should be derived in a generic diagnosis situation. The processing and representation principles of the model must incorporate explanations of cognitive phenomena previously unexplained in information processing terms. Important aspects of processing are the role of pattern recognition, prototype similarity and analogical reasoning. Important aspects of representation are the role of search-independent representations, schema representations, prototype representations and template matching.

From a set of basic principles of processing and representation, a model must be developed that is interpretable in cognitive terms. Whether the developed representations reflect competence theory of financial diagnosis is an empirical question to be investigated, but such investigation must be possible. Furthermore, a methodology must be used that does not assume direct correspondence between measured units and units of the representational system. This should prevent the developed representations from being constrained by methodology so that only certain types of representations can be formed. Thus, the model must be able to develop, for example, symbolic representations and perform "qualification"[1], but the model must not be constrained in a way that *only* allows such representations to develop. Furthermore, the model should allow processing behaviour describable in rule-based terms to develop, but not in a way that makes similarity based processing impossible.

With these presumptions introduced, the selection of a *connectionist* mode of classification should "come natural". These models are considered "among the leading candidates" (Nosofsky, Gluck, Palmieri, McKinley & Glauthier, 1994, p. 366), in cognitive classification models today. The backpropagation model was specially formulated to develop internal representations, a property of a model presumed necessary in the discussion above. If implemented in a backpropagation model[2], the connectionist model may be illustrated as in figure 4.1.

---

[1] That is, transformation from analog to digital form of stimulus dimensions.
[2] An introduction to the structure and algorithms of this model is given in section 3.2.

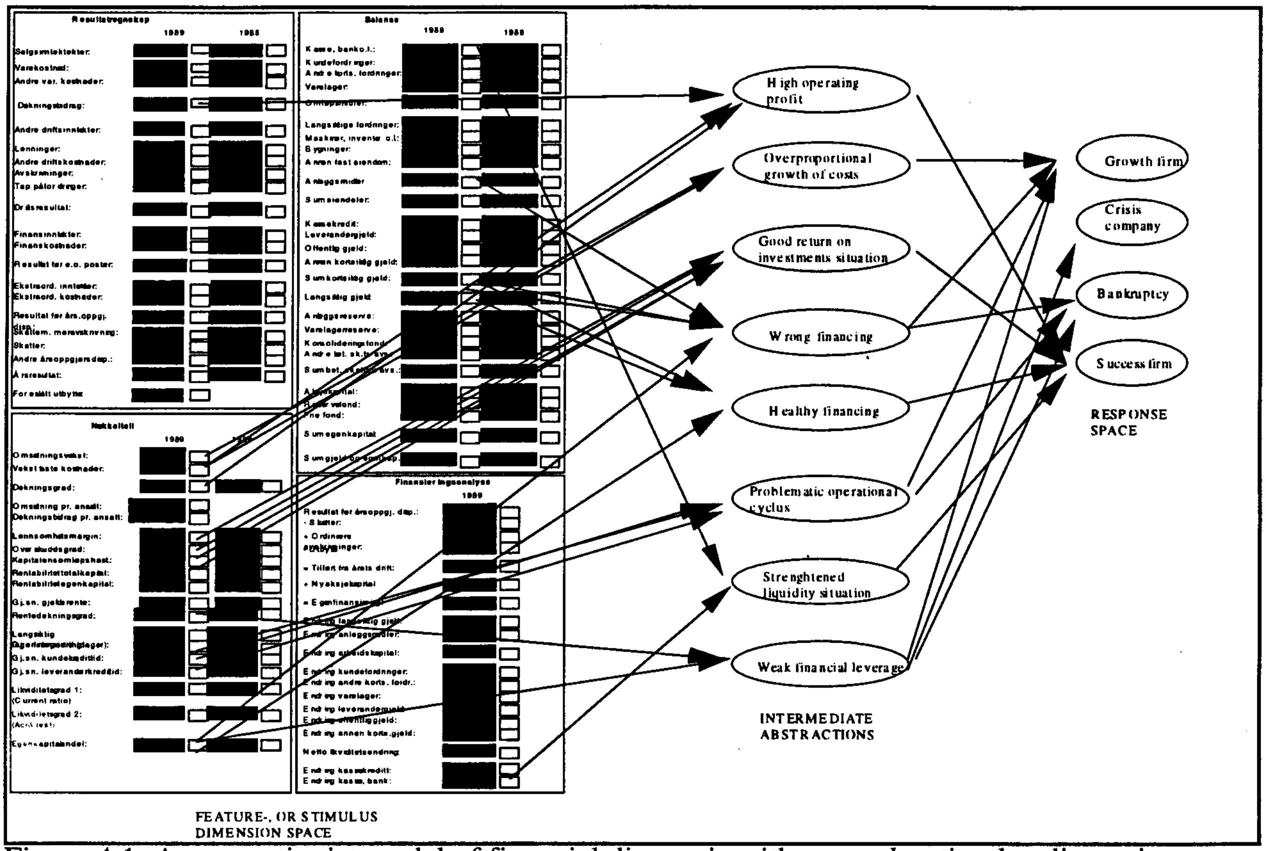

Figure 4.1. A connectionist model of financial diagnosis with *example* stimulus dimensions, intermediate abstractions and response classes

In figure 4.1, the relevant stimuli are found among the financial statement information on the left of the figure. The actual stimulus dimensions or features used by the model are not determined. However, the summary in chapter 2 indicated the theoretical concepts traditionally assumed operationalised in financial diagnosis. Among these, cues covering the theoretical concepts "leverage", "profitability", "liquidity" and "financing"[1] were considered the four most important.

On the right side, some classes of the classificatory diagnosis are shown. In figure 4.1, these are shown as units indicating linguistic terms presumed to be relevant in financial diagnosis. The actual linguistic terms used, must be determined empirically. However, the different task contexts of financial diagnosis shown in chapter 2, suggest a space of relevant responses[2].

If the financial diagnosis is a complex task, intermediate abstractions are presumed functional to the cognitive classification. In figure 4.1, the intermediate abstractions are shown as "feature detectors". Other possibilities of intermediate abstractions exist, such as intermediate

[1] See the summary table 2.7 of concepts assumed operationalised by financial cues in studies of financial diagnosis.

[2] Linguistic terms expressing for example a bankruptcy or going concern classification, or ordinal scales as in bond rating contexts are examples of relevant responses.

abstractions working as variables, or as exemplar detectors identifying similar exemplars belonging to the same subclass. Some examples of different types of intermediate abstractions are shown in figure 4.2. The intermediate abstractions actually developed by the model must be determined empirically, but in chapter 2, most of the higher level concepts presumed relevant as intermediate abstractions were introduced. The intermediate abstractions developed by the connectionist model can be analysed and evaluated against these concepts to assess the cognitive relevance of the model.

In figure 4.1, it is assumed that connections illustrated by arrows, constitute the links between stimulus dimensions and intermediate abstractions, and between intermediate abstractions and the classificatory diagnosis. When implemented in a backpropagation network, these connections have weights indicating the diagnosticity of stimulus dimensions and of intermediate abstractions. Activation of units representing intermediate abstraction and responses are presumed to be a function of the weighted sum of diagnosticity of units in a layer below. By implementing the model with local input units, the model may develop local internal abstractions corresponding directly to concepts of relevance to financial diagnosis. The use of sigmoid output functions of intermediate and output units allow such a development[1], but does not enforce it. To be interpreted as classificatory response, local representation of responses must be used. Parameter estimation is done by using the traditional backpropagation learning rule explained in chapter 3.

From the presumptions above and the backpropagation implementation of our connectionist model of financial diagnosis, some interesting propositions can be made.

A first proposition of a new model, is traditionally made of its capacity to model the investigated phenomena. Since financial diagnosis tasks have been modelled with three different approaches, several different models have been developed that could provide a *benchmark* for the evaluation of our new model. However, studies with a cognitive approach was shown in chapter 2 primarily to have modelled information processing behaviour, and no general model of financial classificatory diagnoses could easily be found suitable as a benchmark. Predictive studies have focused on modelling financial diagnostic criterion variables, and benchmark models from this approach are not suitable as models of judgmental variables, such as financial classificatory diagnoses. However, it was shown in chapter 2 as part of the standard assumptions of the judgement modelling approach, that linear models showed good fit to actual financial diagnostic classifications. Thus, linear models of the judgement modelling approach could be used as benchmarks for a strong test of connectionist model fit. Consequently, we can formulate the following proposition:

[1] See the presentation of the "receptive fields" of such units in a backpropagation network, shown in figure 3.12.

P1: The connectionist model will show better fit to financial diagnostic classifications than linear benchmark models.

This proposition implies a far stronger test of the developed model than a proposition that the model fit should be better than chance, often used in similar tests (e.g. Kryzanowski et al., 1993). The linear benchmark models should be developed following the traditional principles of the financial diagnosis literature reviewed in chapter 2.

If the linear models are outperformed by the connectionist model of financial diagnosis, several reasons can be proposed. First, the functional form of the mapping in the connectionist model is nonlinear, and this nonlinearity may in itself be important. Second, the method of parameter estimation used in our model is the backpropagation learning algorithm, and this estimation method differs from the method traditionally applied to linear models. Third, the intermediate abstractions developed in our connectionist model are considered important to the model's behaviour. Of these three reasons, the intermediate abstractions are what is presumed to give the model its cognitive relevance. Thus, a test of the effect of the network's intermediate abstractions is necessary, and can be formulated in the following proposition:

P2: The connectionist model with intermediate abstractions will show better fit than similar models without intermediate abstractions.

The linear models are estimated with traditional methods[1]. If the connectionist models without intermediate abstractions are estimated using the backpropagation learning algorithm, the difference in model fit between the connectionist models without intermediate abstractions and the models with such representations can be used to test proposition P2. In a backpropagation model, the intermediate abstractions are implemented in hidden units, and consequently, a test of the differences in performance between models with and without hidden units can be used.

The intermediate abstractions of our model can take many forms. Common to all forms is that they can be analysed with connectionist analysis methods (e.g. Hanson & Burr, 1990; Hinton, 1989; Sanger, 1989; Shakey, 1991). Such analysis should reveal the internal representations and their role in processing. For the model to have cognitive relevance, the intermediate abstractions should be interpretable in terms and concepts identified as relevant to the financial diagnosis task. Even if representations of these concepts are distributed, behaviour

[1] For example, ordinary least squares estimation methods.

of the system should be interpretable in such terms. Thus, the following proposition can be made:

P3: Intermediate abstractions in a connectionist model of financial diagnosis have cognitive relevance.

Several intermediate abstractions can be found, and the functional abstractions must be determined empirically. However, based on previous applications of connectionist models, some suggestions can be made. In figure 4.2, examples of intermediate abstractions working as "feature detectors", as variables, and as "subclass identifiers" are shown.

The intermediate abstractions shown in figure 4.1 are "feature detectors" that work as local representations of specific or configural features in the stimulus material. Representations operating as "feature detectors" bear close similarities to rule-based representations, because outputs of such units will typically show bimodal distributions. Another representation is variables, which is also a local representation, but the outputs of units representing such intermediate abstractions are continuous and show unimodal distributions. In the "subclass detectors" shown on the right of figure 4.2, values of particular features or variables are distributed among the units, and a place encoding more similar to what is found in exemplar based models (Kruschke, 1992) is developed. Independent of the actual form or type of intermediate abstractions developed, their cognitive relevance can be determined by evaluating their role in the processing against concepts and terms found to have cognitive relevance in studies of financial diagnosis reviewed in chapter 2.

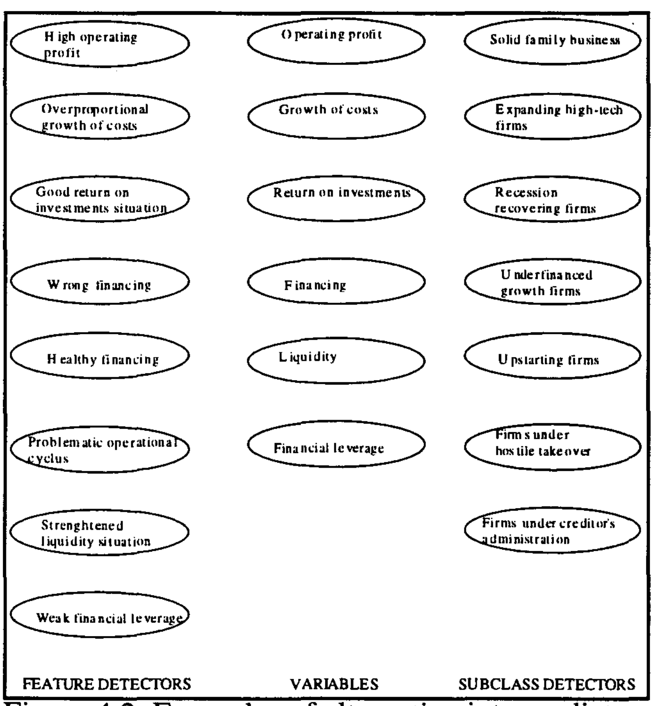

Figure 4.2. Examples of alternative intermediate abstractions with varying degree of locality

With local representations of stimulus and response, the backpropagation model is, in principle, free to develop any representational form of intermediate abstractions. However, previous research (Hanson & Burr, 1990) suggests that internal representations are sensitive to representations of stimulus and responses. In the financial diagnosis task, stimulus representation is constrained by the variables form of the financial cues. However, different response representations can be tested. The form of the response representation is likely to affect the intermediate abstractions in the connectionist model of financial diagnosis.

The way the representation of responses will affect internal representations is also an empirical question. However, we suggest that the variables form of response representation will make similar internal representational forms evolve. Furthermore, class representations of responses, such as linguistic terms, will make a variables form of the internal representations less likely.

Before proceeding to the empirical investigation of our propositions, it should be stressed that the propositions do not have a traditional hypothesis form. A main reason is the problems associated with tests of cognitive models that we addressed in chapter 3. Of particular relevance are the problems with a formal test of proposition P3. The formulation of this research question as a proposition illustrates that we accept that the best we can do is to make P3 probable, more than formally test it.

PART III - METHOD

Chapter 5. Research design

To investigate the propositions set out in chapter 4, a data set of financial diagnoses must be provided. As in any other empirical study, the quality of the data set depends on how threats to validity are treated (Cook & Campbell, 1979; Pedersen, 1989). Consequently, the standards followed in traditional *experimental* settings should be applied independent of the model finally used to simulate the mapping of stimulus to response. Recently, several studies in experimental cognitive psychology have explicitly reported experimental procedures when connectionist simulations are applied (e.g. Estes et al., 1989; Gluck & Bower, 1988a, 1988b; Nosofsky et al., 1994; Taraban & Palacios, 1993). Similar principles are followed in this study.

In addition to the methodological aspects of the experimental study, the *simulation* methodology used in connectionist modelling requires further elaboration of the simulation environment, parameter settings and validation principles applied. Here, the methodology of the empirical study and the simulations are presented in chapters 5 and 6, respectively.

In classical experimental designs, validity of the experimental data is obtained by manipulation, control and randomisation. In a financial diagnosis experiment, *manipulation* is the controlled changes in financial cue patterns presented to the diagnostician. By *control* we mean control of the experimental conditions in which the diagnosis takes place. *Randomisation* can be used to enable comparisons across manipulated levels of an independent variable. Here, it is used to randomise subjects to separate groups given different manipulations. Following these principles, the "ideal" experimental conditions of a financial diagnosis experiment are manipulation of a single, or small number of financial cues of diagnostic importance in an artificially controlled setting with subjects randomised to groups given different values of the manipulated experimental variable. However, such a design has several general problems and practical limitations. To take an example, the manipulated variables in financial diagnosis most likely consist of financial statement cues that covary in a limited number of patterns, so that individual cues can not be manipulated freely. Consequently, some form of quasi experimental design must be used.

At the other extreme of the range of possible research designs is the use of a secondary data set. When providing secondary data sets of financial diagnoses, such as official bond ratings or going concern evaluations, at least two main problems are present. First, in Norway, availability of official information on, for example, bond ratings or similar evaluations, is very limited. At the time of this study, no official sources of such information existed. At the present time, some sources are available[1], but their quality is not easily determined. Even if official data were available, their quality as *experimental* data would be outside the control of the experimenter. Thus, experimental data on financial diagnosis classifications by qualified subjects had to be collected under controlled conditions. An experimental design based upon the general principles of manipulation, control and randomisation, and applied to the financial diagnosis task, was set up.

As explained, the best way to provide controlled conditions is to create an artificial situation attaining free manipulation, full control and randomisation. In financial diagnosis, however, these conditions sensitise other threats to validity. For example, artificial manipulation of financial statement data was applied in previous research (Holt & Carroll, 1980; Lyngstad, 1987; Methlie, 1993), and was easily recognised by the diagnosticians as unrealistic or impossible statements. Thus, cue patterns should be provided from real financial statements even though this could prevent sufficient variation in the manipulated variable.

"Natural" manipulation is provided with selection of stimuli from the real world. Even though not all financial cues are normally distributed, a bell shaped distribution is likely. If selection of financial statements are made at random, many levels of the manipulated variables will be produced. Consequently, a considerable number of diagnosticians is necessary to validate the diagnoses at all levels of the manipulated variables. To illustrate the situation, consider the following two examples: In expert systems development, one traditionally uses one or a few experts' diagnoses of several cases to provide sufficient variation in the financial statement cue pattern for induction to take place. Then, there are many levels of the manipulated variable, but only few subjects to validate the diagnoses. In classical experiments, one traditionally uses two levels of manipulation and several subjects in order to test the effect of the manipulation while controlling for individual variation. In this study, we suggest a design that compromise between these two extremes, so that data, relatively free from between-subject variation, can be provided on a sufficient number of diagnostic cases.

5.1 The stimulus material

The stimulus material in financial diagnosis primarily consists of financial statement cue patterns. Official financial statement content and presentation structure are regulated by the "Accounting Act"[2] and other laws depending on the ownership structure and size of the firm[3].

---

[1] The present sources are the Central Bureau of Statistics, Dun & Bradstreet Soliditet, the Register of Company Accounts at Brønnøysund and some industry related registers. However, the access to, the prices, the contents and the completeness of these registers vary considerably. The contents and availability of financial statement information in these registers are discussed in section 5.1.

[2] In Norwegian; "Regnskapsloven".

[3] Such as in the Norwegian "Companies Act" (Aksjeloven).

However, previous research on financial diagnosis of Norwegian firms (Lyngstad, 1987; Methlie, 1993), and the research reviewed in chapter 2 (e.g. Bouwman, 1983; Bouwman et al., 1987), suggest that official financial statements need careful preparation and reorganisation for financial diagnosis to take place. Studies reviewed in chapter 2, particularly stressed the importance of financial ratios in financial diagnosis. Ratios covering the most widely applied theoretical concepts in financial diagnosis; "profitability", "leverage", "liquidity" and "financing" must be provided in the stimulus material, or they will be calculated by the diagnostician. In addition to the traditional financial statement and the financial ratios, funds flow statements have been suggested important to financial diagnosis (Gentry et al., 1985). In order to secure that most of the material required by the financial diagnosticians was provided, we chose to incorporate income statements, balance sheets and 19 financial ratios of two consecutive years, and a funds flow statement of the most recent year in the stimulus material.

The income statement was organised according to the "Accounting Act" with some alterations recommended in previous research (Lyngstad, 1987; Methlie, 1993, 1994). The alterations were that contribution margin was calculated, and that the year end adjustments were organised for analysis purposes following the regulations given in Norwegian tax laws of the most recent year of the financial statement included. Financial and extraordinary items were presented by sums, thus the income statement cues were rather coarse.

Similar recommendations were followed for the organisation of the balance statement. A few more summarising items and a somewhat more detailed specification of tax related reserves than recommended in the "Accounting Act" were used.

The income statement, the balance statement and the financial ratios were placed in separate sections of the stimulus material. The ratios were selected among the most frequently used financial ratios of chapter 2, and were used by Norwegian financial diagnosticians (see e.g. Lyngstad, 1987; Methlie, 1993 p. 152-153). The selected ratios were presumed to indicate aspects of the theoretical concepts "operations", "productivity", "profitability", "financing", inventory, collection and payable "turnovers", "liquidity", and "leverage" of each firm. The ratios were grouped in blocks in the stimulus material corresponding to these theoretical concepts. The ratio section of the stimulus material included all ratios recommended by two widely applied Norwegian standard textbooks (Eklund and Knutsen, 1994; Kinserdal, 1992) on traditional financial statement analysis of Norwegian companies.

The funds flow statement was organised according to the recommendations of a presumed expert on financial diagnosis (see Lyngstad, 1987). The subjects participating in this study had previously been given an introduction to this particular way of organising the funds flow statement.

Additional information on the number of employees, firm industry and location of the firm had been indicated as useful in previous studies (Lyngstad, 1987; Methlie, 1993, 1994), and was included in the stimulus material. Industry was indicated using industry terms following the industry classification of the Norwegian Central Bureau of Statistics[1], and location was indicated by the home region[2] of the company.

Much of the information provided in the stimulus material is available in official financial statements. However, additional information had to be collected regarding firm industry, and all cues of the ratio section and funds flow statement had to be calculated from the official information.

An example of the introductory text, the stimulus material, and the response form used in this study is presented in appendix B. Formulas used to calculate the cues of the ratio section are shown in appendix D. Even though the subjects were only expected to use parts of this material, *selection* of the relevant parts was left to the subjects. Selection of the relevant parts of the stimulus material is treated as part of measurement development and is presented in section 5.4.

When the format and content of the stimulus material were decided, the firms representing the stimulus *manipulation* had to be selected. This selection is a part of measurement development because it represents operationalisation of the levels of the manipulated variables. However, since the selection is closely related to the content and format of the stimulus material, it is treated in this section. The sample of firms selected to represent the stimulus manipulation is termed the *stimulus sample.*

In Norway, no computerised and official sources of financial statement information similar to, for example, COMPUSTAT are available[3]. Norway is a small economy, and selecting statements from listed companies could result in the diagnosticians identifying the companies. Smaller companies are not generally listed, and are consequently not so easily recognised.

---

[1] The industry group level descriptions of the Norwegian "Standard for næringsgruppering" was used. These descriptions are the Norwegian version of the International Standard Industrial Classification of all Economic Activities (ISIC).

[2] The "regions" used were the Norwegian "fylke".

[3] In Norway, official financial statements are registered at the Register of Company Accounts at Brønnøysund. However, this register is only in paper format and is expensive in use. Later, Dun & Bradstreet-Soliditet have created a register of financial statement information (Dun & Bradstreet-Soliditet, 1995), but this is expensive and so far only delivers information with limited content and in paper format. At the time of data collection this source was not available. Other sources only contain partial financial statements.

Furthermore, official market evaluations of the value of smaller firms do not exist, and thus, diagnoses of such firms depend more upon financial statement information, making financial diagnosis even more relevant to these firms. The selection of small firms was thus considered to further enhance realism of the stimulus situation. A register of small company financial statements had been collected for other analysis purposes (Boye & Kinserdal, 1992) at the Norwegian School of Economics and Business Administration[1]. This register was used as a sampling frame for the stimulus sample in this study. The number of financial statements required was calculated[2]. to 75 The data collection was performed in two stages. At the first stage, financial statements were selected at random from the sampling frame, and checked for reporting and calculation errors[3]. At the second stage, otherwise correct financial statements were excluded because of extreme values or particular incidents in their recent history. When the data collection was finished, 120 statements had reached the first stage, and 85 statements had reached the second stage. A list showing the identity of the 85 firms reaching stage two is supplied in appendix C. Of the 85 statements, three were excluded due to issuance of new shares (AH, AV, AY), four due to extreme values on financial ratios (BO, G, X, Z), one due to an extremely negative situation (AM), one due to extreme growth (CB), and one due to extreme fluctuations (E) in recent history. The final 75 financial statements represented the stimulus sample frame from which the presented stimuli of each subject was selected. Summary statistics of selected items from the financial statements are shown in table 5.1.

Item[4]	Mean	St. dev.	Minimum	Maximum
Sales	9 835 213	4 917 314	2 368 396	24 818 146
Operating profits	314 324	356 050	-295 098	1 946 510
Total assets	4 200 354	2 071 344	788 514	8 711 186
Equity	434 558	385 374	-312 571	1 557 897
Employees	16.9	11.2	2	54

Table 5.1 Summary statistics of firms in the stimulus sample

As can be inferred from table 5.1, the values of the financial cues of these financial statements varied considerably, and was expected to show sufficient variation in financial cue values to represent true stimulus variation. A stimulus situation was created from these data by randomly selecting a financial statement, and asking for diagnosticians evaluation of the financial situation of the firm represented by this information.

---

[1] At the time of the collection, the Register of Small Firms at the Norwegian School of Economics and Business Administration contained 155 randomly selected copies of financial statements of manufacturing and retailing companies with sales less than NOK 25 mill. from selected industries in the Register of Company Accounts. The register was later revised, and the revised version was used in Boye and Kinserdal (1992).

[2] The procedure for calculating the stimulus sample size is explained in section 3.1.3, and depends upon the number of diagnosticians available, the time to perform the diagnosis, the number of diagnoses performed with sufficient concentration on limited time by the diagnosticians, and the size of the composite judge committee considered sufficient.

[3] Reporting errors were evaluated against the report formats specified by the Norwegian Accounting Act, and calculation errors were tested using a database developed for stimulus material production.

[4] All financial items are in NOK.

To measure the dependent variables related to diagnosis and control variables, response forms were added to the stimulus material. The response form is shown in appendix B. In addition, the stimulus material contained an introduction also read aloud by the experimenter, and measures of selected education and experience variables. This material is also shown in appendix B. Finally, a list similar to the one provided in appendix D[1], with the formulas used for calculating ratios, was supplied.

A sample of the stimulus material was used in a pretest with seven PhD students in accounting and finance as subjects. The time used to complete one diagnosis was recorded, and general comments on the stimulus material were given in a debriefing interview with the subjects. The average completion time was 15 minutes, and several suggestions regarding changes in the introductory text were given. Suggested changes were done in the introductory text, but no changes were made to other parts of the stimulus material[2].

5.2 Subjects

In this study there are two samples. The sample of stimuli was introduced in section 5.1. The *sample of subjects* exposed to the stimuli is introduced in this section. A primary requirement of subjects performing financial diagnosis is sufficient knowledge of the task. Beyond this requirement, research reviewed in chapter 2 (e.g. Bonner & Pennington, 1991), indicated small differences between subjects with different experience of the financial diagnosis task. This suggests post graduate students in accounting or auditing may be used as surrogate financial diagnosticians in general.

A full class of post graduate students in auditing[3] was given to be at our disposal for one hour in their final semester. The maximum number of students regularly following lectures was estimated at approximately 150. As will be thoroughly explained in section 5.3, this figure limited the maximum number of diagnoses obtainable. A total of 108 students responded to the stimuli set up in the experimental situation.

The students varied with respect to graduate education, experience and latest professional position. Summary statistics of the subject sample are shown in table 5.2.

---

[1] The list of the stimulus material used Norwegian terms.
[2] The introductory text shown in appendix B is the final text used after changes had been made.
[3] Students at the "Høyere revisorstudium" (HRS), a study qualifying for state-authorisation of auditors in Norway.

Variable	Statistics	
Post graduate experience (numbers)	Experienced:	96
	Inexperienced:	12
Total post graduate experience (years)	Mean:	3.10
	St.dev.	2.99
	Minimum:	0
	Maximum:	20
Post graduate experience excluding inexperienced (years)	Mean:	3.46
	St.dev.:	2.96
	Minimum:	0.5
	Maximum:	20
Education (numbers)	Business graduate:	27
	Auditing graduate:	72
	Law school:	8
	Econ. graduate:	1
Latest position (numbers)	Assistant auditor:	8
	Junior auditor:	53
	Senior auditor:	21
	Supervising auditor:	2
	Accountant:	5
	Other:	7

Table 5.2 Summary statistics of sample subjects

From table 5.2, it can be concluded that the majority of the subjects had prior experience relevant to the financial diagnosis task. Furthermore, subjects without such work experience had relevant graduate education, and had gained experience of financial diagnosis tasks through their education. Consequently, the knowledge criterion of subjects in financial diagnosis was presumed satisfied, even though no subjects necessarily were experts on the task. Controlling for other demographic variables in the sample was considered less relevant because the recruitment to the educational program made subjects rather homogeneous. A homogeneous sample was fortunate to us because we were primarily concerned with the effects of stimulus variation, and not with variation in problem solving strategies or task performance across subjects.

Despite the homogeneity of the sample, subjects were expected to show some individual task variations. Previous research reviewed in chapter 2 (e.g. Chalos, 1985; Iselin, 1991; Libby & Blashfield, 1978), indicates that reduction of individual task variations and better performance[1] can be achieved simultaneously by using composite judge diagnoses.

5.3 Treatments and procedures

A primary purpose of this experiment was to provide a sufficient number of financial diagnoses, while simultaneously controlling for individual variation and other relevant threats to validity. Forming composite judge committees was considered necessary to perform this control, even though it would reduce the total number of diagnoses available. To determine how many diagnoses could be obtained, we started with the constraining conditions given to us. Pretests had indicated that proper financial diagnosis required minimum 15 minutes when full stimulus material was provided. It was further assumed that subjects' attention beyond 45 minutes of intense concentration, was difficult to obtain. Consequently, a maximum of three financial diagnoses and measures of the control variables could be obtained from the subjects for the duration time of the experiment of one hour.

---

[1] In terms of prediction errors.

If the stimulus sets had been utilised to maximise the number of diagnoses, 324 firms[1] could have been evaluated. However, this would imply no control of individual variation in diagnostic behaviour. Instead, we followed the recommendations of the judgement modelling approach, that a composite judge committee of minimum three subjects is sufficient to control for individual variation in financial diagnosis tasks[2]. A further advantage of using composite judge committees was that manipulation check could be performed.

A main problem with using composite judges in this study was that since participation in the experiment was voluntary, the number of participating subjects could not be determined exactly in advance. However, the maximum number of subjects participating was estimated[3] at 150, and the professor in charge of the class estimated the minimum number of subjects participating at 75. Consequently, we required a treatment plan that gave a minimum composite judge committee of three diagnosticians for each firm if the minimum number of subjects participated. Since each subject could maximally perform three diagnoses, the total number of firms that could be evaluated was 75. This figure was estimated early in the semester, and prior to the stimulus sample data collection described in section 5.1.

Furthermore, we wanted our treatment plan to be insensitive to more than 75 subjects participating, and to control for order effects, level effects[4], industry group effects, mortality and communication among subjects. To perform this control, we decided to use randomisation. A booklet of three financial statements was made by drawing, without replacement, from the sample of stimuli. Since this sample consisted of 75 financial statements, 25 booklets could be made from the stimulus sample. Since the maximum number of subjects that could participate was 150, we had to design 150 such booklets. The stimulus sample was used as a sampling frame for the design of the booklets six times. As a consequence of this randomisation procedure, all financial statements were placed in different contexts, the probability that subjects sitting close to each other would be given the same financial statement was low, and mortality effects were expected to be random.

The experiment was performed in an ordinary large classroom. The booklets were distributed from the back of the classroom from two sides of the classroom simultaneously. Each of the two distributors had 75 booklets placed in the order they had been designed. This distribution

[1] 108 subjects performing three diagnoses of different firms gives a total of 324 diagnoses of different firms.
[2] Libby (1981) suggested non-interacting composite judge committees could be used to control individual task behaviour variation and simultaneously increase predictive performance when individual performance error was random. Some experiments on the size of the composite judge committee have been performed (e.g. Libby & Blashfield, 1978, see Libby, 1981, p. 113) suggesting a minimum size of three judges.
[3] The full class consisted of 150 registered full time students.
[4] Level effects can occur when a subject is only presented stimuli likely to give similar responses, for example, when only success firms are presented.

method should provide as many diagnoses as possible of each financial statement, while simultaneously reducing the unfortunate effects if many subjects sitting close to each other decided not to participate in the experiment after the stimulus material had been distributed.

The experiment gave 324 diagnoses of the 75 financial statements, with an average composite judge size of 4.32. Fortunately, only one of the composite judge committees contained less than three subjects. Even though this was below the recommendation of three judges per diagnosis, we decided to incorporate all 75 firms in the resulting analysis.

5.4 Measurement and properties of measures.

Treatments represent different stimuli presumed to give different diagnoses as responses. The selected cues of the financial statement are operationalisations of stimuli and are treated as independent variables. Responses collected by measurement instruments of the response forms are operationalisations of diagnostic responses and are treated as dependent variables. In this section, we first describe the operationalisations of the relevant stimuli - the independent variables. Next, we treat the operationalisations of responses - the dependent variables.

Our definition in chapter 2 presumed that selected financial statement cue patterns were the relevant stimuli of the financial diagnosis task. We provided our stimulus material with most of the cues that could be parts of such relevant cue patterns. Some of the studies reviewed in chapter 2 assumed that the complete financial statement cue pattern was the best operationalisation of relevant stimuli (Bouwman, 1983; Chalos, 1985; Rodgers & Housel, 1987; [illegible]ers & Johnson, 1988), while other used selected parts of the financial statement as relevant stimuli[1] (Casey, 1980a; Libby, 1975). Several methods have been used for the selection of relevant stimuli. In judgement modelling research, two methods have been used. Either, the subjects indicate which parts of the stimulus material are relevant (e.g. Abdel-Khalik & El-Sheshai, 1980; Selling & Schanks, 1989), or statistical analysis is used to select relevant parts of the stimulus material (Kida, 1980; Libby, 1975). In experimental cognitive studies, the experimental designs restrict the stimulus manipulations to a small number, so that relevant parts of the stimulus material are preselected using previous research (see Enis, 1988). In descriptive cognitive studies, subjects are not explicitly asked to indicate relevant parts of the stimulus, but verbal protocol methodology makes posterior selection of the relevant parts possible. Thus, subjects indicate relevant parts of the stimulus indirectly.

---

[1] Even qualitative descriptions of small parts of the information in the financial statements have been assumed as the relevant parts of the stimulus (Schepanski, 1983).

As explained in section 5.1, previous research and pretests were used as a basis for the selection of the *content* of the stimulus material exposed to the subjects. Following the principles used in many judgement modelling and cognitive processing studies, subjects' indication of cue usage guided the selection of *relevant* cues in the stimulus material to be used as the final set of independent variables. The reasons for this two-stage procedure were to provide the sufficient amount of relevant information in the stimulus material, and to secure realism in the manipulations. Using subjects' indications of cue usage to select the final set of independent variables is well established in, for example, previous judgement modelling studies (Abdel-Khalik & El-Sheshai, 1980; Selling & Schanks, 1989). In this study, subjects indicated a maximum of four cues when performing diagnosis in each of five diagnostic areas[1].

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. In our analysis, we comply with this strong evidence by using both values of a cue when it is used in a model, even though this implies the use of highly correlated independent variables.

An illustration of the ten most frequently indicated cues within each of the five diagnostic areas is provided in table 5.3.

The cues indicated as the ten most important in table 5.3 represented 70.8 % of all indicated cues. Thus, the majority of cues used were found among these ten. Several cues were used in more than one diagnostic area. In table 5.3, ratio cues are shaded. These cues were presented as an integrated part of the stimulus material, and were, as expected, the most frequently used cues. The shaded ratio cues represented 78.3 % of the frequencies of cues in table 5.3, and 55.4 % of the frequencies of all cues indicated. Thus, the single most important section of the stimulus material was the *ratio report section*. Previous researchers have made similar findings (e.g. Biggs, 1984; Blocher & Cooper, 1988; Bouwman et al., 1987), but our subjects relied even more on the ratio section of the stimulus.

---

[1] The diagnostic areas are "profitability", "financing", "liquidity", "leverage" and "general situation" diagnosis. These areas are further explained below.

Rating/Area	Profitability	Financing	Liquidity	Leverage	General
1	ROI (19.6/19.6%)	LTINV (19.0/19.0%)	CURR (19.6/19.6%)	BER (50.1/50.1%)	BER (11.5/11.5%)
2	PROMARG (9.4/33.3%)	LTL (29.6)	ACID (18.1/37.7%)	EQUITY (11.5/61.6%)	(ROI (9.4/20.9%)
3	ROE (7.7/42.7%)	BER (9.2/38.8%)	CHLIK (9.0/46.7%)	ICOV (4.6/66.2%)	ICOV (6.5/27.4%)
4	OPMARG (7.6/50.4%)	STL (8.0/46.8%)	CASH (6.2/52.9%)	ROE (4.5/70.7%)	SGROWTH (6.1/33.5%)
5	OPROF (6.6/58.0%)	ICOV (7.5/54.3%)	APT (5.8/58.7%)	ROI (2.8/73.5%)	ROE (6.1/39.6%)
6	CONTPR (6.6/64.6%)	CAPASS (3.7/58.0%)	STL (5.7/64.4%)	TOTCAP (2.7/76.2%)	EQUITY (4.7/44.3%)
7	ICOV (6.5/71.1%)	APT (3.1/61.1%)	CURRASS (4.7/69.1%)	RES (2.0/78.2%)	CONTPR (3.5/47.8%)
8	PLBEI (3.9/75.0%)	AIR (3.0/64.1%)	LTINV (2.9/72.0%)	LTINV (1.9/80.1%)	PROMARG (3.5/51.3%)
9	SGROWTH (3.6/78.6%)	ROI (2.6/66.7%)	CHWORK (2.8/74.8%)	FREERES (1.9/82.0%)	LTINV (3.3/54.6%)
10	CGROWTH (2.7(81.3 %)	CURRASS (2.5/69.2%)	ITURN (2.7/77.5%)	STL (1.9/83.9%)	OPMARG (3.2/57.8%)
Responses	1166	1047	1007	766	489

Table 5.3. The ten most frequently indicated cues used on different diagnostic areas (individual and cumulative probabilities in parentheses)

Further selection of relevant financial statement cues are often made model dependent. If, for example, regression analysis is used, multicolinearity or normality assumptions are important. When we develop the benchmarks for performance comparisons of our models, we take these assumptions into consideration. However, if subjects show sensitivity to, for example, correlated independent variables, this sensitivity should be part of a *cognitive* model (Kruschke, 1993a). Consequently, we used all the resulting ratio cues as independent variables in our connectionist models[1]. With this selection, a majority of the cues indicated used by the subjects were used in the models. In addition, the ratio cues are often presumed to be size independent indicators (White, Sondhi & Fried, 1994, pp. 198-199). For ratio cues with two values, one for each of the two consecutive years of the financial statements, both cue values were applied as independent variables.

Descriptive statistics and correlation matrices illustrating the distributions of, and correlations between, the independent variables are supplied in appendix E and F. From these appendices, it is obvious that several problems when using the 32 variables as independent variables in traditional models were present. First, 20 of the 32 ratios had distributions that differed significantly from the normal distribution, causing problems with models presuming normality or multinormality, such as discriminant analysis. Second, 37.5 % of the 496

[1] The ratio cues "sales per employee" and "contribution margin per employee" were not indicated used by any subjects, and were excluded from further analysis.

relevant correlation coefficients computed between the independent variables were significantly different from zero (α=0.05). This indicated problems with using the independent variables in models presuming no multicolinearity, such as traditional linear regression.

As stated in chapter 2 and 4, traditional linear models, such as discriminant and regression analysis, have shown good fit to financial diagnostic data in judgement modelling studies. Consequently, such models were suggested as good benchmarks for evaluating the connectionist models' performance. However, these models must be developed taking care of their normality and multicolinearity presumptions. The recommended procedure in judgement modelling studies, and in predictive studies, is to use factor analysis to select independent variables with as few of these problems as possible (see Libby, 1975, p. 153), or to carefully analyse the original independent variables for normality and mulicolinearity problems (Karels & Prakash, 1987). The first of these procedures was used here. Two benchmarks were developed, and their performance results are reported in section 6.3. However, the factor analysis performed when developing these benchmarks provided interesting information on the structure of the independent variables, and is reported here.

The factors of the first benchmark, termed A, were extracted using principal components analysis with varimax rotation on all the 32 independent variables. The analysis revealed nine factors with eigenvalue greater than 1.00, explaining 83.3 % of the variance in the original variables. The rotated factor matrix is shown in appendix G. Interpretation of the factor loadings showed that the nine factors represented "liquidity", "year one profitability", "year two profitability", "assets turnover", "collection and payable turnovers", "leverage", "operations", "interest coverage", and "growth", respectively. The factors roughly corresponded to theoretical concepts found in previous studies of financial statement cue patterns (Gombola & Ketz, 1983; Pinches et al., 1973), and illustrated in table 2.7. Somewhat different from these was the "growth" factor with high factor loadings on the ratios computed as changes in sales and costs over the two consecutive years. Only the "liquidity", the "interest coverage" and the "growth" factors showed distributions slightly different from the normal distribution.

The factors of the second benchmark, termed B, were extracted using a similar principal components analysis with varimax rotation on the *averages* of the two values of each selected cue. Average values are often recommended for financial analysis (Kinserdal, 1992), but information on change is lost. The analysis revealed five factors with eigenvalue greater than 1.00, explaining 76.6 % of the variance in the original data. The rotated factor matrix is shown in appendix H. Analysis of the factor loadings showed that the five factors represented "profitability", "liquidity", "inventory vs. assets turnover", "collection and payable turnovers", and "leverage/coverage", respectively. Again, most of the factors corresponded rather well to theoretical concepts discussed in chapter 2. Only the "profitability" and the "leverage/coverage" factors showed distributions slightly different from the normal distribution[1].

Whether the independent variables developed in the factor analysis were diagnostic, was an open question which partly had to be revealed by evaluating their relevance in linear benchmark models, by evaluating their theoretical relevance, and by comparing them with intermediate abstractions developed in other models, such as connectionist models. However, the factor analysis revealed a structure of the independent variable very similar to what had been found in the predictive literature of chapter 2 and in research on financial statement cue patterns (Gombola & Ketz, 1983; Pinches et al., 1973). Thus, it seemed reasonable to assume that our stimuli represented realistic manipulations of the financial situations of firms, and that the variations in financial statement cue patterns were representative of variations found when exposed to financial statements in daily diagnostic work.

In our definition of financial diagnosis in chapter 2, we presumed a judgement of the financial situation of the firm was the theoretical concept that was operationalised by the response variable in the financial diagnosis task. We assumed in chapter 3 that this response took the form of a classification, leading us to argue that the financial diagnosis task was a classification task. Several operationalisations of this theoretical concept were found. In chapter 2, we showed how different task contexts demanded different response operationalisations. Within the task context of bankruptcy prediction, the classification of a firm as bankrupt or not was the most frequently used operationalisation (Libby, 1975). Judgements of the probability of bankruptcy have also been used within this context (Simnett & Trotman, 1989). A similar operationalisation was used both in the going-concern context of financial diagnosis (Kida, 1980) and in the loan decision context (Chalos & Pickard, 1985). In bond rating tasks, the classification of the financial situation of the firm took place on an ordinal scale (Lewis et al., 1988). All operationalisations of response in these studies were done with simple measurements using one ore a few indicators. At the other extreme, we found operationalisations of the judgement of the financial situation performed with complete verbal reports or in linguistic terms. These operationalisations were traditionally used in descriptive cognitive studies of financial diagnosis (Bouwman et al., 1987). We argued in chapter 3, that judgements of the financial situation could be expressed in the form of a classification.

[1] Similar analysis have been performed on year two data showing almost identical results.

In this study, we used several operationalisations of the judgements of a financial situation, with both ordinal classifications and subjects' own linguistic terms as basis. The subjects performed judgements of the level and trend of four diagnostic areas and a general situation diagnosis on a five point ordinal scale. In addition, the subjects were asked to describe the financial situation of the firm in their own linguistic terms[1]. From these indicators, three separate measures of the "judgement of the financial situation of the firm" were developed. The measurements used in the ordinal classification-operationalisation of the financial situation concept are explained first. Next, the measurements developed from the linguistic terms used in subjects' own classifications are described.

Following the recommendations of Libby[2] (1981), random diagnosis error can be reduced by calculating composite judgements. These judgements were calculated as the average diagnosis of the members of the composite judge committees. The average number of such committee members was 4.32, and for our five point measures, the traditional arithmetic mean was used. A composite judge diagnosis was computed for all 75 firms on the level and trend of profitability, financing, liquidity, leverage, and the general financial situation indicator. Summary statistics of the indicators are supplied in appendix I. The composite judgement transformations give the indicators favourable measurement properties, transforming them from ordinal classifications to almost interval scale measures. The distribution properties of the indicators also improve from the averaging performed in composite judgements.

Supported by the findings referred to in chapter 2, diagnosis of level and trend were treated separately, giving two sets of approximately interval scaled measures of subjects' diagnoses of the financial situations. Following traditional recommendations for measurement development (e.g. Nunnally, 1978), Cronbach's α (Cronbach, 1951) and item to total correlations were calculated for these measures. Furthermore, factor analyses[3] of the measures were performed. The results of these computations are shown in tables 5.4 and 5.5 for the level and trend measures, respectively.

[1] The response form is shown in appendix B.
[2] See also Ashton (1982), p. 42-43.
[3] Traditional principal components analysis was used.

Factor analysis:		Chronbach's α
No. of factors:	1	0.91
Variance explained:	73.8 %	
	Factor loadings:	Item to total correlations
Profitability	0.7473	0.7630
Financing	0.9056	0.9031
Liquidity	0.8171	0.8079
Leverage	0.8487	0.8470
General level	0.9605	0.9585

Table 5.4 Measurement statistics of the level measure of the financial situation

As shown in table 5.4, the α was high, and the factor analysis extracted one factor with eigenvalue greater than 1.00. By inspection of the factor loadings and item to total correlations, it was obvious that the general level measure of the financial situation was so highly correlated with the total summed score[1] that it represented a single and simple measure of the concept presumed operationalised. Since the distribution properties of this simple level measure of the financial situation diagnosis were acceptable, we used this level measure as the first dependent variable of our simulations.

Factor analysis:		Chronbach's α
No. of factors:	1	0.93
Variance explained:	73.8 %	
	Factor loadings:	Item to total correlations
Profitability	0.7880	0.8090
Financing	0.9181	0.9090
Liquidity	0.8451	0.8342
Leverage	0.9134	0.9085
General	0.9493	0.9524

Table 5.5 Measurement statistics of the trend measure of the financial situation

As can be seen in table 5.5, very similar results were found for the trend measure of the financial situation. The α was even higher, and the factor analysis extracted one factor with eigenvalue greater than 1.00. By inspection of the factor loadings and item to total correlations, similar conclusions as for the level measure could be drawn. A somewhat higher value of the Kolmogorov-Smirnov statistic illustrated that the distribution of this variable deviated somewhat from the normal distribution[2]. However, the advantages of using a simple single measure with good measurement properties lead us to conclude that the simple general trend measure of financial situation was sufficient[3]. Thus, this simple trend measure of the financial situation diagnosis was used as our second dependent variable.

A third operationalisation of the financial situation concept was designed by asking the subjects to indicate a noun that in their opinion best described the firms' financial situation. A list of example nouns was supplied in the introductory text, but the subjects were

---

[1] And the common factor of the factor analysis.

[2] See appendix I.

[3] An interesting property of measures developed by using summed scales or scales derived from factor analysis is that they inherit linear properties by their construction. Since both linear and nonlinear methods were to be used in this study, we preferred to use simple and unweighted measures.

recommended to use their own terms. The linguistic terms were analysed by the author, and a classification system of nine focal linguistic terms was used. These terms, and the number of diagnoses classified using each term are listed in table 5.6.

In table 5.6, the nine most frequently applied linguistic terms are ordered by the average value of the level measure of the financial situation presented above. This value was calculated as the average level value for the diagnoses where each focal linguistic term was used, irrespective of which firm was diagnosed. The table illustrates how linguistic terms were ordered. In addition, the level value of each linguistic term had a considerable standard deviation illustrating the "fuzziness" of the terms.

Linguistic term	Number	Average level value
Success	8	4.25
Solid[1]	6	3.83
Growth[2]	53	3.62
Normal	7	3.14
Risk[3]	5	3.00[4]
Stagnation[5]	6	2.83
Problem	10	2.10
Crisis	38	1.97
Bankruptcy	21	1.33
Other terms	99	
No characteristic given	71	

Table 5.6 Most frequently used linguistic terms ordered by average level value

As can be seen in table 5.6, the linguistic term "bankruptcy" is somewhat special, with an extremely low average level value. Thus, the "bankruptcy" term seemed well suited to identify a firm that was presumed to be in a particularly unfortunate financial situation. Furthermore, a focus on bankruptcy classification could make comparisons with studies in the bankruptcy classification task context of chapter 2 possible. Consequently, we chose to use a bankruptcy classification measure as our third dependent variable.

To classify firms as bankrupt firms based upon the linguistic terms used by the members of a composite judge committee, a rule had to be set up. We chose to classify the firm as "bankrupt" when the *majority* of the members of the composite judge committee had used the linguistic term "bankruptcy[6]"in *any form* in their description of the financial situation of the firm. The application of this rule resulted in the classification of 12 of the firms as "bankrupt". A test of the validity of this rule was performed by presenting four individual coders with the linguistic descriptions of all members of the composite judge committee for each firm. The coders where asked to select from the linguistic terms "bankrupt", "problem", "normal" and "success", the term they found best unifying the different descriptions given by

[1] The Norwegian term "solid" was used by the subjects.
[2] The Norwegian term "vekst" was used by the subjects.
[3] The Norwegian term "risiko" was used by the subjects.
[4] The term "risk" has the highest standard deviation.
[5] The Norwegian term "stagnasjon" was used by the subjects.
[6] The Norwegian term "konkurs" was used.

the composite judge committee members. The coders agreed[1] on 13 of the firms classified as "bankrupt". Among these were all the 12 firms covered by the original rule[2]. This result supported the applicability of the simple majority vote rule of the composite judgement classification of "bankrupt" firms used as our third dependent variable.

To summarise, we developed three dependent variables, measuring different aspects of the subjects' judgement of the financial situation of the firms. Two variables were approximately interval scaled, and were derived from composite judgements of ordinal level and trend measures of the financial situation. The last variable was a dichotomous variable indicating the classification of a firm as "bankrupt" by a majority of the members of a composite judge committee.

Since the composite judge committees were rather small and the number of firms was rather large, a fairly strong manipulation check could be performed by testing if all mean diagnoses were equal irrespective of stimulus. Simple one-way analysis of variance was performed, and the results are shown in appendix J. Both the level and trend variables showed significant differences in diagnosis value explained by stimulus, supporting our suggestion made above that the stimulus material had a significant manipulation effect.

---

[1] All coders selected the "bankrupt" term.

[2] On the thirteenth firm, discrepancy was found between the linguistic terms used by the subjects and the average level and trend values indicated. Consequently, we chose not to include this firm in the class of "bankrupt" firms.

Chapter 6. Simulation design

In modelling cognitive phenomena, operationalisation by model is traditionally applied. The principles of operationalisation by model are different for operationalisations of information processing and connectionist models. As shown in figure 3.10, connectionist model operationalisations are done by selecting a connectionist model, designing a simulation environment, and applying the model in the environment on collected stimulus-response data. The model is evaluated by analysing its generalisation performance and its representations. In this section, the *simulation environment* set up to develop the backpropagation model of section 4.2 on the data presented in chapter 5 is introduced.

6.1 Methodology of connectionist simulations

In principle, decisions on the simulation environment must be made regarding all the first seven functional components of a connectionist model[1] (Rumelhart & McClelland, 1986). The operating environment of the model is sufficiently constrained by the research design set up to provide the stimulus-response data.

Our selection of a particular connectionist model was primarily theory-driven, and was explained in part II of this thesis. However, different parameter settings of particular connectionist models can give varying performance results and final representations. As opposed to, for example, linear models, most connectionist models contain random elements requiring repeated simulations to determine their validity. As an example, a backpropagation model is sensitive to initial weights, and this sensitivity must be controlled to evaluate the performance of the model. The setting of simulation parameters is the first group of methodological decisions that must be made in connectionist simulations.

Of the seven relevant functional components of a connectionist model[2], parameters settings must be determined for the representational transformations of inputs and outputs[3], network topology, output and activation principles, and learning principles. To give some examples, representational transformations must usually be performed to fit the output functions of the selected connectionist model. Network topology decisions are relevant if, for example, network topology changes during the course of learning. Output and activation function forms are traditionally determined with the selection of a particular connectionist mode, but , for example, choice of maximum and minimum values, or the use of symmetric or asymmetric functions are methodological decisions. Similarly, the main principles of the learning rule are

---

[1] The functional components of a connectionist model are presented in section 3.2.1.1.
[2] See section 3.2.1.1.
[3] Indirectly determining the number and type of processing units.

determined by choice of model, but decisions on, for example, epoch or example based learning must be considered methodological.

The second group of methodological decisions that must be made is regarding the programming environment or design tools to be used in the model implementation. Most connectionist learning rules are computationally demanding. To optimise learning speed, specially written algorithms are often used and implemented in traditional programming languages. However, the experimenter's opportunity to interact with and alter such models, is limited. Several design tools for connectionist models have been developed to allow easy interaction with the models[1], and these are frequently used by connectionist modellers.

The third group of methodological decisions that must be made is related to how the model fit, or performance, of a connectionist model is to be evaluated. As explained in section 3.2.2, the approximation potential of some connectionist models (Hornik et al., 1989), require that proper evaluation of model fit must be performed by testing the *generalisation* properties of the model. However, generalisation can be evaluated by different *measures* and with different *test sample designs*. The selection of relevant generalisation measures and the organisation of learning and test samples for generalisation evaluations are methodological decisions. These decisions do not affect the principles of the selected connectionist model[2], but they affect the way performance is evaluated. To give some examples, generalisation measures are often not explicitly given (Gluck & Bower, 1988a, 1988b), or very simple test sample designs are set up to measure generalisation[3] (Schanks, 1992). However, modern statistical resampling techniques (Efron & Tibshirani, 1993) provide test sample designs for proper evaluation of generalisation in models with strong approximation potential.

The first and second group of methodological decisions are closely related, and in the following they are treated simultaneously in section 6.2. The choice of principles for evaluating model performance is treated in section 6.3 and 6.4.

6.2 Simulation environment parameters and design tools

When we chose to apply the backpropagation model as a model of financial diagnosis, most of the eight functional components of a connectionist model (Rumelhart & McClelland, 1986) were indirectly determined. However, simulation environment parameters, modelling tools

---

[1] For a review of some of these tools see James (1994).

[2] Different learning and test sample organisations also affect the representations of the estimated model.

[3] In experimental psychology, the term "transfer" is traditionally used about generalisation of a model to new stimulus-response patterns.

and performance measurement principles must be determined. The most important issues to be settled are (Gallant, 1993):

Representational issues:
-- Input representation
-- Output representation

Configural issues:
-- Choice of output function
-- Selection of network topology
-- Initialisation of weights

Learning rule issues:
-- Learning rate settings
-- Momentum term settings
-- Use of epoch or example based learning

Modelling tool issues

In the following, our decisions regarding each of these issues are explained and discussed. For many of the simulation parameters of connectionist models in general, and for backpropagation models in particular, no consistent formal basis for value selection has yet been found[1] (McKay, 1992, p. 450; Ripley, 1993; Smith, 1993). Many of the parameters values used in this study were selected after a series of initial experiments not reported here.

The representation of financial cues as input patterns and financial diagnostic classifications as output variables must be adjusted to fit the output functions of a backpropagation network[2]. The design of the input patterns and the alternative output variable operationalisations are explained in chapter 5, and will not be treated in greater detail here. The financial cue patterns consisted of 32 input variables in the range [-8.02, 474.96]. The bankruptcy classification operationalisation of the response variable was a dichotomous variable in the {0, 1} range, and the level and trend response variables were in the range [1.25, 4.33] and [1.00, 4.33], respectively. The distribution properties of these variables are reported in chapter 5, and illustrated in appendices E and I. To fit the output functions of the backpropagation model, the variables were scaled. A [0,1] scale is traditionally used for asymmetric output functions and a [-1,1] scale is used for symmetric output functions. We chose to use the standard logistic output function as reported in equation (3.6). Consequently, our input and output variables were scaled linearly[3] to the [0,1] scale. Except from the scaling

[1] Analytical derivation of optimal simulation parameter values is currently investigated in the artificial neural network community (MacKay, 1992).
[2] In connectionist modelling, the terms "input pattern" or "input variables" are often used of the set of independent variables, and the term "output variables" is often used of the dependent variables. Consequently, these terms are also applied here.
[3] Linear scaling is performed by subtracting the minimum value of a variable from each value, and dividing the result by the range of the variable.

of input and output variables, no other changes in their representational form were made. Thus, all input and output representations were local variables representations.

With the choice of a local variables representation of the input and output variables, the number of input and output units was determined. The resulting topology decisions were selection of the number of hidden units and the connectivity pattern of the network. In our simulations, determination of the number of hidden units was part of the simulation setup. A constructive procedure was developed, where the number of hidden units was gradually increased while performance was measured. The procedure is explained in greater detail in section 6.3. Traditionally, a backpropagation model is set up with full connectivity between layers. However, several ways of constraining network topology have been shown to improve generalisation properties of the network (Le Cun, 1989). Initially, all our networks were set up with full connectivity between layers, but experiments with constrained connectivity were performed, and these are reported explicitly in part IV.

When a backpropagation model is set up, the initial value of the weights must be determined. If all weights are initialised to zero, the learning rule of backpropagation will develop hidden units with equivalent weight patterns[1] (Smith, 1993, p. 96). Thus, weights are often initialised at small random values to enable the hidden units to develop different internal representations. Final representations are somewhat sensitive to initial weights, and learning speed is even more sensitive to initial weights. In addition, the optimal initialisation range will depend upon the size of the network and the chosen output function. However, we wanted an initialisation range that could be kept constant with variations in network topology. To find this range, a series of initial simulations was set up while monitoring learning speed. On an average, the best and least sensitive learning performance was found for initial weights 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. These ranges were used in all simulations.

As reported in section 3.2.1.3, the learning rate, η, is an important parameter in a backpropagation model. It must be set to provide fast learning, but to avoid oscillation[2] and saturation[3]. Modifications in standard backpropagation have been suggested to obtain fast learning and avoid the problems above (Fahlman, 1989; Jacobs, 1988), but after initial learning experiments, oscillation and saturation problems did not seem particularly relevant in our simulations. Thus, the standard backpropagation algorithm was used. More stable learning was obtained by setting η_h somewhat larger than η_o. Thus, η_h was set to 0.5 and η_o

---

[1] And consequently, with equivalent representations.

[2] As an example, oscillation will occur if the learning rate is too large in an error landscape with steep valleys.

[3] As an example, saturation will occur if the learning rate is too large in an error landscape with a very flat error surface.

to 0.4 in all simulations reported. Initial experiments were also performed with gradually decreasing η with the number of learning iterations, but no increase in performance was found, further supporting the assumption that saturation in local minima was not a major problem in our simulations.

As shown in equation (3.16), the original backpropagation learning rule can be modified by using a momentum term, α, to smooth weight changes. The term is used to prevent too large changes in weights from one learning iteration to the next. Initial experiments suggested a momentum term could be used to increase learning speed and prevent initial oscillation of the model. Consequently, a momentum term α of 0.093 was used in all simulations[1].

The original formulation of backpropagation referred to in section 3.2.3.1 (Rumelhart, Hinton & Williams, 1986), presumed epoch learning, but also showed that exemplar based learning approximated gradient decent in total output error if the learning rate was kept small. Initial learning experiments showed that learning was only marginally faster with example based learning. Thus, to stay as close to the original formulation as possible, epoch learning was used in all simulations.

Lacking a formal basis for the selection of several simulation parameters, most practical applications of connectionist models start with a series of initial learning and testing experiments in order to find proper configurations and simulation parameters. In this phase, flexibility in the programming environment to allow changes is important. When proper configurations have been found, the experimenter is often more than willing to trade flexibility for learning speed. Rapid configural changes can often best be obtained in connectionist model design tools, but at the expense of computational power and learning speed. However, several fast programming environments for connectionist models now exist. Surveys of the different programming environments for implementing neural networks and connectionist models have been reported by several authors (Eberhart and Dobbins, 1990; James, 1994 ; Nesvik, 1993), and these will not be reviewed or compared here. However, some of these systems were tested, and at the time of our choice (1990), only few of them satisfied our needs. Initially, we wanted the programming environment to cover the whole range of design tools from standard development tools to implementation tools with sufficient efficiency. The interface should be user friendly and the system should be able to run on a variety of platforms. We choose Neural Works Professional II as our programming environment. It provided a set of predefined connectionist models in addition to design tools

---

[1] Actually, this is a normalised momentum term found by dividing a momentum term found useful in example based learning by the square root of the epoch size. An initial momentum of 0.8 was found useful in simulations of exemplar based learning. Since epoch size was 74 in our simulations, the normalised momentum term in epoch based learning was set to $0.8 / \sqrt{74} = 0.093$.

that could be used to modify the standard models. In addition, the system had several options for controlling and measuring performance during processing. The standard system and its additional components are well documented in Neural Ware (1993a, 1993b, 1993c).

6.3 Generalisation measures

The criterion for evaluating the performance of a connectionist model is its ability to generalise to new and unseen stimulus-response patterns. With parametric methods, one can make inferences about this property while using all empirical stimulus-response patterns for the estimation of the model. In connectionist terminology, this means that all stimulus-response patterns are used in the learning phase of the model development. When measuring performance using the error on the learning sample, a *resubstitution* error rate is calculated (Ripley, 1993). Generally, but especially when the training sample is small, it is likely that this error rate is biased downward (White, 1990, p. 539).

With universal approximators, this problem is even more evident. The problem is thoroughly treated in section 3.2.2, and will only be briefly mentioned here. There is likely to be a trade-off between accuracy in the learning sample and generalisation (Ripley, 1993, p. 70). Geman et al. (1992) have named this trade-off the bias/variance dilemma. The point they make is that for a universal approximator, the estimator is likely to be unbiased when the learning sample is sufficiently large. When the training sample is smaller, overfit is likely to occur and the unbiased model is likely to have high variance. As explained in section 3.2.2, the bias/variance dilemma is attempted controlled by some way of "smoothing" the complexity of the mapping performed by the model. The complexity of the mapping should be increased with larger training samples, but not so fast that overfit may occur. Another suggestion is that learning should be stopped before overfit occurs[1], thus controlling complexity growth (Smith, 1993). To summarise, in a backpropagation model, the network training can be stopped well before convergence is reached, or the number of parameters in the network can be controlled during learning (Geman et al., 1992, p. 32).

There are two ways of controlling the number of parameters. One is by controlling the number of hidden units[2]. Another method is to introduce some way of penalising the complexity of the network as it learns[3]. Pruning[4] is the most widely used method for penalising complexity (Karnin, 1990; Weigend et al., 1991). However, even if all these methods were applied, no assurance could be given that resubstitution error equates true

[1] We will name this method the optimal stopping rule.
[2] We will name this method the optimal hidden unit rule.
[3] We will name this method the complexity penalty rule.
[4] See sections 3.2.1 and 3.2.2.

model prediction or classification error. Furthermore, no formal principles exist on how the different rules should be applied (Ripley, 1993; Smith, 1993).

As an alternative to the resubstitution error rate, different variations on holdout procedures can be used. As explained in section 3.2.2, several combinations of training and test samples can be used; a method often named cross validation (Stone, 1974). True cross validation is based on a "leave-one-out" assessment of generalisation error (Geman et al., 1992, p. 34). This procedure is extremely computationally demanding, but gives an error rate estimate with small bias (Ripley, 1993, p. 71). Consequently, an almost unbiased estimator of prediction or classification error exists (Cheng & Titterington, 1994, p.20; Efron & Gong, 1983, p.37; Moody, 1993; White, 1990, p. 539), and this estimator can be used to determine optimal complexity of the connectionist model (Geman et al., 1992, p. 33-34). By using the estimator while increasing or decreasing complexity, overfit can be avoided, and the model with optimal[1] complexity can, in principle, be determined.

Initial learning and test sample splits showed that it was possible to obtain very low average squared errors (MSE) on the test sample even when learning and test samples were selected at random[2]. In some of these test samples, the estimate of MSE is likely to be biased downward. In other splits of learning and test samples, the validated MSE was not better than chance, indicating severe overfit by the model. Consequently, resampling techniques had to be used to control the complexity of the mapping function if generalisation ability was to be properly estimated. In this study, prediction error was measured by using a true cross validated average squared error measure. Complexity was controlled by a gradual increase in model complexity. Cross validation was used to control both complexity parameters of the mapping; the stopping point and the number of hidden units. The cross validated average squared error applied in this study can be explained by using the error measure of the backpropagation model. As shown in section 3.2.1.3, a backpropagation model minimises:

$$E = \tfrac{1}{2}\sum_{\forall p}\sum_{\forall j}(t_{pj} - o_{pj})^2 . \tag{6.1}$$

However, this sum will depend upon the number of training cases. The most widely used measure of error is MSE. In this study, MSE is defined as:

$$MSE = \frac{1}{N}\sum_{\forall p}\sum_{\forall j}(t_{pj} - o_{pj})^2 . \tag{6.2}$$

---

[1] Random elements of the models will prevent *the* optimal model to be found, but close approximations are likely.

[2] This observation illustrates the high variance of the estimator.

Variations of this measure have been developed for different purposes (see Moody, 1993), but since this measure is close to the actual error minimised by the backpropagation algorithm, it was used here with the only modification that the error was averaged over the test cases in an N-fold cross validation procedure. This error has been termed cross validated average or mean squared error. The procedure of cross validation and the selected measure given above have been recommended by leading authors in the neural network literature:

> *The data-driven methods are based on cross-validation measure of network performance, cross-validated average squared error, that we advocate for general use in evaluating network performance. This is not the only appropriate or useful measure, but it offers considerable improvement over naive methods. (White, 1990, p. 544)*

Leading authors have also suggested that cross validated average squared error should be used in constructive or pruning algorithms (see Moody, 1993). However, the computational requirements of such procedures have led to the development of simplified, approximated methods (Moody and Utans, 1995). Despite the computational demands put on the procedure, we chose to use the cross validated average squared error measure in a *constructive algorithm.* Our procedure was very similar to the procedure developed by Moody[1] (1993), but we started with a set of input and output units derived from the methods applied in traditional measurement development.

Complexity was controlled by the number of hidden units and the stopping point of the backpropagation learning algorithm. The number of hidden units was controlled by starting with a model without hidden units. Next, hidden units were added until no obvious improvement in cross validated average squared error was obtained[2]. During learning, we measured the cross validated error, so that the optimal stopping point could be found a posteriori. Using this method, the optimal complexity of a model with a given input configuration was found. However, smaller cross validated average squared errors may be obtained by altering the input configuration. As in Moody's (1993) study, effects of changing input configurations were only evaluated for the set of models previously found optimal. A sensitivity analysis was used to estimate the effects of changing input configurations. Input units representing cues of small sensitivities were eliminated, and a new series of cross validated average squared errors was computed while network complexity was growing.

[1] The procedure developed by Moody (1993) is a pruning based algorithm, while we use a constructive algorithm. The Moody (1993) procedure was not known to us until the Moody and Utans (1995) paper was published.

[2] In our simulations, hidden units were added until the model contained 14 hidden units. At this point it was obvious that no improvement in cross validated average squared error could be obtained with more hidden units.

For the simulations set up with the bankruptcy classification response variable, only the constructive algorithm was used. For the simulations set up with the level and trend variables, both the constructive algorithm and the sensitivity analysis were used to determine the optimal model. The results of using these procedures are reported in chapters 7, 8 and 9.

For classificatory response variables, such as our bankruptcy classification variable, cross validated average squared error is not the only measure of error. A similar procedure can be used to estimate the cross validated *classification* error. This measure was used in addition to evaluate the model using the bankruptcy classification response variable.

6.4 Benchmarks

As explained above, evaluation of the models was done with a cross validated average squared error measure or a cross validated classification error. However, there are several ways to use this measure. A simple method is to evaluate whether the cross validated average squared errors or the cross validated classification errors are better than chance. A similar principle is often used in traditional significance tests of, for example, causal models (e.g. Rodgers & Housel, 1987), and it has also been applied to neural network models (e.g. Kryzanowski et al., 1993). In chapter 2, we showed how linear models have been used to model financial diagnosis with high outcome accuracy in judgement modelling studies. Thus, a stronger test of the connectionist models seemed necessary, since linear models have already shown to be significantly better than chance in modelling financial diagnosis. Based on the proposition P1 set out in chapter 4, strong support for the connectionist models could be provided if they could fit financial diagnosis behaviour significantly better than the models previously applied in judgement modelling studies.

Research on financial diagnosis from a predictive perspective has stressed the limitations of linear models when applied to financial data. In particular, the multinormality presumptions of discriminant analysis (Altman et al., 1981) and the normality and multicolinearity presumptions of regression analysis (Karels & Prakash, 1987), have been stressed. Traditional models should be used and designed with careful consideration of these problems and presumptions (see Libby, 1975, p. 153). We chose to design benchmark models of the same kind as in judgement modelling studies, but with consideration for the problems with applying these models to financial data. As a benchmark model for the bankruptcy classification variable, we chose to use logistic regression instead of the traditional linear discriminant analysis applied in many judgement modelling studies[1]. A main reason for this choice was that the financial cues showed distributions deviating from the normal

[1] This is a choice also made by other authors (e.g. Hopwood et al., 1994). However, classification results for the cross validated discriminant analysis models are shown in footnotes.

distribution, making multinormality presumptions even more speculative. As a benchmark for the level and trend variables, we chose to use traditional regression analysis. However, because of multicolinearity presumptions, all benchmark models were developed following the recommendations of both judgement modelling and predictive studies of financial diagnosis.

A two stage procedure recommended by several authors (e.g. Libby, 1975, p. 153; Zavgren & Friedman, 1988) was followed. First, factor analysis was used to obtain independent variables with acceptable probability distributions and low multicolinearity. Traditional principal components analysis was used, and the factors with an eigenvalue larger than 1.00 were rotated by the varimax method, and used in the further analysis. Two sets of variables were used in the factor analyses. One analysis used all 32 input variables also used in the connectionist models, and gave 9 factors. The other analysis was based upon averaged values of the 32 input variables over the two consecutive years, and gave 5 factors. Interpretations of these factors comparing our independent variable measures to previous analyses of financial cue patterns (e.g. Pinches et al., 1973), are reported in chapter 5.

Next, the factor scores obtained for each financial statement were used in benchmark models; regression analysis for the level and trend response variables, and logistic regression analysis for the bankruptcy classification variable. The analysis reported in section 5.4 suggested that the independent variables of these benchmarks had probability distributions and correlation matrices that made them suitable in regression and logistic regression analysis. However, since these models were used as benchmarks for the connectionist models, the same measures of performance were used. Thus, all measures of classification error or average squared errors were cross validated in an N-fold cross validation procedure similar to the one used for the connectionist models. The performance results of the benchmark models are shown in tables 6.1 and 6.2 for the 9-factor (A) and 5-factor (B) benchmark models respectively.

Dep.var. / Measure.	Bankruptcy classification	Level	Trend
Correct classifications[1]	88.00 %[2]		
MSE[3]	0.120[4]	0.232	0.354
Corr. target		0.059	0.083
Corr. distance from target		0.252*	0.221

Table 6.1 Performance results of the first (9-factor) benchmark (A) (* indicates significant at $\alpha = 0.05$)

[1] The percentage of correct classifications was calculated for all firms collectively, and results are shown for optimal cut-off values without consideration of misclassification costs.
[2] The percentage correct classifications for the corresponding discriminant analysis was 89.33 %.
[3] Cross validated average squared error.
[4] The mean squared error for the corresponding discriminant analysis was 0.072.

In table 6.1, several measures of benchmark performance are reported. A column of relevant measures is shown for each dependent variable measuring the subjects' judgements of the financial situation. For the bankruptcy classification variable, only cross validated classification error and average squared errors are shown. For the variables level and trend, cross validated average squared errors and correlation measures illustrating the distribution of the error terms are shown. These measures showed that error was somewhat correlated with distance from target, suggesting that the models made the larger errors on the more extreme diagnoses. This was not surprising, since it indicated that the models were somewhat "regressive".

Similar results were found for the 5-factor benchmark B shown in table 6.2. Performance results were somewhat better for the bankruptcy classification measure, but the results on the other variables were considerably worse for this benchmark. In particular, the results for the trend variable were considerably worse for benchmark B. This was not surprising, since the averaging of the independent variables made the model lose all information of the change in the independent variables. In fact, the performance on the trend variable was worse than what could be obtained by guessing the average trend value on each diagnosis. Consequently, it seemed as if different information was used for different judgements of the financial situation, but this suggestion was very preliminary.

Dep.var. Measure.	Bankruptcy classification	Level	Trend
Correct classifications	89.33%[1]		
MSE	0.103[2]	0.357	0.759
Corr. target		0.110	0.215
Corr. distance from target		0.382**	0.777**

Table 6.2 Performance results for second (5-factor) benchmark (B) (** indicates significant at $\alpha = 0.01$).

Despite some weaknesses of these benchmarks models, the best performance results from each model should provide a strong test of the connectionist models. If the connectionist models significantly outperform the best benchmarks, strong evidence would be provided in support of the connectionist models.

To provide further benchmarks for the connectionist model, regression analysis and logistic regression analysis were used as benchmarks of each connectionist model with the full set of independent variables used in the connectionist models. In addition, traditional models with

[1] The percentage correct classifications for the corresponding discriminant analysis was 88.00 %.
[2] The mean squared error for the corresponding discriminant analysis was 0.073.

stepwise procedures applied to select the best independent variables were also used as benchmarks in each connectionist model simulation of part IV[1].

[1] We use these benchmarks to provide as many of the traditional benchmarks of the judgement modelling approach as possible. The cross validation procedure applied in this study makes conclusions on the performance of these benchmarks legitimate, despite the normality and multicolinearity problems reported above.

PART IV - SIMULATIONS AND RESULTS

Chapter 7. A connectionist model of classificatory response

In this part, a series of simulation experiments are reported that were set up to model the stimulus-response data reported in chapter 5. This part consists of three chapters. Chapter 7 reports simulations of a connectionist model using the bankruptcy classification variable of chapter 5. Chapter 8 reports simulations of a connectionist model using the level and trend variables. In both these chapters, the full set of independent variables collected from the ratio section of the stimulus material, and reported in chapter 5, was used. In chapter 9, the principles explained in chapter 6 were used to reduce model complexity, and the results of these constrained models are reported using the level and trend variables. Each chapter is divided into three main sections. The first section of each chapter reports model performance results and comparisons with benchmarks. This section explores and tests the propositions P1 and P2 made in chapter 4, for each of the proposed models. The second section of each chapter reports analyses of the connectionist models' representations using the principles of Hinton (1989), Hanson and Burr (1990) and Sharkey (1991). The purpose of this section is to explore and evaluate the proposition P3, made in chapter 4. The last section in each chapter summarises the main findings of each simulation.

To understand the models and the parameter settings of the simulations, the main principles of the backpropagation model are explained in chapter 3. The measures of the stimulus and response variables are explained in chapter 5, whereas the setting of simulation parameters and the choice of performance measures are explained in chapter 6. A short summary of the conclusions that could be drawn from these simulations is found in section 10.1.

The number of stimulus-response pairs was small compared to the number of input variables when all 32 input variables were used. This suggested that most stimulus-response pairs should be used for learning. By using cross validation (see chapter 6, and Efron & Tibshirani, 1993), now suggested by several authors (Moody, 1993; Ripley, 1993; White, 1989), but originally suggested by Lachenbruch (1967) for linear discriminant analysis[1], and by Stone (1974) for regression and analysis of variance, a low biased estimate of generalisation error could be found (Ripley, 1993, p. 71). A favourable side effect of the cross validation procedure was that the effects of different initial weight values were randomised over all the 75 simulations.

[1] The method proposed by Lachenbruch (1967) is often characterised as a "jack-knife" procedure (e.g. Altman et al., 1981, p. 154), but it is actually a cross validation procedure since the error estimate is based on the "left" out observation.

To allow comparisons with previous judgement modelling and predictive approaches to financial diagnosis, a series of simulations was set up using the bankruptcy classification response variable in chapter 5. As a consequence of the cross validation procedure, a total number of 600 backpropagation models was simulated using 32 input units to represent financial cues and one output unit to represent the bankruptcy classification variable. The bankruptcy classification task context is the most widely used context in the judgement modelling and predictive studies reported in chapter 2, and the bankruptcy classification variable applied here was similar to the variables applied in these studies.

In all simulations, the relevant measures explained in chapter 5 and the relevant simulation environment settings explained in chapter 6, are used. This chapter reports the performance results of the connectionist model simulations in section 7.1 and the analysis of the connectionist model representations in section 7.2. Some of the main conclusions of the simulations are summarised in section 7.3

7.1 Performance results

Using both the cross validated average squared error and the cross validated classification error measures of model fit, the results of the connectionist models of bankruptcy classification are shown in tables 7.1 and 7.2. The tables also show the corresponding performance results of simple benchmark models, using logistic regression and stepwise logistic regression[1] on the original financial cue variables. As can be seen, the results of these simple benchmark models were comparable to the benchmarks developed in chapter 6.

Iterations:	5000	10000	15000	20000	25000	30000	Regr.[2]
Model:							
Log. regr. (all)							0.147
Log. regr. (stepw.)							0.099
HID0	0.091	0.096	0.100	0.104	0.104	0.106	
HID2	0.101	0.063	0.061	0.061	0.062	0.065	
HID4	0.096	0.063	0.061	0.063	0.063	0.062	
HID6	0.078	0.066	0.066	0.064	0.063	0.063	
HID8	0.080	0.069	0.067	0.067	0.067	0.065	
HID10	0.081	0.066	0.065	0.064	0.065	0.064	
HID12	0.075	0.070	0.068	0.068	0.067	0.066	
HID14	0.077	0.070	0.069	0.072	0.066	0.065	

Table 7.1. Mean squared error (MSE) for diagnosis of bankruptcy (N=75)

---

[1] Stepwise logistic regression was used as an additional benchmark despite the fact that these models used a constrained set of independent variables. It can be argued that these models should be compared to the constrained models of chapter 9 only, but since cross validated performance of these models were favourable, their results are consistently reported throughout the dissertation.

[2] Performance results of the logistic regression models are placed in a separate column. These models were estimated using traditional maximum log likelihood methods.

The results of the connectionist models are shown for different stopping points in the learning process. Initial experiments indicated that 30000 iterations were above the optimal stopping point, and learning was terminated at this point. Furthermore, the tables show the performance for increasing numbers of hidden units. The connectionist models are termed "HID", followed by the number of hidden units. At 14 hidden units, complexity was regarded so high that adding hidden units above this number would only degenerate the generalisation due to a considerable overfit. Additional experiments had been run with 20 and 30 hidden units confirming this assumption, but the results are not reported here.

The cross validated average squared error showed an expected pattern. It was lower for models with hidden units than for benchmarks and models without hidden units. Furthermore, all models with hidden units showed very similar MSEs. A small increase in MSE was found for an increasing number of hidden units, indicating overfit with increased complexity. The MSE decreased with an increase in the number of learning iterations for all models up to a minimum. Learning beyond this point increased the MSE, again indicating overfit.

The cross validated classification error is a much coarser measure of model fit. In addition, it depends on the setting of a cut-off value, and must combine two types of error, the error of wrongly indicating "bankruptcy", and similarly indicating "non-bankruptcy".

Iterations:	5000	10000	15000	20000	25000	30000	Regr.
Model:							
Log. regr. (all)							85.33
Log. regr. (stepw.)							89.33
HID0	89.33	90.67	88.00	86.67	86.67	88.00	
HID2	85.33	92.00	92.00	94.67	94.67	94.67	
HID4	88.00	93.33	93.33	93.33	93.33	94.67	
HID6	89.33	92.00	93.33	93.33	94.67	94.67	
HID8	90.67	93.33	93.33	94.67	94.67	94.67	
HID10	90.67	93.33	93.33	94.67	94.67	96.00	
HID12	92.00	92.00	92.00	93.33	94.67	94.67	
HID14	90.67	92.00	93.33	93.33	94.67	96.00	

Table 7.2. Correct classifications of bankruptcy and non-bankruptcy (N=75)

The cross validated classification errors shown in table 7.2, are reported for optimal cut-off values found by iteration. Misclassification costs were not taken into consideration, so the results are comparable to the results of the benchmarks in chapter 6.

As indicated in tables 7.1 and 7.2, performance was better for the best connectionist models than for any benchmark model. The best connectionist model evaluated by cross validated squared error was the model with two hidden units. The cross validated average squared error was 0.061 compared to the 0.099 of the best benchmark[1]. Even though the difference was considerable and in the proposed direction, it was, however, *not significant* at α=0.05 (t=1.33, d.f.=74)[2] .

A closer look at the performance measures showed that the results of the best benchmark model and the connectionist model without hidden units were very similar. Thus, increased performance was not caused by the functional form of the simple connectionist model or the estimation method (learning algorithm). Nor was it caused by the careful control of overfit in our procedure. Even though not significant, the increase in model fit obtained in the models with hidden layers, was caused by the internal representations built by the hidden units.

The best connectionist model evaluated by cross validated classification error, was either the model with 10 or with 14 hidden units. There was not a unique correspondence between minimisation of average squared error and classification error. However, all connectionist models showed better performance than the best benchmark, but a McNemar test of the difference in classification probabilities showed that the difference in favour of the connectionist models was not significant[3] at α=0.05.

Consequently, both measures of performance showed somewhat better performance for the connectionist models, but the differences were not significant. The lack of significant findings can be explained in several ways, and will be further discussed in section 7.3. However, the results were in the right direction, and must be considered promising. To improve the results, and to understand the mapping performed by the connectionist models, an analysis of the connectionist model representations was performed.

7.2 Analysis of model representations

To study the representations of the model and the tasks performed by the hidden units, a number of methods can be applied (see chapter 3). Here, sensitivity analysis and analysis of Hinton diagrams (Hinton, 1989) were applied.

---

[1] Notice that this performance is obtained with the standard stepwise procedure implemented in SPSS (SPSS Inc., 1990) applied to the original 32 independent variables of the connectionist model, and not with the benchmarks of chapter 6.

[2] Unless otherwise explicitly indicated in the text, the levels of significance are indicated using two-sided tests. Since the propositions P1 and P2 are formulated in favour of the connectionist models, one-sided tests could have been used. However, one-sided tests are only discussed when there are discrepancies between the conclusions that could be drawn from one-sided and two-sided tests.

[3] The binomial formula gives a probability that the two probabilities are similar of 0.125, thus the hypothesis must be rejected at α=0.05.

When judged by cross validated average squared error, the simplest connectionist models with a hidden layer showed the best results. The advantage of simplicity is important if connectionist representations are to be analysed. Because the final representations of the models are sensitive to initial weights, a series of 10 models was developed with different initial weights. To study the representations of the models, we want to utilise all stimulus-response patterns in the learning sample. However, this could easily cause overfit if the optimal stopping point found in section 7.1 was used. A stopping point, assumed to give approximately similar performance results when all stimulus-response patterns were included in the learning sample, was suggested for half the optimal stopping point of the cross validation procedure[1]. Thus, 10 versions of the model with two hidden units were developed, and learning was stopped at 7500 learning iterations[2].

The 10 models, termed A, B, C, D, E, F, G, H, I and J, showed an average squared error of 0.037 and a standard deviation of the MSEs of only 0.003. This indicated that the performance of the models was better than the cross validated models. However, the increased performance was largely due to overfit, even though learning had been stopped very early. The small standard deviation of the MSEs indicated that the performance was very similar in models developed with different initial weights.

The simplest analysis of the *mapping* performed by a connectionist model is done using some form of sensitivity analysis (Moody, 1993). Traditionally, this analysis is done by performing a fixed percentage variation in the inputs to the model and observing the effects on model response. This analysis is limited to the complete mapping of the model, and does not reveal the internal structure and processing of the model. Still, sensitivity analysis can be used to detect important input units.

A sensitivity analysis was performed by varying the input unit values by 5 % and observing the effect on the response variable. The fixed percentage variation in a value of an input unit is often termed "jogging" the unit. The "jogging" was performed by adding 5 % to the value of each input variables for each of the stimulus-response patterns, in each of the 10 models[3]. Thus, 750 observations of the effect of variable changes were recorded for each financial cue used as input to the model. A summary table of these effects is shown in table 7.3.

---

[1] Based upon initial, small scaled experiments.
[2] Optimal stopping point in section 7.1 was between 15000 and 20000.
[3] Initial experiments had shown small differences in the result with different *small* percentage variations in the input variables. Consequently, a 5 % positive change in the input variables was used.

Input unit	Mean effect	Standard deviation	Rank order
SGROWTH	-9.14	5.33	10
CGROWTH	-1.54	1.27	10
CONTPR1	0.94	1.27	15
CONTPR2	-4.46	2.85	15
PROMARG1	2.48	2.00	5
PROMARG2	-21.03	12.48	5
OPMARG1	7.58	4.96	3
OPMARG2	-24.32	14.41	3
ASSTURN1	-3.44	2.54	16
ASSTURN2	1.16	1.19	16
ROI1	3.83	2.28	2
ROI2	-26.15	15.38	2
ROE1	-2.30	1.62	13
ROE2	-6.69	3.97	13
AIR1	1.96	1.99	6
AIR2	20.05	11.96	6
ICOV1	5.08	3.46	14
ICOV2	-1.54	1.92	14
LTINV1	-12.16	7.34	9
LTINV2	-13.25	7.82	9
ITURN1	-7.28	4.50	12
ITURN2	1.94	2.03	12
ART1	0.69	1.58	7
ART2	-19.58	11.45	7
APT1	3.83	2.73	1
APT2	30.92	18.11	1
CURR1	3.44	2.62	11
CURR2	-7.97	5.14	11
ACID1	5.44	3.65	4
ACID2	-22.48	13.16	4
BER1	-0.16	1.15	8
BER2	-15.61	9.30	8

Table 7.3. The effects of jogging input values 5 % in the bankruptcy classification model (N=750)

In table 7.3, the different financial cues used as input variables are indicated. Explanations of the different cue abbreviations are found in the nomenclature. The year of the financial statement from which each financial cue is collected, is indicated by either "1" or "2" for each cue. The rank order column indicates which financial cue had the largest effect on the bankruptcy classification response variable when both values of each cue were joined together. The mean effects of a cue should be interpreted as the mean effect of increasing the value of the input unit on the bankruptcy classification response unit. Thus, increasing the ROI2 by 5 % decreased the value of the bankruptcy classification response unit by 26.15 %. The three most important input units were units representing accounts payable period, return on total assets, and operating margin. This was not surprising, and corresponded well to what could be expected. Furthermore, the most important units represented financial cue values from the most recent year's financial statement. However, the analysis was limited. First, standard deviations were large. Second, the analysis revealed nothing of the internal representations of the model.

A closer look at the *representations* of the models can be done by investigating the connection weights (Hanson and Burr, 1990). A simple Hinton diagram (Hinton, 1989) can be used to illustrate the weights of the model. The weights from the hidden units to the output unit are shown in figure 7.1.

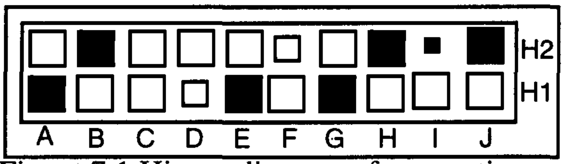

Figure 7.1 Hinton diagram of connections between the hidden and output layer of 10 bankruptcy classification models with 2 hidden units

In the Hinton diagram of figure 7.1, the size of the squares illustrates the absolute value of the weights, and the colour illustrates their sign. The connections go from hidden units number one or two, to the output units of the models A to J. The dark squares represent connections that excite the bankruptcy classification output unit, and the light squares represent connections that inhibit the unit. As indicated, the models varied considerably with respect to how the internal representations in the hidden units were used to form the bankruptcy classification. One reason for these differences in connection weights was the bias of the output units[1], and another reason was that equal representations could be obtained by reversing the sign of the corresponding connections, such as in models A and B. However, these reasons could not explain all the differences observed in figure 7.1.

To further illustrate the representations formed by the hidden units, a Hinton diagram of the connections between the input layer and hidden layer was investigated. The connections of all the 10 models, with two hidden units each, are shown in figure 7.2.

In figure 7.2, the connections are shown vertically for each model. Each of the input units are shown by their corresponding financial cue in the left column. The weights from the bias units are also shown. In the top row, model indicators are shown, and in the bottom row, the number of the corresponding hidden units are indicated. By visual inspection, we found that each model had at least one hidden unit with a set of large-valued weights. Furthermore, these hidden units seemed to be present in each model, and their weight patterns were very similar. The most frequently occurring unit of this type had a negative bias, but exceptions were found in the models E, G, and H. Model C seemed to have two very similar large-valued hidden units. Since these hidden units were shared by all connectionist models, we can term them "shared" hidden units. We do not, however, know whether such units are shared by all our connectionist models of financial diagnosis.

[1] Also termed the "threshold". See section 3.2.1.3.

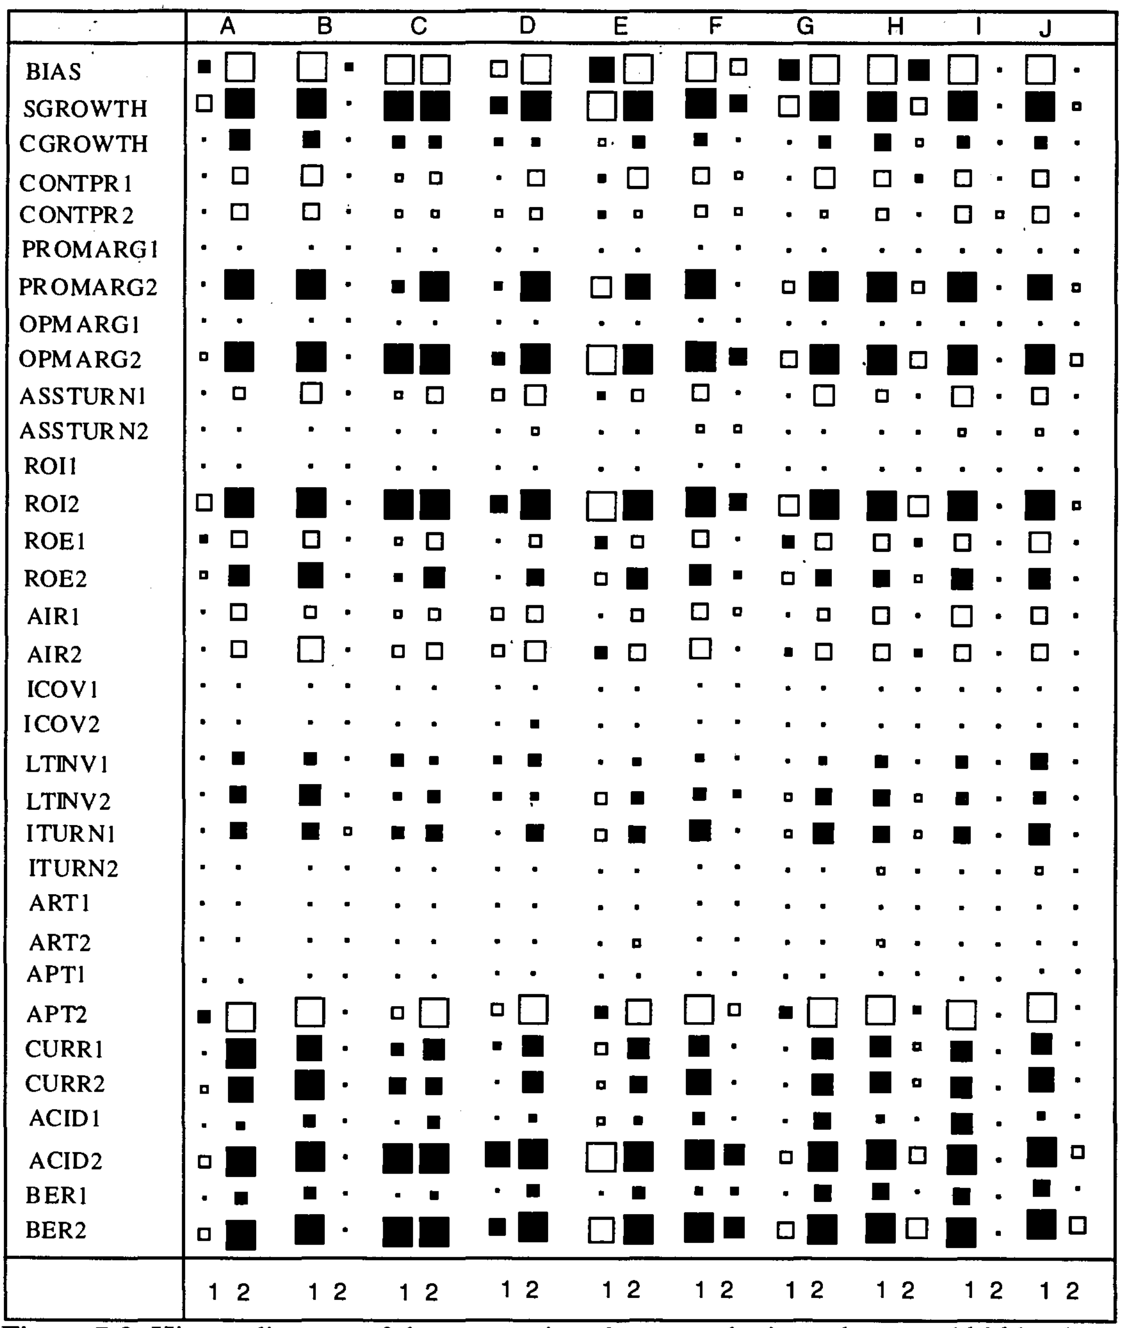

Figure 7.2. Hinton diagram of the connections between the input layer and hidden layer of 10 bankruptcy classification models with 2 hidden units

The size of a weight in a Hinton diagram can be interpreted as the importance of a unit in exciting or inhibiting another unit. In figure 7.2, the weights of the hidden units indicated the importance of a financial cue in exciting a hidden unit. Consequently, the "shared" hidden units seemed the most important in performing the bankruptcy classification. Some interesting observations could be made from a visual inspection of the Hinton diagram shown in figure 7.2. First, the "shared" hidden units had large input weights from a series of input units. Thus, the unit did not specialise on any traditional theoretical concept used in financial diagnosis, such as "profitability" or "liquidity". The units seemed to represent a variable indicating a very general "condition" concept, representing several diagnostic concepts simultaneously. Second, the signs of the weights were generally in the expected direction. Third, the largest weights were weights from the units representing cues of the most recent financial statements, corresponding well to the findings of the sensitivity analysis. Thus, the "shared" hidden unit focused on the "current general situation". Fourth, two exceptions were found. The sales growth unit had a very large weight to the "shared" hidden units, and the return on equity units had a plus/minus pattern of weights for the two consecutive years, both findings indicating that aspects of change in the value of cues from two consecutive years were relevant. Consequently, the representations formed by the "shared" hidden unit were rather complex.

From visual inspection it was not easy to detect common aspects of the task performed by the other hidden units in the models. Further analysis of the connectionist model representations were left to models with *significantly* better performance results than their corresponding benchmarks.

7.3. Conclusions

The connectionist model simulations of the bankruptcy classifications showed results close to what could be expected. The traditional forms of the learning and cross validated error curves were replicated[1]. Similarly, the expected dilemma of sufficient complexity and overfit were illustrated both by the learning overfit and the hidden unit overfit findings.

The recommended measure of model fit; cross validated average squared error, showed lower values for the connectionist models with hidden units than for similar models without hidden units and for all benchmarks, even though the difference was not significant. Lack of significance could be explained in several ways. First, the distribution of the errors of the stepwise logistic regression model[2] was very different from the errors of the connectionist model. This was caused by a set of relatively few large error values, giving a very large standard deviation for the benchmark model, and consequently, a very large standard error of the t-test. The distribution of errors of the connectionist model was very different, and showed response values around the cut-off value for all cases with large errors. This difference must be considered an advantage of the connectionist models, even though it weakens the significance of tests of differences[3]. Second, there was a mismatch between the values of the response variable, the error measure used in both our tests, and the error measure used by

[1] Learning error is the error of the model during learning. Generally this error decreased as an inverse exponential function of learning time, but since generalisation is focused in this dissertation, no separate reports were made of learning error. However, an example of the typical relationship between learning and cross validated error is shown in appendix K.

[2] The best benchmark.

[3] Due to a large pooled standard deviation in tests of the differences.

backpropagation. However, this difference will be present in all posterior probability estimator based connectionist models[1]. A suggestion is that better correspondence between the response variable measure and the error measure of the models should be created. Third, the bankruptcy classification operationalisation used in this simulation was derived from a measure containing more information on the subjects' judgement of the financial situation than a simple classification could reveal. A suggestion is that the connectionist models would benefit from taking this information into consideration. These suggestions are further explored in chapters 8 and 9. To summarise the performance results, no significant support for the propositions P1 and P2 was found.

The analysis of the connectionist models performed in this chapter was very limited. We only introduced some of the most widely applied methods for analysing connectionist model representations, and left the deeper analysis to connectionist models proving significantly better performance than their benchmarks. Despite the limitations, sensitivity analysis and visual inspection of the Hinton diagrams indicated that the hidden units of the connectionist models developed complex variable representations[2]. These representations proved useful in diagnostic performance, but evaluation of their cognitive relevance seemed difficult. Thus, no support for proposition P3 made in chapter 4 was found.

However, the results of the bankruptcy classification simulations were promising, and modifications of the inputs, topologies and outputs to better control complexity and overfit were encouraged. In particular, modifications better utilising the error correction algorithm of backpropagation and the available measures of diagnostic response should be made. The set-up and results of simulations taking these modifications into consideration are reported in chapters 8 and 9.

---

[1] As long as the connectionist models minimise squared errors.

[2] For an explanation of different representational types, see sections 3.2.3 and 4.2.

Chapter 8. A connectionist model of continuous response

The first and simplest modification made to improve the results of chapter 7, was using a measure of the judgement of the financial situation of the firm that contained more information than the simple bankruptcy classification measure. The measures of level and trend developed in chapter 5 were presumed to satisfy this demand.

To estimate the generalisation properties of a connectionist model with continuous response, a series of simulations was set up. A model with continuous response deviates from the bankruptcy classification model of chapter 7. While the bankruptcy classification model could be interpreted as a posterior probability estimator, the continuous response models of this section have more in common with expectancy estimators and prediction models. However, this change in family resemblance is a result of the transformation of a classificatory variable by composite judge averaging, and not the result of a change in cognitive task. Thus, the models must still be interpreted as cognitive classification models.

This chapter reports the performance results of the continuous response connectionist model in section 8.1. An analysis of the model representations is reported in section 8.2, and the main conclusions of the simulations are summarised in section 8.3.

8.1 Performance results

As mentioned in chapters 3 and 6, a number of parameters influence the learning and generalisation properties of a connectionist model. Hanson and Burr (1990) suggested that unit complexity (output function), architecture (hidden units) and learning rule affect the representation. Traditionally, weight initialisation, number of hidden units, learning cycles, and learning parameters are investigated for their effect on learning and generalisation (e.g. Nesvik, 1993). In most studies, the effect on learning error is in focus. Here, we focus mainly on the generalisation ability of the connectionist models. Thus, cross validation simulations were set up following the principles explained in chapter 6. The learning rates, momentum terms and intervals of initialised weights were set as described in chapter 6. As in chapter 7, the maximum number of learning iterations was set to 30000. Initial tests indicated that 30000 was a number of iterations somewhat larger than the point of optimal fit.

By monitoring the generalisation error for an increasing number of learning iterations and hidden units, we had the opportunity to use both the optimal hidden unit rule (Le Cun, 1990) and the optimal stopping rule (Smith, 1993) to find the best model. As a consequence of the cross validation procedure, a total number of 1800 backpropagation models were simulated in this chapter. In all simulations, the full set of 32 input variables were used. Three configurations of the output units were used. The first configuration used one unit representing the level variable. The second configuration used one unit representing the trend variable, and the third configuration used two output units representing the level and trend variables.

The results of the level diagnosis models are shown in table 8.1.

Iterations:	5000	10000	15000	20000	25000	30000	Regr.[1]
Model:							
Regr. (all)							0.435
Regr. (stepw.)							0.275
HID0	0.224	0.242	0.255	0.266	0.273	0.281	
HID2	0.206	0.175	0.185	0.195	0.201	0.208	
HID4	0.194	0.183	0.191	0.199	0.210	0.214	
HID6	0.181	0.179	0.189	0.201	0.210	0.222	
HID8	0.183	0.184	0.197	0.205	0.214	0.224	
HID10	0.183	0.188	0.197	0.211	0.221	0.220	
HID12	0.185	0.193	0.200	0.210	0.219	0.225	
HID14	0.186	0.194	0.205	0.215	0.228	0.235	

Table 8.1. Mean squared error (MSE) of the level diagnosis (N=75)

Table 8.1 shows how the number of hidden units affected the generalisation properties of the models. For the maximum number of epochs, we found the best generalisation measure in the connectionist models with two hidden units. This model also had the lowest MSE of all models. This suggested that two hidden units should be used in the model. In addition, the generalisation errors seriously increased when the hidden layer was removed, making the performance of the connectionist model close to the best regression model.

In addition to the two benchmark models of chapter 6, we estimated two comparable regression models. First, a complete regression model with the same number of independent variables as the connectionist models was tested. Next, a model developed by stepwise regression, using the standard stepwise procedure of SPSS (SPSS Inc., 1990), was used. The cross validation procedure was followed for the regression models in the same way as for the connectionist models. The stepwise procedure model performed significantly better than the full model. Furthermore, the best connectionist model significantly outperformed the multiple regression analysis. This is illustrated in table 8.2, showing the t-values of pairwise comparisons of the best regression model and the best connectionist model.

[1] Performance results of the regression models are placed in a separate column. These models were estimated using traditional OLS methods.

Iterations:	5000	10000	15000	20000	25000	30000
HID2	1.53	2.57*	2.45*	2.11*	1.91	1.60

Table 8.2. T-tests of best connectionist model vs. best stepwise regression model for the level diagnosis at increasing number of iterations (* indicates significant at α=0.05, d.f.=74)

A similar test of the difference in performance between the best connectionist model and the best benchmark of chapter 6, also showed a significant difference in favour of the connectionist models. This is illustrated in table 8.3, showing a t-test of the difference in MSE between the best benchmark (A) of chapter 6 and the connectionist model of the level diagnosis.

Iterations:	5000	10000	15000	20000	25000	30000
HID2	0.93	2.43*	2.03*	1.40	1.09	0.80

Table 8.3. T-tests of best connectionist model vs. best benchmark model (A) of chapter 6 for the level diagnosis at increasing number of iterations (* indicates significant at α=0.05, d.f.=74)

In addition to performing worse than the connectionist model, the standard deviation and the MSE of the regression models were correlated. This indicated that the variance of the model errors increased with increasing mean error. The same tendency was found in the maximum error of the models, which was larger for the regression analysis and for the connectionist models without hidden units than for the connectionist models with hidden units. Consequently, strong support was found for proposition P1 of chapter 4 for the level diagnosis.

To investigate proposition P2, the difference between the performance of the connectionist models with and without hidden units was compared. Table 8.4. illustrates this in a t-test of the differences between means.

Iterations:	5000	10000	15000	20000	25000	30000
HID2	0.73	2.86**	2.80**	2.60*	2.49*	2.55*

Table 8.4. T-tests of connectionist model with 2 hidden units vs. connectionist model without hidden units for the level diagnosis at a comparable number of iterations (* and ** indicates significant at α=0.05 and 0.01 respectively, d.f.=74)

A test of the difference in performance between the two best connectionist models with and without hidden units showed that the connectionist model with two hidden units significantly outperformed the model without hidden units at α=0.01 (t=2.69, d.f.=74). This finding strongly supported proposition P2 for the level diagnosis.

To further investigate the distribution of the errors for the different models, correlations between error terms and targets were computed. The correlations of the cross validated squared errors (SE) with the target value of the composite judge level diagnoses are shown in table 8.5.

Iterations:	5000	10000	15000	20000	25000	30000	Regr.
Model:							
Regr. (all)							0.179
Regr. (stepw.)							0.147
HID0	0.085	0.141	0.162	0.173	0.183	0.181	
HID2	-0.107	-0.166	-0.120	0.003	0.012	0.041	
HID4	0.037	0.005	0.028	0.040	0.052	0.040	
HID6	0.028	0.012	0.017	0.025	0.021	0.040	
HID8	0.003	0.009	0.028	0.033	0.046	0.068	
HID10	0.035	0.026	0.046	0.061	0.071	0.072	
HID12	0.058	0.035	0.046	0.063	0.079	0.064	
HID14	0.042	0.043	0.063	0.065	0.070	0.069	

Table 8.5. Correlations of SE and target for the level diagnosis (N=75)

None of the correlation coefficients in table 8.5 are significantly different from 0 (α=0.05), but typically they are larger for the regression and connectionist model without hidden units than for the other models. The high correlation of MSE and target suggested errors were larger for larger targets. The connectionist models did not seem to make this type of error. This can be explained by an equal distribution of the errors along the target value, or by a distribution with larger errors for targets distant from the mean target. To test the last suggestion, correlations of the cross validated squared errors with the differences from target means were calculated. The results are shown in table 8.6.

Iterations:	5000	10000	15000	20000	25000	30000	Regr.
Model:							
Regr. (all)							0.010
Regr. (stepw.)							0.095
HID0	0.199	0.132	0.087	0.065	0.041	0.031	
HID2	0.253*	0.082	-0.002	-0.044	-0.099	-0.100	
HID4	0.225	0.098	0.033	-0.005	-0.034	-0.078	
HID6	0.214	0.107	0.059	0.022	-0.001	-0.030	
HID8	0.198	0.084	0.049	0.020	-0.022	-0.027	
HID10	0.199	0.117	0.057	0.007	-0.016	0.001	
HID12	0.195	0.133	0.097	0.057	0.017	-0.032	
HID14	0.175	0.132	0.111	0.048	0.033	-0.013	

Table 8.6. Correlations of SE and distance from mean target for the level diagnosis (N=75)(* indicates significant at α=0.05)

Table 8.6 shows very small and insignificant correlations for the "saturated" models. This suggests the errors were somewhat correlated with the target value for the simple models, but not for the connectionist models with hidden units. However, early in the learning process, the correlations with the distance from targets were high. This can primarily be explained by the probability density function of the targets. With the majority of targets close to the mean, these diagnoses will be learned first, and the model will specialise on learning the distant targets later.

To further investigate the errors of the connectionist model, we computed the correlations of the standard deviation of the composite judge diagnosis, indicating inter-judgmental disagreement on the diagnosis, with the squared errors of the models. If there was disagreement among the diagnosticians on a stimulus, one might assume the errors of the models should also be large on these stimulus-response pairs. This was not the case. All correlations were low and not significantly different from 0 (α=0.05).

One might further ask if the different models failed on the same stimulus-response pairs. The correlations of the SEs could be used to answer this question. The correlation matrix of all model errors was calculated. They indicated that all correlations were significantly different from 0 (α=0.01) and about 0.9 in magnitude. A small difference was detected between the connectionist models, with the lowest mean error correlation being 0.89, and the regression analysis, with the lowest correlation being 0.76. This meant that the connectionist models all failed on the same stimulus-response pairs in similar patterns. The regression models also failed on the same pairs, but the structure of the error was somewhat, but not significantly, different.

The procedures and analyses of the level diagnosis described above were also set up for the trend diagnosis. At least from a cognitive perspective, some trend information can be produced simply by comparing two or more figures without reference to an internalised standard. Thus, one may assume that the diagnosis of trend is simpler than the diagnosis of level. However, this may not be the case for models that are given situational cues as inputs. These models have to develop some notion of a "trend" concept as resulting from a difference in two or more input cues. To further complicate the concept, the trend and level diagnoses are correlated. There is a higher probability of being in a positive trend in a good situation than in a bad one. These aspects imply that the "trend" concept might be more complex than first assumed, and that an intermediate abstraction of parts of the concept may be helpful in diagnosis.

The cross validation procedure was set up in the same way and with the same parameters as in the level diagnosis simulations. The cross validated results for the model of the trend diagnosis are shown in table 8.7.

Iterations:	5000	10000	15000	20000	25000	30000	Regr.
Model:							
Regr. (all)							0.503
Regr. (stepw.)							0.431
HID0	0.344	0.352	0.357	0.363	0.366	0.370	
HID2	0.479	0.347	0.356	0.358	0.358	0.374	
HID4	0.442	0.351	0.357	0.356	0.362	0.351	
HID6	0.390	0.347	0.351	0.344	0.348	0.347	
HID8	0.366	0.355	0.352	0.362	0.362	0.362	
HID10	0.358	0.349	0.353	0.360	0.357	0.350	
HID12	0.361	0.359	0.349	0.355	0.360	0.355	
HID14	0.350	0.354	0.359	0.359	0.359	0.361	

Table 8.7. Mean squared error (MSE) of the trend diagnosis (N=75)

Table 8.7 shows almost the same pattern as for the level diagnosis models. The connectionist models significantly outperform the regression analysis in cross validated squared error. Pairwise t-tests of the performance differences between the best connectionist model and the regression model are shown in table 8.8.

Iterations:	5000	10000	15000	20000	25000	30000
HID6	0.92	2.54*	2.26*	2.25*	2.04*	1.94

Table 8.8. T-tests of best connectionist model vs. stepwise regression for trend diagnosis at increasing number of iterations (* indicates significant at α=0.05, d.f.=74)

The results in table 8.8 show that the connectionist model was significantly better than the best regression model, with the most significant difference in means for 10000 iterations. The same pattern was found in maximum error differences of the connectionist and the regression models.

T-tests of the difference in performance between the best benchmark model (A) of chapter 6 and the best connectionist model showed a small difference in the same direction as found in table 8.8, but the difference was not significant at α=0.05 (t=0.26 , d.f.=74). Consequently, the connectionist models did not significantly outperform the best linear benchmark model (A) on the trend diagnosis.

The larger number of hidden units in the optimal connectionist model of the trend diagnosis than the model of the level diagnosis, could indicate that the trend concept was more complex than the level concept. However, four things are worth mentioning. First, the error surface was very flat and almost independent of the complexity of the model used. Second, the major improvement in model performance occurred when shifting from regression model to connectionist model without hidden units. For a given number of iterations, the connectionist model with 6 hidden units was better than the model without hidden units, but the difference was not significant. The corresponding differences were significant for the level diagnosis models above. Third, the model without hidden units had a low MSE even for a small number of iterations. This suggested that simpler models could be used if no other weaknesses[1] were discovered. Fourth, all MSEs were larger for the trend diagnosis models than for the level diagnosis models. The last finding suggested that "trend" was a more complex[2] concept to model, but the other findings suggested this interpretation should be made with caution.

Correlations with the trend target were computed to study the distribution of the errors over the different stimulus-response pairs. The results are shown in table 8.9.

Iterations:	5000	10000	15000	20000	25000	30000	Regr.
Model:							
Regr. (all)							0.166
Regr. (stepw.)							0.072
HID0	0.107	0.091	0.082	0.071	0.067	0.045	
HID2	0.097	0.079	0.086	0.086	0.075	0.043	
HID4	0.156	0.102	0.073	0.078	0.071	0.043	
HID6	0.156	0.090	0.062	0.060	0.032	0.019	
HID8	0.146	0.092	0.077	0.054	0.036	0.039	
HID10	0.121	0.076	0.056	0.044	0.028	0.033	
HID12	0.132	0.082	0.067	0.039	0.028	0.038	
HID14	0.099	0.069	0.068	0.058	0.021	0.018	

Table 8.9. Correlations of SE and target for the trend diagnosis (N=75)

The correlations showed approximately the same pattern as for the level diagnosis models. As opposed to the level diagnosis models above, the connectionist models without hidden units now showed a correlation comparable to the other connectionist models. Models with small MSEs had small correlations with target. Again, this can be explained by an even distribution of the errors, or by errors correlating with distance from target means. By correlating the SEs with distance from the mean trend targets, we got the results shown in table 8.10.

The results shown in table 8.10 were very different from the level model results illustrated in table 8.6. There were significant correlations of SEs with distance from mean targets for the stepwise regression and for the connectionist models after few learning iterations. This implied that the models missed most on the patterns classified as "good" or "bad". This was somewhat disturbing, since we particularly wanted the model to fit these cases correctly.

---

[1] Such as unfavourable correlations of error with targets or distance from targets
[2] Or noisy.

Iterations:	5000	10000	15000	20000	25000	30000	Regr.
Model:							
Regr. (all)							0.114
Regr. (stepw.)							0.264*
HID0	0.264*	0.208	0.194	0.185	0.172	0.170	
HID2	0.720**	0.240*	0.197	0.185	0.172	0.155	
HID4	0.587**	0.231*	0.197	0.167	0.172	0.143	
HID6	0.467**	0.228*	0.194	0.182	0.172	0.143	
HID8	0.460**	0.205	0.185	0.188	0.162	0.142	
HID10	0.369**	0.227*	0.208	0.191	0.182	0.167	
HID12	0.357**	0.241	0.188	0.183	0.180	0.151	
HID14	0.298**	0.206	0.186	0.167	0.145	0.135	

Table 8.10. Correlations of SE and distance from mean target for the trend diagnosis (N=75) (** and * indicates significant at α=0.01 and 0.05 respectively)

For the level models, a similar effect was eliminated by learning. For the trend diagnosis models, it was not. For the level diagnosis models, one was led to assume that the optimal learning point and the even distribution of errors were found simultaneously by the learning rule. However, the results of the trend diagnosis indicated that the optimal learning point and favourable distribution of errors were not found simultaneously by the learning rule. The lowest correlation was found for the full regression model.

One reason why the model did not eliminate the errors by learning, could possibly be found in the distribution of the trend variable. The standard deviation of the trend variable was not larger than the level variable, but the Kolmogorov-Smirnov test indicated a small deviance from the normal distribution. If this explanation was correct, a lacking robustness of connectionist models previously unattended (Cheng and Titterington, 1994; Ripley, 1993), could have been detected. Another possible explanation was that the models' complexity was insufficient to capture all the properties of the "trend" concept. The third possible explanation was that the models were too complex, and thus, always overfit. This suggested that the number of free parameters in the model was too large to be set by the relatively small samples of 74 patterns each. However, the stepwise regression procedure resulted in a model with considerably less free parameters, but with similar error distribution problems. Thus, a sensitivity of the connectionist models to the distribution of the response variable seemed to be the most reasonable explanation.

One conclusion that could be drawn from these results is that only investigating MSE when evaluating a model's performance may be too limited. The results suggested an evaluation of the generalisation ability of a connectionist model should consider both MSEs and measures of the distribution of errors. The suggestion that models with less free parameters could be used to avoid some of the error distribution problems detected here, is explored in chapter 9.

To further investigate the errors of the connectionist model, we correlated the squared errors with the standard deviation of the trend diagnosis of the composite judge groups. The analysis showed no significant correlations. As for the level diagnosis, this indicated that the model errors were not correlated with disagreement in the composite judge groups. Finally, we calculated the correlation coefficients of all the errors, and all correlations were significantly different from 0 (α=0.01), and about 0.9 in magnitude. A small difference was detected between the connectionist models, with the lowest mean error correlation being 0.96, and the mean error correlation of the regression models, with the lowest correlation being 0.78. This implied that the connectionist models failed in similar patterns. The structure of the errors was somewhat different for regression analysis, but the difference in the model error correlations was not significant.

Other studies (e.g. Bounds, Lloyd and Mathew, 1990; Chakaborty et al., 1992), have documented improved test results by modelling more than one output variable simultaneously, when the output variables are correlated. To test the performance of such a *combined* model, a cross validation procedure similar to the previous simulations was set up. The simulations were run with the same parameter values. The only change in models was the introduction of an additional output unit. The performance results for the two diagnostic variables are shown in tables 8.11 and 8.12.

Iterations:	5000	10000	15000	20000	25000	30000
Model:						
HID0	0.239	0.249	0.264	0.273	0.278	0.283
HID2	0.194	0.180	0.179	0.181	0.189	0.194
HID4	0.197	0.178	0.181	0.188	0.195	0.205
HID6	0.189	0.187	0.201	0.210	0.209	0.218
HID8	0.197	0.188	0.193	0.198	0.208	0.212
HID10	0.187	0.187	0.193	0.197	0.201	0.205
HID12	0.191	0.193	0.199	0.199	0.211	0.218
HID14	0.192	0.197	0.200	0.205	0.213	0.215

Table 8.11. Mean squared error (MSE) of level diagnosis in a combined model (N=75)

Iterations:	5000	10000	15000	20000	25000	30000
Model:						
HID0	0.350	0.358	0.363	0.368	0.368	0.365
HID2	0.428	0.347	0.349	0.350	0.357	0.365
HID4	0.399	0.336	0.347	0.353	0.360	0.367
HID6	0.396	0.339	0.352	0.360	0.369	0.379
HID8	0.361	0.342	0.344	0.352	0.352	0.360
HID10	0.359	0.328	0.339	0.348	0.354	0.358
HID12	0.341	0.337	0.347	0.355	0.367	0.377
HID14	0.350	0.344	0.353	0.351	0.379	0.373

Table 8.12. Mean squared error (MSE) of trend diagnosis in a combined model (N=75)

The results of the combined model were very similar to the results of the separate models. A small improvement could be noted for the combined models, particularly for small numbers of hidden units. Compared to the separate models, a combined model with two hidden units stopped at 15000-20000 iterations performed very well. However, no *significant improvement* in model fit was found using a combined model[1]. However, for analysis purposes, the advantages of a simple, combined model should not be underestimated.

To test if a combined model had the same error distribution problems as the separate models, similar correlations as reported above were computed for the combined models. The same pattern as for the separate models was found for all correlations. However, the correlations of the SEs with differences from mean targets were somewhat smaller for the trend diagnosis variable in the combined model. This is illustrated in table 8.13.

Iterations:	5000	10000	15000	20000	25000	30000
Model:						
HID0	0.274*	0.227*	0.213	0.192	0.182	0.171
HID2	0.468**	0.227	0.169	0.149	0.142	0.140
HID4	0.386**	0.229*	0.185	0.155	0.120	0.105
HID6	0.415**	0.237*	0.210	0.186	0.165	0.167
HID8	0.349**	0.225	0.212	0.177	0.178	0.163
HID10	0.361**	0.234*	0.214	0.174	0.168	0.163
HID12	0.334**	0.239*	0.213	0.177	0.178	0.156
HID14	0.290*	0.221	0.194	0.187	0.166	0.142

Table 8.13. Correlations of SE and distance from mean target for the trend diagnosis in a combined model (N=75) (** and * indicates significant at α=0.01 and 0.05 respectively)

The correlations of the trend diagnosis errors with distance from mean target of the simple connectionist model with two hidden units was comparable to the best separate trend model with six hidden units. Consequently, there seemed to be no advantages in modelling level and trend diagnoses separately. Three important lessons were learned from these simulations. First, significantly better model fit was found for the level diagnoses with connectionist models. Second, the improved fit of the connectionist models was restricted to models with hidden units. Third, the connectionist models' fit should be evaluated by investigating more than a simple measures of cross validated squared error.

The promising results for the connectionist models with hidden units made analysis of the internal representations of these models particularly relevant.

[1] The t-value of a test of difference in cross validated average squared error of the trend diagnosis between the best combined connectionist model and the best benchmark (A) of chapter 6 is somewhat higher than for the separate model. However, it is still not significant at α=0.05 (t=0.84, d.f.=74).

8.2 Analysis of model representations

When studying the representations of the connectionist models, we wanted to utilise the full learning sample. A *combined* connectionist model with 32 input units and two output units was trained with the full learning sample. To get a picture of the variations in the representation, the number of hidden units was varied from two to four units. Overfit was controlled by stopping learning at a point lower than the optimal fit point found and reported in section 8.1. This procedure was used to control *early* overfit resulting from including all stimulus-response pairs in the learning sample. Even though the stopping point was as low as 10000 iterations, the mean squared errors of these models were lower than the cross validated average squared error reported above. The effect of different initial weights was controlled by developing 10 different models for each number of hidden units. Each model was developed with different randomised initial starting weights. The performance results of these models are shown in table 8.14.

Model	Average level MSE	St dev. of level MSE	Average trend MSE	St. dev. of trend MSE	Correlation of level with diff. from target	Correlation of trend with diff. from target	Runs with common hidden unit
HID2	0.110	0.012	0.203	0.021	0.028	0.183	10
HID3	0.101	0.003	0.187	0.005	0.026	0.164	10
HID4	0.101	0.004	0.183	0.004	0.023	0.153	10

Table 8.14. Results of 10 combined models with full learning sample and randomised initial weights (figures are averages of the 10 runs)

From table 8.14, we see that the MSE was generally lower than in the cross validation simulations. This was expected since all cases were included in the learning sample. The low standard deviation of the 10 simulations with different initial weights indicated that the performance of the models were relatively independent of the initial weights. The performance results of the connectionist models with two hidden units were somewhat disturbed by an MSE of 0.137 and 0.253 for the first simulation (version A2) on the level and trend diagnoses respectively. The previously found pattern of small correlations with difference from the target for the level diagnosis and large correlations for the trend diagnosis, was also found here. Consequently, the properties of the errors in these models were similar to the cross validated errors analysed in section 8.1.

Even though performance results were very similar in the models in table 8.14, the final weights were very different. Summary statistics illustrating the weight distributions of each model are shown in table 8.15. The reasons for the differences in final weights could be many local optima, or a flat error surface around the optimal solution. Since error continued to decrease as learning continued, the first explanation was excluded. It was more likely that the model error surfaces with respect to the weights were very flat around the saturation area.

Theoretically, the error should approach zero along one, or a few, of the weight axes if the process had not been terminated. Therefore, depending on the initial weights, each representation studied was one among several different representations with similar performance results. For each model, their current representation was one of many performing the same mapping, but each of the representations was different when it came to *how* the mapping was performed.

For a combined connectionist model, we were interested in whether or not the hidden units specialised on the level and trend diagnoses. All models with two, three and four hidden units had hidden units with large similar connections to *both* output units. Consequently, one of the hidden units was used to represent common aspects of the two diagnoses. What tasks were performed by the rest of the hidden units was, however, a somewhat more difficult question to answer and this will be treated in section 8.2.2

To study the representations of a connectionist model and the tasks performed by the hidden units, a number of methods can be applied (see chapter 3). Here, analysis of the weight distributions and outputs of the hidden units, Hinton diagrams (Hinton, 1989), and cluster analysis (Gorman & Sejnowski, 1989; Hanson & Burr, 1990) were applied. These methods are well known, but applying them to our models[1] demanded careful modification and adjustment.

To illustrate the hidden unit weights of the models, a Hinton diagram of the 30 connectionist models with two, three or four hidden units is shown in figure 8.1. Only the weights between hidden units and output units are shown. As in figure 7.1 and 7.2, the size of the squares represents the absolute value of the weight. Dark and light squares represent positive and negative values of the weights.

Figure 8.1 illustrates the differences between the connectionist model weights resulting from randomisation of initial weights. The Hinton diagram indicated presence of a common hidden unit in all the models. Most models had one positive common hidden unit, but models with one negative and more than one common hidden units were also found.

---

[1] To study the cognitive relevance of our model representations.

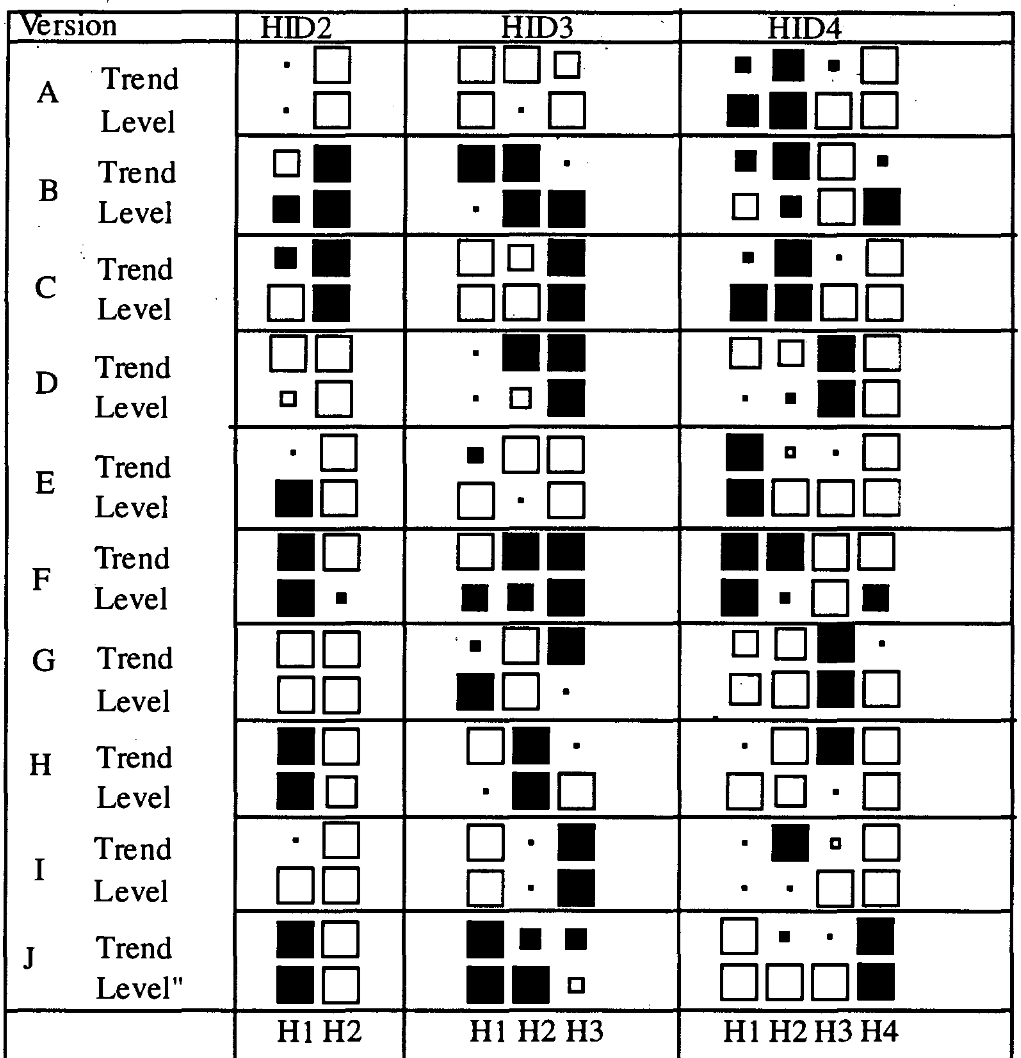

Figure 8.1. Hinton diagram of the weights between hidden layer and output layer of 10 versions of the combined models with 2, 3 and 4 hidden units

Comparing the different versions of each model, no recurring weight pattern was found. Of the models with two hidden units, versions H and J were somewhat similar, and of the models with four hidden units, versions A and C were somewhat similar. Except for these examples, the representations did not seem to have a small number of local solutions (attractors).

To further investigate the differences in representations of the models, summary statistics on the distributions of the connection weights were calculated. Some statistics illustrating these distributions are shown in table 8.15.

Iterations:	HID2 (N=72)		HID3 (N=107)		HID4 (N=142)	
Version::						
A	Mean:-0.17 Skew:-1.22	Stdev:0.72 Kurt:3.61	Mean:-0.13 Skew:-0.54	Stdev:0.68 Kurt:0.86	Mean:0.03 Skew:0.88	Stdev:0.56 Kurt:4.01
B	Mean:0.10 Skew:0.84	Stdev:0.80 Kurt:2.38	Mean:0.08 Skew:0.41	Stdev:0.71 Kurt:1.95	Mean:-0.01 Skew:-0.49	Stdev:0.57 Kurt:2.92
C	Mean:0.01 Skew:1.10	Stdev:0.79 Kurt:2.37	Mean:-0.03 Skew:-0.80	Stdev:0.65 Kurt:2.22	Mean:-0.01 Skew:-0.62	Stdev:0.56 Kurt:2.24
D	Mean:-0.12 Skew:-0.83	Stdev:0.81 Kurt:1.29	Mean:0.03 Skew:1.19	Stdev:0.70 Kurt:4.28	Mean:0.01 Skew:-0.84	Stdev:0.56 Kurt:3.07
E	Mean:-0.02 Skew:-1.36	Stdev:0.75 Kurt:3.17	Mean:-0.12 Skew:-0.98	Stdev:0.67 Kurt:2.84	Mean:-0.12 Skew:-0.63	Stdev:0.56 Kurt:3.70
F	Mean:0.10 Skew:-0.84	Stdev:0.81 Kurt:2.48	Mean:0.07 Skew:0.58	Stdev:0.68 Kurt:2.35	Mean:0.01 Skew:-1.04	Stdev:0.57 Kurt:3.36
G	Mean:-0.14 Skew:-0.58	Stdev:0.82 Kurt:0.64	Mean:0.01 Skew:-0.69	Stdev:0.63 Kurt:2.74	Mean:-0.12 Skew:0.89	Stdev:0.56 Kurt:4.23
H	Mean:0.07 Skew:0.43	Stdev:0.76 Kurt:1.78	Mean:0.01 Skew:1.03	Stdev:0.66 Kurt:3.61	Mean:-0.11 Skew:-0.63	Stdev:0.58 Kurt:1.62
I	Mean:-0.16 Skew:-0.99	Stdev:0.75 Kurt:2.00	Mean:-0.08 Skew:-0.05	Stdev:0.65 Kurt:3.95	Mean:-0.12 Skew:-1.09	Stdev:0.56 Kurt:4.07
J	Mean:-0.09 Skew:-0.03	Stdev:0.78 Kurt:1.66	Mean:0.06 Skew:0.48	Stdev:0.70 Kurt:1.73	Mean:-0.05 Skew:-0.94	Stdev:0.56 Kurt:4.14

Table 8.15. Statistics illustrating the differences in representations between the 10 versions of each combined model[1]

In table 8.15, the mean, standard deviation, skewness and kurtosis of the weights of each model are shown. All model connections including bias weights were included in these statistics. Due to the larger number of weights, the standard deviation of the weights was reduced using more hidden units in the model. For the model with two hidden units, the standard deviation ranged from 0.72 to 0.82, the skewness from -1.36 to 1.10, and the kurtosis from 0.64 to 3.61. The small variance of the standard deviation indicated that even though the weight pattern was different from model to model, the variance of the weights was about the same. Models with negative mean weight values were left skewed, and models with positive mean weight values were right skewed. Somewhat surprising was the variance in the kurtosis. All models had a somewhat more peaked distribution of weights than the normal distribution with a similar mean and standard deviation. The models with more hidden units showed a smaller standard deviation, a smaller range of the skewness (-0.69 to 1.19 and -1.04 to 0.89) and generally, a larger kurtosis (0.86 to 4.28 and 1.62 to 4.23). This implied that the distributions were gradually more peaked as the number of weights increased. For the small connectionist models, the necessary large weights caused only a small deviance from the normal distribution. As the models "grew" larger, these weights were not eliminated, but their significance in "normalising" the distribution of the weights was reduced. A preliminary conclusion was that the representations was generally distributed, but that some weights had a highly local representation[2].

---

[1] Mean, standard deviation, skewness and kurtosis. Skewness and kurtosis are standardised.

[2] The differences between local and distributed representations are explained in section 3.2.3.

Cluster analysis can be used to study how distributed a representation is (Hanson & Burr, 1990). Cluster analysis of the connections between the input and hidden units are usually performed to detect local units when the number of hidden units is large. By applying the same procedure to all hidden units in the 10 versions of the models with two, three and four hidden units, we could detect units with similar weights. Absolute values of the model weights were used in the analysis. A dendrogram illustrating the clusters is shown in figure 8.2[1].

In the dendrogram of figure 8.2, the hidden units are identified with a label consisting of the model version (A-J), a figure indicating the number of the hidden unit (1-4), and a figure indicating the total number of hidden units in the model (2-4). Three observations were made from the dendrogram. First, the dendrogram separated the weights in two large clusters with weights very different from each other. Second, the hidden units located in the top cluster of the dendrogram were the hidden units termed "common" in the analysis of figure 8.1. Third, these units were located at the top of the dendrogram because they had weight patterns with large variances. The hidden units in this cluster were the units with the largest absolute weight values and consequently, they could be characterised as forming a *local* representation. The same pattern was found in separate cluster analyses of the 10 versions with two, three and four hidden units, respectively. Of the 34 hidden units in the upper cluster, 30 units were hidden units previously classified as representing common aspects of the level and trend diagnoses. The last four also represented common aspects, but these units were all the second common hidden unit in models containing more than one such unit.

Two important findings had been made. First, all connectionist models developed a common hidden unit with a local representation detecting the common parts of the correlated diagnostic variables level and trend. In most models, this task was performed by one hidden unit. Depending on the bias weights and the weights between the hidden and output layer, the input weights to these hidden units had one of two different patterns[2]. Second, the rest of the hidden units had a much more distributed representation, and in most of the cases, these units specialised on other specific features in the input material. As a consequence, two analyses were necessary. The common hidden units were analysed first, and the results are reported in section 8.2.1. Next, the separate hidden units were analysed. This analysis is reported in section 8.2.2.

---

[1] All cluster analyses in this study were performed with the default settings of the "CLUSTER" procedure of SPSS using squared Euclidian distances and average linkages.

[2] The weights were different for the hidden units with two large negative weights and two large positive weights, but they had almost similar absolute values.

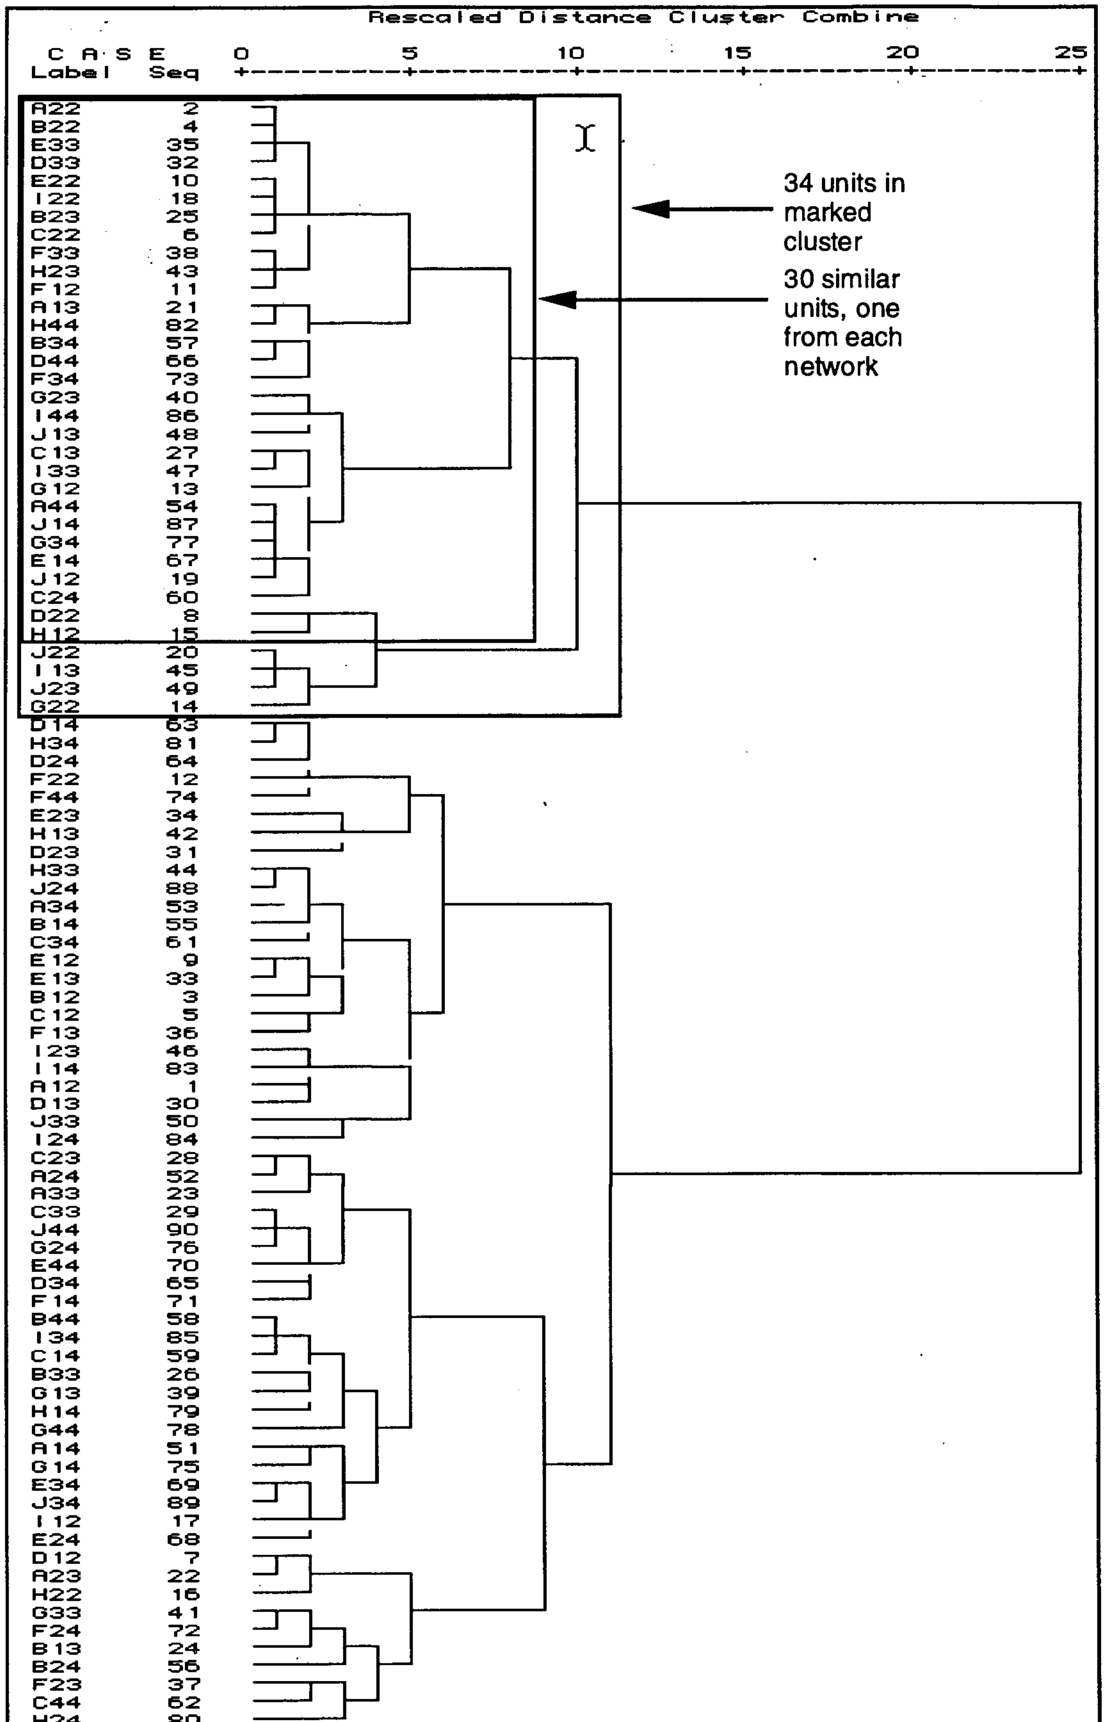

Figure 8.2. Cluster analysis dendrogram of input to hidden weights in all the 30 versions of the combined connectionist model with 2, 3 and 4 hidden units

8.2.1 Common hidden unit analysis

The feature detecting properties[1] of the common hidden units can be analysed by looking at their weight pattern. The weight patterns of three common hidden units can be studied in figure 8.5. Notice the sign differences in the connectionist models with two and three hidden units, from that of four hidden units. To generalise this analysis principle, all the common hidden units in the 30 models were investigated.

In our models, we found both excitatory and inhibitory common hidden units. An inhibitory common hidden unit is activated by turning the hidden unit on, and letting this unit inhibit the output units already turned on by a positive bias weight. Thus, the sign of the weights coming into an inhibitory common hidden unit can be turned to make it an excitatory common hidden unit. In our analysis, inhibitory common hidden unit signs were turned and the weights were averaged over the 30 models. The excitatory common hidden units were: B22, B23, C22, C24, D33, E14, F12, F33, G34, H12, H23, I33, J12, and J13. The inhibitory units were: A13, A22, A44, B34, C13, D22, D44, E22, E33, F34, G12, G23, H44, I22, I44 and J14. Some statistics on how these units were connected to the input units are shown in table 8.16.

The small standard deviations in table 8.16 illustrate the similarity of the weight pattern of the common hidden units. Variation in the weights among the 30 hidden units was not large enough to prevent the t-values of table 8.16 from being very high. Of most interest were the weights where maximum and minimum values were on the same side of the origo. This was the case for units representing SGROWTH, CGROWTH, CONTPR1, PROMARG1 and 2, OPMARG1 and 2, ROI2, ROE1, AIR2, ICOV1, LTINV2, ART1, APT1 and 2, CURR2, ACID2 and BER2. The value of the weight was an indication of the importance[2] of an input unit in turning the common hidden unit on. The input units represented indicators of "profitability", "financing", "liquidity" and "leverage". The *sign* of the weight values indicated positive or negative influence on the common hidden unit. The *values* of the weights were comparable across indicators because the input values had been transformed linearly to the [0,1] scale. However, different bias weights to the hidden units prevented a direct interpretation of the values as measures of importance. It is also important to remember that the basic nonlinear relationship between input value and hidden unit output in connectionist models, prevents simple importance interpretations of the values of weights. Despite these limitations, a first interpretation of the common hidden units was that they

---

[1] For an explanation of the term "feature" we refer to chapter 3. The term "feature" is traditionally used of discrete aspects of the stimulus, while the term "stimulus dimension" is used of continuous aspects of the stimulus. Until it has been determined what aspect of the stimulus the hidden units use, we apply the term "feature".

[2] See chapter 7.

detected aspects relevant to all four financial diagnostic areas[1] with a selected, but broad set of indicators.

Input unit	Mean weight	St. dev. of weight	d.f.	t-value
SGROWTH	0.453	0.098	29	25.42**
CGROWTH	-0.358	0.087	29	-22.69**
CONTPR1	-0.332	0.120	29	-15.12**
CONTPR2	-0.005	0.091	29	-0.32
PROMARG1	-0.629	0.205	29	-16.84**
PROMARG2	1.022	0.177	29	31.67**
OPMARG1	-0.937	0.237	29	-21.67**
OPMARG2	1.581	0.225	29	38.57**
ASSTURN1	-0.046	0.114	29	-2.22*
ASSTURN2	0.175	0.324	29	2.96**
ROI1	-0.448	0.268	29	-9.16**
ROI2	2.034	0.317	29	35.10**
ROE1	-0.893	0.104	29	-47.01**
ROE2	0.137	0.124	29	6.03**
AIR1	0.032	0.267	29	0.65
AIR2	-0.960	0.167	29	-31.50**
ICOV1	-0.435	0.189	29	-12.58**
ICOV2	0.124	0.170	29	4.00**
LTINV1	0.158	0.169	29	5.13**
LTINV2	0.413	0.118	29	19.15**
ITURN1	-0.059	0.116	29	-2.79**
ITURN2	0.138	0.079	29	9.61**
ART1	0.613	0.148	29	22.67**
ART2	-0.027	0.137	29	-1.08
APT1	0.299	0.078	29	20.93**
APT2	-0.379	0.089	29	-23.27**
CURR1	-0.261	0.175	29	-8.16**
CURR2	0.626	0.141	29	24.36**
ACID1	-0.393	0.222	29	-9.70**
ACID2	0.566	0.200	29	15.48**
BER1	-0.238	0.232	29	-5.62**
BER2	1.385	0.301	29	25.20**

Table 8.16. Mean weight values, standard deviation and t-value of test of $\mu=0$ for the weights between input units and the *common* hidden units (** and * indicates significant at $\alpha=0.01$ and 0.05 respectively)

The next thing to remark about the weight pattern of the common hidden units, was the signs of the weights. We found a pattern of one positive and one negative mean weight for the units representing each consecutive year of the following cues: PROMARG, OPMARG, ASSTURN, ROI, ROE, AIR, ICOV, ITURN, ART, APT, CURR, ACID and BER. All these patterns except the pattern of ART were easily interpretable. First, the negative and positive values were as expected. Favourable indicators had a different weight pattern from

[1] "Profitability", "financing", "liquidity" and "leverage".

unfavourable indicators[1]. Second, the pattern indicated that the common hidden units were turned on by one of the units representing each financial cue value and off by the other. Since these units represented cue values from different years of the financial statement, the resulting effect was that the hidden unit was more excited by a large *change* in the value of a financial cue from one year to the next than a small change. The difference in weight values indicated the effect of the *change*, and the absolute value of the weight indicated the importance of the *level* of the financial cue. Third, the typical pattern was that the mean value of the weight from the most recent value of the financial cue was the largest, and thus, contributed most to the activation of the hidden unit.

In conclusion, we found that the common hidden unit was a complex "feature detector". However, the "feature" detected was not discrete. Rather, the units seemed to give a response varying continuously with the value of the detected stimulus dimension. Thus, the common hidden units detected a *new stimulus dimension* in the stimulus material by a complex transformation of the original stimulus dimension values. The weight pattern of the common hidden units suggested that the new stimulus dimension represented a rather complex and merged concept. *Complex,* because it was formed by using indicators of many financial diagnostic areas, and *merged*, because it involved an evaluation of both level and change aspects of the financial cues. The complexity of this concept made configural processing[2] necessary. This was not unexpected, since the diagnostic variables the connectionist models were set up to fit, were highly correlated[3]. It seemed that the common hidden units formed a representation of a "general condition variable".

8.2.2 Analysis of non-common hidden units

A problem existed in analysing the representation of the connectionist models further, because the representations formed by the resulting hidden units seemed relative to the whole connectionist model's weights. One consequence was that further analysis had to be performed on a selected model or a group of selected models. We had groups of connectionist models with two, three and four hidden units. The first step in the analysis was to investigate the three groups to see if the models within each group could be further divided into functionally different groups. A cluster analysis of the unscaled *output* of the connectionist models with two hidden units was performed. The results are shown in figure 8.3.

---

[1] Compare, for example, BER to AIR.
[2] See section 2.4.
[3] Similar findings had been done by Chakaborty et al. (1992).

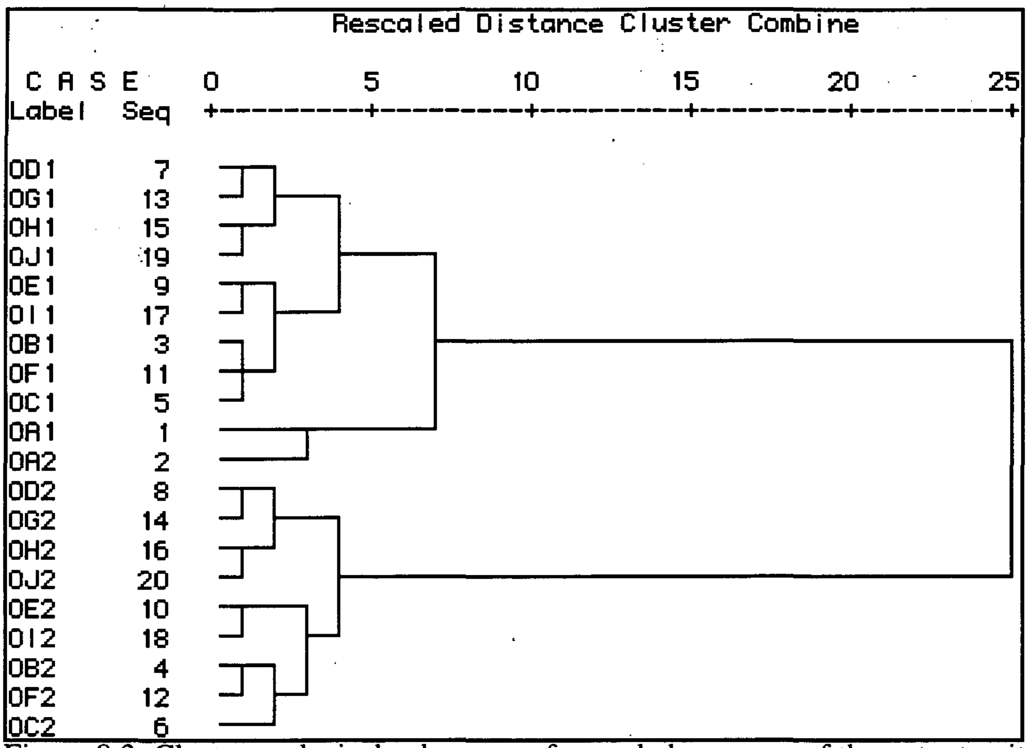

Figure 8.3. Cluster analysis dendrogram of unscaled response of the output units in all combined models with two hidden units[1].

The dendrogram shows the response of the 20 output units in the 10 connectionist models with two hidden units. We found a split in two main clusters. As expected, the main clusters consisted of the level diagnosis output units in the upper cluster, and the trend diagnosis units in the bottom cluster. Model A distorted the picture somewhat, but this was the model with high MSE referred to in section 8.2 as creating the large standard deviation of the MSEs among the models with two hidden units. We chose to ignore this model. Within the main clusters, the other response units were organised in two sub-clusters; similar for both level and trend diagnoses. We interpreted this result as an indication that the connectionist models were organised in two functionally different ways. The models E, I, B, F and C performed the tasks in a way functionally different from the models D, G, H and J. Table 8.17. shows the correlations of the hidden unit outputs with the targets in these models.

[1] The labels indicate O for output, a character showing which network the unit belongs to and a number showing the output unit number (1 - level, 2- trend).

Table 8.17 shows the common hidden units (marked C) and the output pattern of the other hidden units. The common hidden units of the models in the first cluster (DGHJ) were not as jointly correlated with both level and trend as the common hidden units in the models of the second cluster (EIBFC). In addition, we found high correlations of the separate hidden units in all the models in the EIBFC cluster with the difference between the level and trend diagnosis values. A study of the outputs of the models in this sub-cluster indicated that the common hidden unit detected the common aspects of level and trend, and that the separate hidden unit responded to the *"difference"* between the two diagnoses. These models seemed to implement a heuristic with a common hidden unit assuming level and trend were almost perfectly correlated, and a hidden unit specialising on detecting the stimulus patterns that were exceptions to this rule. Model F was most representative of implementing this heuristic.

Unit	Level	Trend	Difference	
D1	-0.3016**	-0.6700**	0.5917**	
D2	-0.9091**	-0.6806**	-0.2353*	C
G1	-0.5421**	-0.8001**	0.4541**	
G2	-0.8782**	-0.5691**	-0.3605**	C
H1	0.9050**	0.6662**	0.2513*	C
H2	-0.3253**	-0.6768**	0.5691**	
J1	0.5424**	0.7993**	-0.4525**	
J2	-0.8866**	-0.6034**	-0.3205**	C
E1	0.7387**	0.3273**	0.5317**	
E2	-0.8222**	-0.8444**	0.1322	C
I1	-0.8045**	-0.4400**	-0.4529**	
I2	-0.8413**	-0.8325**	0.0879	C
B1	0.6888**	0.2555*	0.5705**	
B2	0.8687**	0.8154**	-0.0241	C
F1	0.9034**	0.7652**	0.0998	C
F2	0.3627**	-0.1356	0.7079**	
C1	-0.5104**	-0.0233	-0.6731**	
C2	0.8501**	0.8262**	-0.0660	C

Table 8.17. Correlations of the hidden unit outputs with targets and difference between targets (** and * indicates significance at α=0.01 and 0.05 respectively)

In the first cluster (DGHJ), the "common" hidden units showed more focus on the level diagnosis than on trend, while the second hidden unit focused more on trend diagnosis. In these models, the hidden units were more *specialised* with regard to our target concepts. Model G was most representative of implementing this functionality. The differences in representations between the models of the two sub-clusters can be exemplified by looking at how the two hidden units in each model responded to variation in level and trend diagnosis. This can be illustrated by a simplified output function in each of the hidden units. The two most representative models of each sub-cluster were selected and are shown in figure 8.4. In the F model shown in the upper half of figure 8.4, hidden unit F1 implemented the common factor, and hidden unit F2 detected the exceptions to the rule that level and trend were correlated. In the G model, the hidden unit G1 detected the common aspects and trend, whereas hidden unit G2 detected the common aspects and level. Both the hidden units of the G model were inhibitory and thus, low level and trend gave high hidden unit response. In figure 8.4, predicted outputs were used. If using real outputs, noise made the output distributed around the "planes" drawn in figure 8.4.

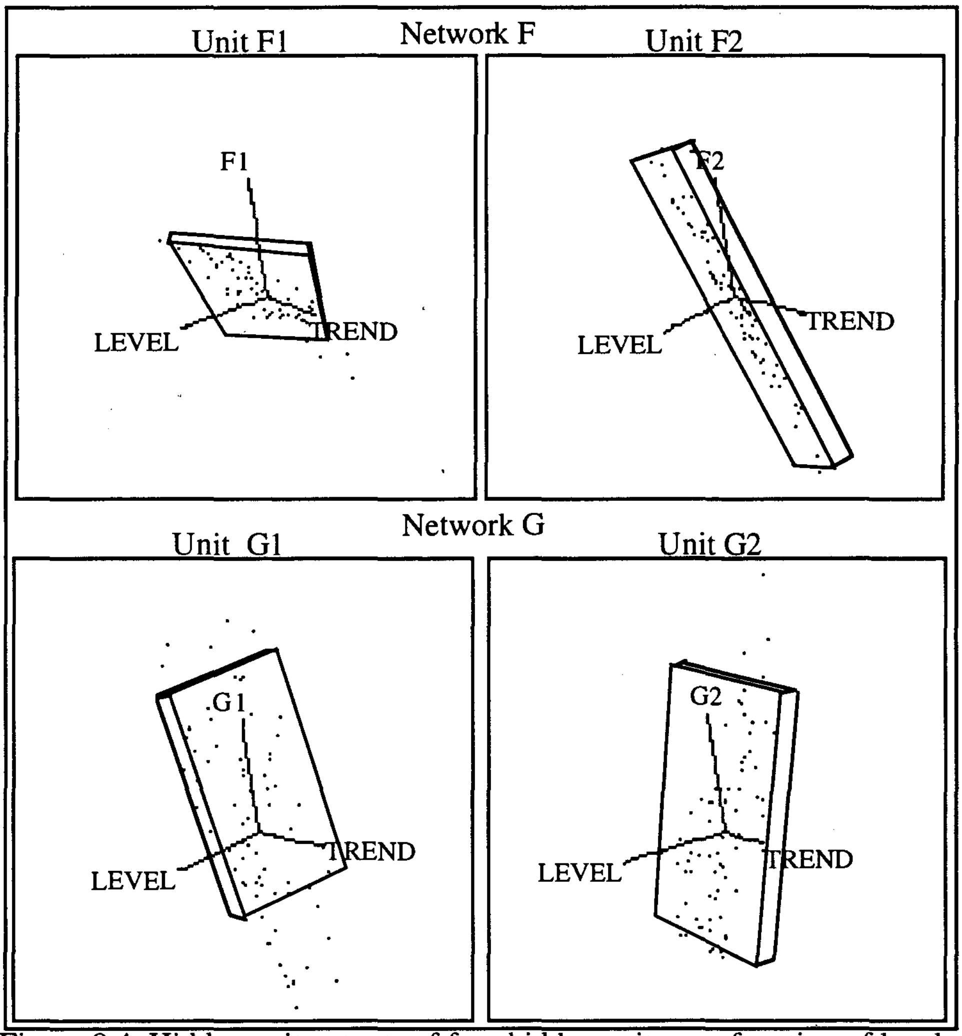

Figure 8.4. Hidden unit outputs of four hidden units as a function of level and trend. The response functions are simplified and illustrated as a "plane" in three dimensional space.

The outputs of the hidden units separated the cases in categories by constructing regions, but by using graded decision bounds. For illustration purposes, we restricted this analysis to the models with two hidden units. For the models with three and four hidden units, the cluster analysis split the diagnosis of level and trend in two main clusters, but functional groups of the same kind as found in figure 8.3 could not be detected. One reason may be that the larger models organised their response in the same functional way, or that they used almost as many functional organisations as there were models.

To proceed with the analysis of the representations any further, we assumed that a single model had to be selected. To investigate the representational differences between the smaller and larger models, we started by selecting the three models termed B2, B3 and B4 . These models seemed "representative" of their group. Another reason for selecting the B models, was that they seemed to have a representative common, and interesting separate hidden units. A Hinton diagram of all the weights in the three models is shown in figure 8.5.

We started the analysis by investigating the Hinton diagram of the connections between *hidden and output* units. Model B2 implemented the "difference" heuristic previously explained. This could be inferred from the pattern of hidden unit 2, forming a common hidden unit, and from the pattern of hidden unit 1, detecting the exceptions to the rule that level and trend were correlated. A study of the Hinton diagram of model B3 revealed a different organisation. In addition to the common hidden unit 2, this model had developed two specialised hidden units. Hidden unit 1 specialised on trend diagnosis, while hidden unit 3 specialised on level diagnosis. In model B4, we noticed that the common hidden unit was inhibitory, but it had the same structure as the other common hidden units. In addition, we noticed that the organisations of both model B2 and B3 were implemented. Hidden units 2 and 4 implemented the trend and level diagnosis detectors found in model B3, while hidden unit 1 implemented the difference heuristic from model B2, but with reversed signs.

When looking at the connections from the *input units to the hidden units*, the picture was more complex. The pattern of connections had previously been compared in the cluster analysis dendrogram shown in figure 8.2. The trend-oriented hidden units, B13 and B24, were grouped together in the bottom of the dendrogram of figure 8.2. These units had a weight pattern similar to each other, and were previously classified as units specialising on the trend diagnosis. Next, B44 and B33 were grouped together in the lower middle of the dendrogram. These units were classified as units specialising on level diagnosis. Finally, B14 and B12 were grouped in the middle of the dendrogram. Both these units implemented the "difference" heuristic, but with reversed signs.

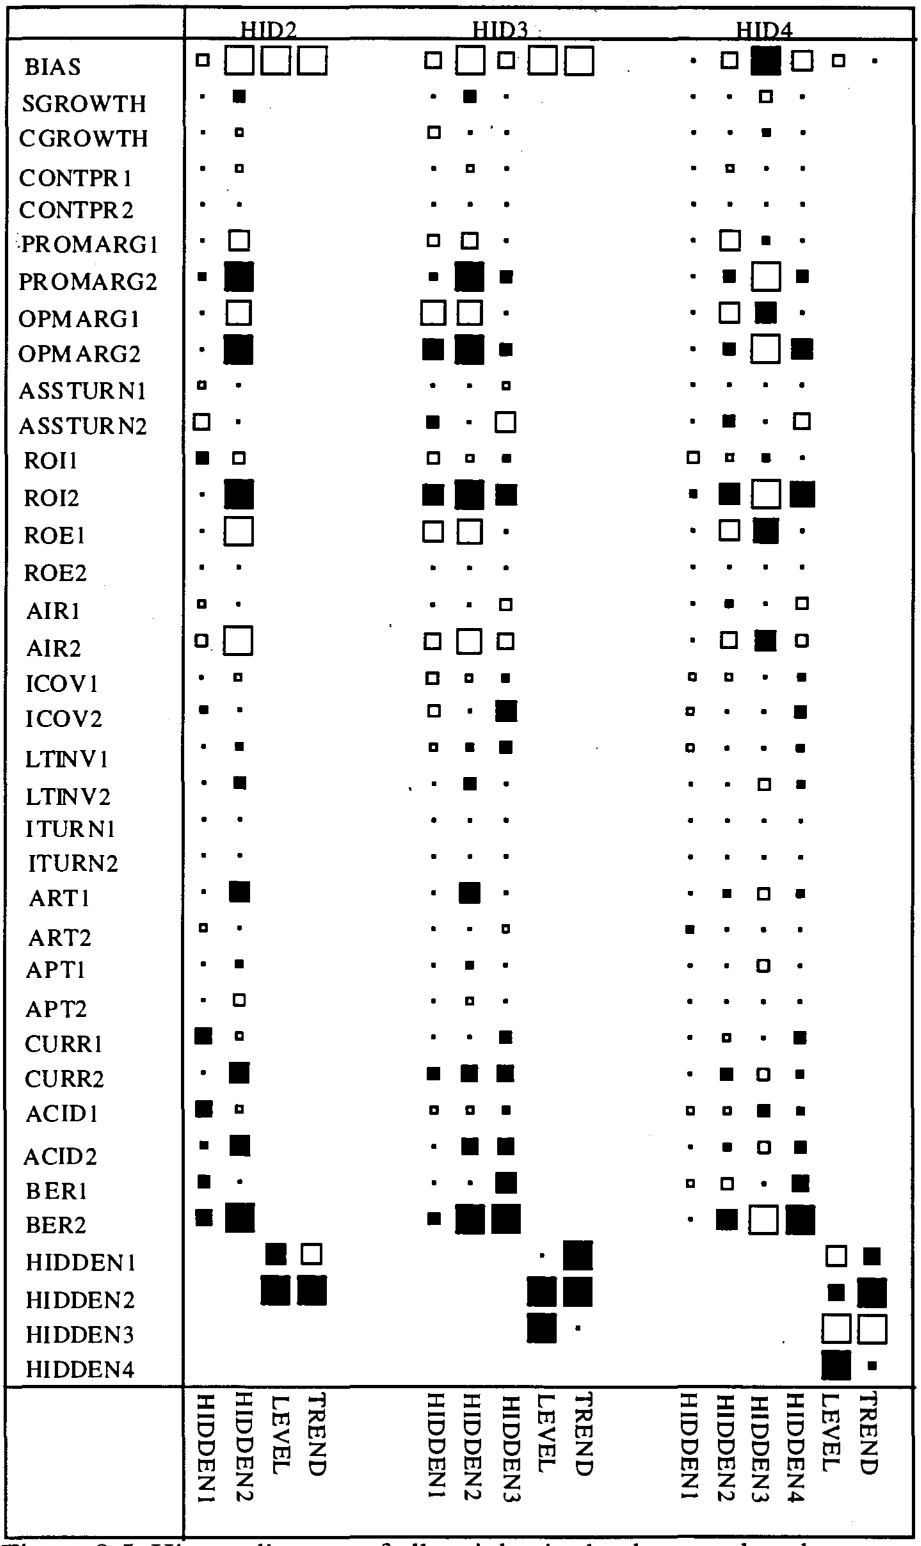

Figure 8.5. Hinton diagram of all weights in the three analysed combined models

Consequently, the cluster analysis of figure 8.2, proved useful for analysis of both common *and separate hidden units*. Analysis of how a task was performed could be studied by averaging the weight pattern of the hidden units belonging to the same cluster in a way similar to what was applied to the analysis of the common hidden units. We had identified clusters of units specialising on trend, level and "difference" detection. By studying the average weight pattern of the units in these clusters, we could explain *how* the units specialising on trend, level and "difference" detection performed their tasks.

The average weight pattern of the trend-oriented hidden units is shown in table 8.18.

Input unit	Mean weight	St. dev of weight	d.f.	t-value
SGROWTH	0.244	0.077	9	10.00**
CGROWTH	-0.269	0.116	9	-7.32**
CONTPR1	-0.232	0.122	9	-5.97**
CONTPR2	-0.027	0.093	9	-0.93
PROMARG1	-0.582	0.128	9	-14.30**
PROMARG2	0.428	0.153	9	8.86**
OPMARG1	-0.840	0.160	9	-16.52**
OPMARG2	0.883	0.208	9	13.41**
ASSTURN1	0.116	0.123	9	2.98*
ASSTURN2	0.470	0.251	9	5.91**
ROI1	-0.541	0.221	9	-7.74**
ROI2	1.039	0.267	9	12.29**
ROE1	-0.663	0.117	9	-17.91**
ROE2	0.071	0.120	9	1.87
AIR1	0.233	0.167	9	4.41**
AIR2	-0.496	0.108	9	-14.52**
ICOV1	-0.489	0.184	9	-8.40**
ICOV2	-0.142	0.163	9	-2.75*
LTINV1	-0.072	0.148	9	-1.54
LTINV2	0.159	0.137	9	3.65**
ITURN1	0.067	0.133	9	1.59
ITURN2	0.063	0.061	9	3.30**
ART1	0.281	0.097	9	9.14**
ART2	0.045	0.125	9	1.15
APT1	0.179	0.079	9	7.15**
APT2	-0.207	0.077	9	-8.46**
CURR1	-0.294	0.146	9	-6.36**
CURR2	0.374	0.111	9	10.65**
ACID1	-0.401	0.173	9	-7.32**
ACID2	0.232	0.101	9	7.27**
BER1	-0.299	0.162	9	-5.82**
BER2	0.641	0.141	9	14.30**

Table 8.18. Mean weight values, standard deviation and t-value of test of $\mu=0$ for the weights between input units and 10 trend-oriented hidden units (** and * indicates significant at $\alpha=0.01$ and 0.05 respectively)

In table 8.18, the sets of input units representing financial cues from two consecutive years with mean weights significantly different from zero ($\alpha=0.01$), are marked[1]. The absolute values of the mean weights indicated the importance of the respective input units. The values

[1] SGROWTH and CGROWTH are treated similarly to the financial cues of two consecutive years.

were comparable because of standardisation of input values. Cues presumed to represent "profitability" information had the highest absolute values, while the cues presumed to represent "financial structure" information had insignificant values. When comparing the pattern of weight values to the common hidden units' weight pattern, they showed a remarkable similarity. However, the pattern of the trend-oriented hidden units had smaller absolute weight values. This suggested that these units were a corrective to the common hidden units.

Input unit	Mean weight	St. dev of weight	d.f.	t-value
SGROWTH	0.170	0.064	12	9.56**
CGROWTH	-0.003	0.090	12	-0.13
CONTPR1	-0.068	0.132	12	1.86
CONTPR2	0.058	0.074	12	2.81*
PROMARG1	0.108	0.081	12	4.80**
PROMARG2	0.421	0.129	12	11.72**
OPMARG1	-0.021	0.132	12	0.58
OPMARG2	0.651	0.160	12	14.67**
ASSTURN1	-0.217	0.170	12	-4.62**
ASSTURN2	-0.441	0.272	12	-5.83**
ROI1	0.294	0.094	12	11.30**
ROI2	0.758	0.127	12	21.45**
ROE1	-0.207	0.091	12	-8.18**
ROE2	-0.088	0.110	12	-2.92*
AIR1	-0.315	0.210	12	-5.39**
AIR2	-0.469	0.208	12	-8.13**
ICOV1	0.315	0.146	12	7.75**
ICOV2	0.501	0.154	12	11.72**
LTINV1	0.428	0.134	12	11.50**
LTINV2	0.250	0.084	12	10.72**
ITURN1	-0.076	0.102	12	-2.68*
ITURN2	0.108	0.103	12	3.78**
ART1	0.233	0.112	12	7.45**
ART2	-0.250	0.118	12	-7.65**
APT1	0.078	0.094	12	3.02*
APT2	-0.069	0.081	12	-3.06**
CURR1	0.458	0.180	12	9.16**
CURR2	0.497	0.103	12	17.36**
ACID1	0.331	0.139	12	8.55**
ACID2	0.493	0.156	12	11.36**
BER1	0.545	0.172	12	11.43**
BER2	1.030	0.195	12	18.99**

Table 8.19. Mean weight values, standard deviation and t-value of test of $\mu=0$ for the weights between input units and 13 level-oriented hidden units (** and * indicates significant at α=0.01 and 0.05 respectively)

Another difference was that the trend-oriented hidden units had more similar values for the weights from units representing each of the two consecutive years' cue values. This meant that trend detection was performed by subtracting the values from the two consecutive years with almost equal weighting. Thus, the trend detector focused exclusively on change in the values of the financial cues from one year to the next. With the signs organised to indicate positive influence on the trend diagnosis for positive weight values, we found the sign patterns consistent with our expectations of how a "positive trend" detector should work. For example, we found that an increase in AIR worked in the opposite direction of an increase in ROI.

The next cluster of units was interpreted as representing the level-oriented hidden units. The average weight pattern of these units is shown in table 8.19.

Following the marking rule of the trend-oriented hidden units, we were left with the marked average weights in table 8.19. These weight values were very different from the weight values of the trend-oriented hidden units. First, the stimulus dimension "change" was not formed by these units. Second, cues presumed to represent information on "financing" were significant. The signs were the same for weights from units representing financial cue values of both of the consecutive years. Not surprising, this suggested it was the *magnitude* of these values that was important in correcting the output of the common hidden units to obtain a valid level diagnosis. Furthermore, the sign patterns were as expected. For example, high ROI consistently indicated a better level diagnosis while high AIR consistently indicated the opposite.

The third cluster of separate hidden units was presumed to indicate exceptions to the rule that level and trend diagnoses were strongly correlated. The average weight pattern of these "difference" units is shown in table 8.20. There were two types of exceptions to the rule that level and trend diagnoses were linearly correlated. One exception was when the level diagnosis was much higher than what could be expected from the trend diagnosis. The opposite was when the trend diagnosis was much higher than what could be expected from the level diagnosis[1]. On an average, the subjects' level diagnosis value was 0.15 higher than the trend diagnosis value. In table 8.20, weights of hidden units representing both these exceptions are shown. Some of the hidden units implemented the first exception, for example unit B12, and some implemented the second exception, for example unit C12. It was somewhat surprising that these exceptions were implemented by units belonging to the same cluster in the dendrogram, and thus, had a similar weight pattern. Upon further inspection, it was clear that the weight pattern of the two types was similar in weight values, but had opposite sign patterns[2]. In table 8.20, the signs have been turned so that the pattern illustrated by mean weights, is the pattern of a "difference" unit responding to the second exception.

---

[1] Or stated differently, when level diagnosis was lower than what could be expected from the trend diagnosis.
[2] Since the cluster analysis was performed on absolute weight values, the units were placed in the same cluster.

Since the typical pattern was level diagnosis being somewhat higher than trend, this was an interesting exception. These units usually responded highly to stimulus patterns where the level diagnosis was not good, but where there was a positive trend. We termed this exception the "high trend" exception.

Input unit	Mean weight	St. dev of weight	d.f.	t-value
SGROWTH	0.074	0.176	23	2.06*
CGROWTH	-0.141	0.098	23	-7.01**
CONTPR1	-0.171	0.131	23	-6.34**
CONTPR2	-0.103	0.137	23	-3.70**
PROMARG1	-0.261	0.198	23	-6.44**
PROMARG2	0.027	0.207	23	0.65
OPMARG1	-0.308	0.261	23	-5.79**
OPMARG2	0.061	0.270	23	1.12
ASSTURN1	0.206	0.222	23	4.55**
ASSTURN2	0.447	0.306	23	7.14**
ROI1	-0.415	0.186	23	-10.90**
ROI2	0.114	0.409	23	1.38
ROE1	-0.141	0.202	23	-3.43**
ROE2	0.165	0.149	23	5.41**
AIR1	0.273	0.271	23	4.93**
AIR2	0.024	0.213	23	0.56
ICOV1	-0.390	0.231	23	-8.29**
ICOV2	-0.284	0.193	23	-7.20**
LTINV1	-0.224	0.157	23	-6.97**
LTINV2	-0.106	0.151	23	-3.44**
ITURN1	0.051	0.153	23	1.66
ITURN2	-0.045	0.108	23	-2.05
ART1	0.030	0.185	23	0.81
ART2	0.214	0.188	23	5.58**
APT1	-0.005	0.142	23	-0.20
APT2	-0.125	0.104	23	-5.87**
CURR1	-0.300	0.203	23	-7.25**
CURR2	-0.096	0.223	23	-2.11*
ACID1	-0.352	0.193	23	-8.95**
ACID2	-0.130	0.169	23	-3.77**
BER1	-0.379	0.202	23	-9.17**
BER2	-0.223	0.334	23	-3.28**

Table 8.20. Mean weight values, standard deviation and t-value of test of μ=0 for the weights between input units and 24 "difference" units (** and * indicates significant at α=0.01 and 0.05 respectively)

From table 8.20, we first found that a unit responding to the "high trend" exception had positive weights to the ASSTURN units. This was not the case for any of the other units investigated so far. We could interpret this as an indication that high sales to assets was part of a positive trend, possibly creating expectations of an improved future level diagnosis. Second, we found that many of the weights were negative, and almost all weights to input units representing financial cues of the first year were negative. Not surprisingly, this indicated that it was difficult to be in an exceptionally positive trend when cues of the first year had high values.

Nine hidden units in the middle cluster of the dendrogram had, so far, not been classified. The pattern of incoming weights to the units showed close resemblance to the level-oriented hidden units, but investigation of the weight patterns between hidden units and output units by looking at the Hinton diagram of figure 8.1, showed that these units also strongly excited the trend diagnosis units. However, the weights to the level diagnosis units were the largest. The average weight pattern of these units is shown in table 8.21, with the signs of the weights of the inhibitory units turned.

Input unit	Mean weight	St. dev of weight	d.f.	t-value
SGROWTH	0.301	0.114	8	7.91**
CGROWTH	-0.065	0.071	8	-2.74*
CONTPR1	-0.051	0.093	8	-1.65
CONTPR2	0.061	0.107	8	1.72
PROMARG1	-0.124	0.150	8	-2.48*
PROMARG2	0.617	0.159	8	11.62**
OPMARG1	-0.223	0.107	8	-6.23**
OPMARG2	0.888	0.156	8	16.98**
ASSTURN1	-0.163	0.105	8	-4.66**
ASSTURN2	-0.134	0.225	8	-1.79
ROI1	0.001	0.088	8	0.05
ROI2	1.120	0.116	8	28.84**
ROE1	-0.387	0.095	8	-12.13**
ROE2	0.021	0.113	8	0.57
AIR1	-0.170	0.191	8	-2.66*
AIR2	-0.613	0.094	8	-19.38**
ICOV1	0.043	0.188	8	0.69
ICOV2	0.303	0.125	8	7.22**
LTINV1	0.349	0.119	8	8.74**
LTINV2	0.381	0.084	8	13.63**
ITURN1	-0.111	0.094	8	-3.55**
ITURN2	0.025	0.111	8	0.67
ART1	0.324	0.085	8	11.45**
ART2	-0.133	0.120	8	-3.33**
APT1	0.073	0.070	8	3.15*
APT2	-0.093	0.090	8	-3.12*
CURR1	0.159	0.156	8	3.05*
CURR2	0.479	0.094	8	15.30**
ACID1	0.136	0.161	8	2.53*
ACID2	0.563	0.111	8	15.17**
BER1	0.274	0.198	8	4.16**
BER2	1.116	0.137	8	24.28**

Table 8.21. Mean weight values, standard deviation and t-value of test of $\mu=0$ for the weights between input units and 9 units in the last cluster (** and * indicates significant at $\alpha=0.01$ and 0.05 respectively)

Inspection of the weight pattern in table 8.21, revealed that the units had a weight pattern close to the weight pattern of the level hidden units. In fact, the sign pattern of the two types of units was the same except for weights from the unit representing PROMARG1. In absolute values, the pattern was somewhat different. A larger difference between the values of the second and the first year cues was typical for these units. This meant the units also detected the difference between the financial cue values of the two consecutive years, and propagated this difference further. This may explain the positive connection these units had to the trend diagnosis unit. Except from the interpretation as a correction to the diagnoses, the representations formed by these units were difficult to interpret. Another suggestion may be that these units actually were redundant. Six of the nine hidden units were from the largest connectionist models, and the rest were from the models with three hidden units. These models did not have a significantly better fit than the smaller models, supporting the redundancy suggestion.

8.3 Conclusions

In this section, the main conclusions of the simulations are summarised with reference to the propositions made in chapter 4. First, main conclusions relating to the *performance* of the connectionist models are presented. Next, the main conclusions from the analysis of the connectionist model *representations* are summarised.

Compared to the simulation results of the bankruptcy classification measure of the financial diagnosis reported in chapter 7, the results for the level and trend diagnosis variables were much more promising. For the level diagnosis variable, performance results significantly better than the benchmark models' results were found. Both evaluated by cross validated average squared errors and by distribution properties of the error terms, the connectionist models outperformed the benchmarks for the level diagnosis. For the trend diagnosis variable, the results of the connectionist models were better than the benchmarks' results, but the difference in performance did not prove significant. These findings support proposition P1 of chapter 4.

For the connectionist models showing significantly better performance results than the benchmarks, a significant difference was also found between connectionist models with and without hidden units. This finding supported the proposition P2 of chapter 4 that the internal representations built by the multilayered connectionist models are the main reason for their better performance results.

The improved performance for combined models of correlated response variables hypothesised by several authors (e.g. Chakaborty et al., 1992) was not found significant in this study. A small improvement in performance was found for these models, but it was not significant. The simplicity of the combined models was, however, regarded to be a "significant" advantage when representational analysis was performed.

Except for these findings directly related to the propositions of chapter 4, a set of findings of "internal" relevance to connectionist modelling research was also done. All connectionist models showed an expected relationship between learning error[1], cross validated error, learning time and complexity. Learning error continuously decreased with learning time, while cross validated error was a U-shaped function of learning time[2]. Similarly, learning error decreased continuously with increased complexity of the models, while cross validated error was a U-shaped function of complexity[3]. Both these findings were explained by model overfit. Furthermore, models with a larger number of hidden units overfitted earlier during learning, and never reached the low minimum error of the simpler models. This indicated that using the optimal stopping point rule as an alternative to the more intricate optimal hidden units rule, as suggested by several authors (e.g. Smith, 1993), was not satisfactory. Connectionist and artificial neural network simulations traditionally report squared errors in some form as performance measure. However, our findings showed that additional analysis of the distribution of error terms was equally important if a satisfactory evaluation of performance was to be made. A further finding was that this evaluation was particularly relevant for response variables that had distributions deviating from the normal distribution.

Due to the small number of cases, the cross validation procedure was not used for generating the connectionist models that were used for representational analysis. A combined model showed no loss in performance and was selected for further analysis. The representations of the models were sensitive to initial weights, and to control for this instability, functional properties shared by many versions of a connectionist model were studied. All connectionist models used a common hidden unit to represent common aspects of the level and trend diagnosis . The common hidden units worked as "general condition detectors" and used cues presumed to represent all diagnostic concepts; "profitability", "financing", "liquidity" and "leverage", to form its internal representations.

Even though the common hidden units were interpreted as forming *local* representations[4], they did not specialise on one specific diagnostic concept. Consequently, diagnostic concepts, such as "profitability" and "financing", were *distributed* in the representation of the common

[1] See section 7.3.
[2] Because focus in this dissertation is on the generalisation properties of connectionist models, extensive reporting of learning error results was not made. However, an example of the typical relationships between learning and cross validated errors is shown in appendix K.
[3] See previous footnote.
[4] Due to their large variance in weights.

hidden unit. None of the separate hidden units were found to specialise on diagnostic areas. Instead, the units implemented *heuristics* that detected exceptions to the rule that level and trend diagnoses were positively correlated, with level diagnosis somewhat larger than trend diagnosis in value. Three groups of interpretable separate hidden units were found. The "difference" detectors had large output values in cases with positive trend in an unfavourable situation and low output values in cases with negative trend in favourable situations. Thus, the heuristic implemented by these units was rather complex. In most models with more than two hidden units, specialised "trend" and "level" detectors were found. These units worked as detectors of a particularly favourable or unfavourable trend or level diagnosis in cases where the common hidden unit "predicted" otherwise. Consequently, these units also worked as exception detectors. Like the common hidden units, the specialised units did not focus on specific diagnostic areas. Rather, the representation consisted of a more distributed[1] weight pattern with high valued weights to input units presumed to represent several diagnostic concepts. The concepts represented by these units were clearly definable by analysis of representations and responses of the model, but traditional diagnostic concepts were distributed over these complex representational units.

Tests were run on models restricting the network by allowing connections only in patterns that corresponded to the expected specialisation of hidden units on particular diagnostic concepts. Such restricted[2] models have fewer free parameters and should implement a representation with hidden units detecting properties of the four diagnostic concepts; "profitability", "financing", "liquidity" and "leverage". The performance results of a restricted model with four hidden units[3] are shown in table 8.22.

Iterations:	5000	10000	15000	20000	25000	30000
Model:						
Restricted level	0.182	0.204	0.211	0.232	0.238	0.246
Restricted trend	0.401	0.358	0.373	0.387	0.404	0.418

Table 8.22. Mean squared error (MSE) of the level and trend diagnoses in a restricted model (N=75)

The results of the restricted model were somewhat worse than the combined model of section 8.1, but they were still surprisingly similar. A Hinton diagram of the restricted model is shown in figure 8.6.

[1] Distributed with respect to the diagnostic concepts "profitability", "financing", "liquidity" and "leverage".
[2] The term "restricted" has been used by several authors (e.g. Haykin, 1994, p. 25) of a network where prior knowledge of the task or problem is used to reduce the number of free parameters (weights and biases).
[3] Corresponding to the four diagnostic concepts; "profitability", "financing", "liquidity", and "leverage".

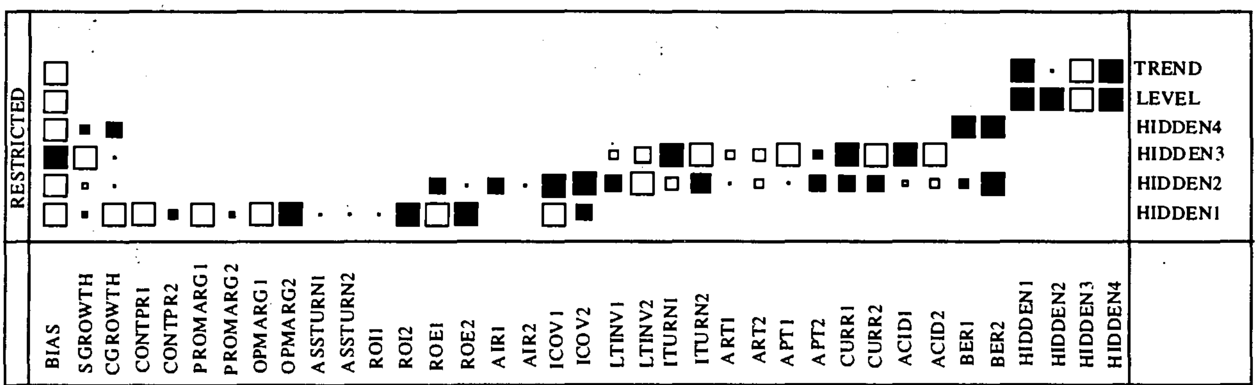

Figure 8.6. Hinton diagram of weights in a restricted connectionist model

Figure 8.6 shows a weight pattern somewhat different from what was expected. Hidden units 1, 2, 3 and 4 were manually connected to units assumed to represent particularly relevant cues to the diagnosis of "profitability", "financing", "liquidity" and "leverage" respectively. From figure 8.6, we found that hidden unit 2 worked as a level-oriented hidden unit described above. One reason for drawing this conclusion was that it was not very likely that financing was only relevant to level diagnosis. A positive-negative weight pattern of the input to hidden unit connections were found for hidden units 1 and 3 , suggesting that these units worked more as common hidden units than as units focusing specially on "profitability" and "liquidity" respectively.

To test the hypothesis that the restricted model did not partition its representation by diagnostic concepts, we computed the correlations between hidden unit outputs and composite judge diagnoses of the respective diagnostic areas[1]. These correlations are shown in table 8.23.

Diagnostic area	HIDDEN1	HIDDEN2	HIDDEN3	HIDDEN4
PROFLEVEL	0.644**	0.328**	-0.288*	0.355**
PROFTREND	0.871**	-0.026	-0.199	0.111
FINLEVEL	0.350**	0.453**	-0.454**	0.562**
FINTREND	0.607**	0.132	-0.426**	0.353**
LIQLEVEL	0.193	0.342**	-0.516**	0.565**
LIQTREND	0.395**	0.085	-0.543**	0.387**
LEVLEVEL	0.374**	0.602**	-0.235*	0.832**
LEVTREND	0.659**	0.239*	-0.383**	0.461**

Table 8.23. The restricted model's correlations of hidden unit outputs with composite judge diagnoses of the four diagnostic areas "profitability", "financing", "liquidity" and "leverage" for level and trend respectively (** and * indicates significantly different from 0 at α=0.01 and 0.05 respectively) (N=75)

[1] Separate measures of subjects' evaluation of these diagnostic areas had been collected. See chapter 5 for an introduction to the measures applied.

Hidden unit 1 had its highest correlations with the diagnostic variables of "profitability". As suspected, hidden unit 2 worked as a level-oriented hidden unit, having high correlations with the level variables in all diagnostic areas. Correlations with the expected variables of "financing" were not higher than correlations with other level diagnosis variables. Hidden units 3 and 4 had their highest correlations with the "liquidity" and "leverage" diagnosis variables, respectively. Consequently, the restricted model partitioned its task to a certain extent by diagnostic area, but "financing", as a diagnostic variable, was treated by hidden units 1, 3 and 4 in common. This organisation left hidden unit 2 to implement a level-oriented hidden unit similar to the one found in section 8.2.2. For the connectionist model, this way of partitioning the financial diagnosis task between the hidden units was more effective than the partitioning suggested by theory. Despite the imposed restrictions, the connectionist model had so many degrees of freedom that it was possible to perform the task differently from what was proposed by theory.

Except for the conclusions on the representations of the connectionist models directly related to proposition P3 of chapter 4, some findings of "internal" relevance to connectionist modelling research was also done. The weight pattern of a model varied considerably at the end of the learning phase, even though the performance of the different versions was comparable. The main reason was that different initial weight values gave different representations at the end of the learning phase. Consequently, to understand the representations of a connectionist model, several versions of a model should be studied. One way to analyse the common aspects of different versions of a model was to use cluster analysis of the weights of several versions of a model simultaneously. With this method, hidden units with similar representations in several versions of a model were detected. From this analysis, the average weight values of hidden units implementing the same functional organisation could be computed and analysed further.

To summarise the findings on the connectionist model representations, we showed that the original stimulus dimensions of the financial cues were not used directly by the connectionist models to form the level and trend diagnoses. Instead, abstracted and transformed stimulus dimensions were "detected" by the hidden units. These intermediate abstractions should be termed stimulus dimensions because the hidden units showed continuous outputs. With a continuous output, the variables form seemed best suited to describe the representational form of the hidden units. Furthermore, the stimulus dimensions abstracted by the hidden units were configural. They were composed by transforming whole patterns of cues representing different diagnostic concepts in the stimulus material. An interpretations of the molar functional organisation of the financial diagnosis task performed by the connectionist models, could best be explained by applying a rule-plus-exception heuristic.

It is not possible to *test* the cognitive relevance of these representations. However, the connectionist models seemed to implement a rule and exception principle of organising their internal representations. Similar principles have been proposed by several authors in classification research (e.g. Nosofsky, Palmieri & McKinley, 1994) as cognitively relevant. Even though the representations developed by the connectionist models were not as expected, the claim that they were cognitively relevant may still seem justifiable.

Despite the interesting conclusions drawn from the simulations of the level and trend diagnosis models above, several problems still existed. The number of cases was too small to effectively constrain the large number of free parameters in the models. This was particularly evident for trend diagnosis. Limiting the number of hidden units in the models was not enough to reduce the number of free parameters. Another way to limit the number of free parameters is by selecting a smaller set of diagnostic cues from the stimulus material. Theoretically, such models should have improved generalisation properties, since they will make overfit less likely. Empirically, such *constrained* models have shown considerable success in some application areas (e.g. Le Cun, 1990). We now turn to the development of such models.

Chapter 9. A constrained connectionist model of continuous response

In chapter 8, we found that the connectionist models had too many free parameters even with few hidden units. One way of reducing the number of free parameters is by decreasing the number of input units. At least two principles can be used in deciding which input units to exclude. By relying on the information given by the *subjects* regarding the importance of a cue for a particular diagnosis, unimportant cues can be excluded. Results of models using this method of cue selection and parameter reduction are reported in section 9.3.

A second method is to use some *quantitative measure* of the importance of a cue in producing the specific output. The last method is used by traditional linear methods, such as stepwise regression. For neural networks, the relationship between input and output is complex and nonlinear. No *obvious* [1] measure of the importance of a cue exists (Garson, 1991), but some form of sensitivity analysis can be used (Moody, 1993; Refenes et al., 1995).

By altering the input values for each variable on each case, the sensitivity of the output values to changes in input values can be estimated. A 5 % change in the input values was performed following the principles used and explained in chapter 7. The percentage change in output values caused by these changes is shown in table 9.1[2].

In table 9.1, the effects of changing the input values are shown for both the level and trend diagnoses. In the table, the highest effect on either level or trend diagnosis is ranked for each set of input units grouping the financial cue values of two consecutive years together. SGROWTH and CGROWTH were treated similarly. Setting the threshold of implementing a cue in the model at 10 % effect on any of the two diagnosis variables, leaves us with six cues of the two consecutive years. However, this rule excluded any traditional "liquidity" indicator. Consequently, we included the highest scoring "liquidity" indicator and dropped the lowest scoring indicator of the highly correlated cues; PROMARG and OPMARG. The resulting cues of a constrained model were consequently; OPMARG, ROI, ROE, AIR, CURR and BER.

---

[1] We showed in chapter 7 and 8 how both sensitivity analysis and analysis of average weight values could be used to evaluate the importance of an input unit.

[2] Because minimum cross validated average squared error was found very similar for models with 2 and 4 hidden units, 10 versions of models with 2, 3 and 4 hidden units are used in the sensitivity analysis. Thus, the number of observations is 3*10*75=2250.

Jogged unit	Mean effect on level diagnosis	Mean effect on trend diagnosis	Highest mean effect	Rank order of highest effect
SGROWTH	5.37	6.71	6.71	11
CGROWTH	-3.22	-5.41	-5.41	11
CONTPR1	-2.97	-5.21	-5.21	15
CONTPR2	0.68	-0.21	0.68	15
PROMARG1	-5.24	-10.14	-10.14	4
PROMARG2	13.00	14.03	14.03	4
OPMARG1	-8.92	-14.98	-14.98	2
OPMARG2	19.78	22.39	22.39	2
ASSTURN1	-2.41	0.49	-2.41	13
ASSTURN2	-1.79	5.30	5.30	13
ROI1	-1.89	-8.22	-8.22	1
ROI2	24.95	28.63	28.63	1
ROE1	-10.13	-13.44	-13.44	5
ROE2	0.08	2.48	2.48	5
AIR1	-2.32	2.36	2.36	6
AIR2	-12.80	-13.23	-13.23	6
ICOV1	-1.85	-7.86	-7.86	10
ICOV2	4.57	0.48	4.57	10
LTINV1	4.78	1.54	4.78	12
LTINV2	6.26	5.64	6.26	12
ITURN1	-1.45	-0.41	-1.45	16
ITURN2	1.87	1.74	1.87	16
ART1	7.55	8.46	8.46	9
ART2	-2.48	0.38	-2.48	9
APT1	3.29	4.32	4.32	14
APT2	-3.80	-5.27	-5.27	14
CURR1	0.48	-4.69	-4.69	7
CURR2	9.48	8.84	9.48	7
ACID1	-1.27	-6.88	-6.88	8
ACID2	9.23	7.67	9.23	8
BER1	1.61	-4.53	-4.53	3
BER2	20.78	19.04	20.78	3

Table 9.1. The effects of jogging input values 5 % in the combined model (N=2250)

The number of free parameters in a connectionist model with one hidden layer is partly determined by the number of weights. The number of weights can be calculated as:

$$w = h(i + o + 1) + o, \tag{9.1}$$

where w is the number of weights, h is the number of hidden units, o is the number of output units and i is the number of input units. The relevant values of w are shown in table 9.2 for combined models with two output units.

Input units\Hidden units	0	2	4	6	8	10	12	14
32	33	72	142	212	282	352	422	492
12	13	32	62	92	122	152	182	212

Table 9.2. Number of weights in the different combined models

From table 9.2, we see that reducing the number of input units to 12, gives a notable decrease in the number of free parameters. In both connectionist and traditional models, this is assumed to improve the generalisation ability of the model (Ripley, 1993). Other authors (e.g. Smith, 1993) have argued that stopping the training at the overfit point should give the same result as reducing the number of free parameters. However, we showed in chapter 8 that the two principles were not equivalent. The reduction of free parameters by individual weight elimination have been suggested by several authors (Karnin, 1990; Weigend et al., 1991), and the sensitivity analysis based "pruning" performed in this study, has been shown to produce similar results (Moody, 1993; Moody & Utans, 1995). Pruning can be performed both during and after learning. Some other pruning techniques were briefly introduced in chapter 3.

The results for the models derived with sensitivity based reduction of parameters are reported in section 9.1. An analysis of the representations of these *constrained* models is reported in section 9.2. A test of similar models constrained by using subjects' evaluations of cue importance is reported in section 9.3. The main conclusions drawn from the simulations of all the constrained models are summarised in section 9.4.

9.1 Performance results

The constrained model was tested using procedures and parameter values similar to what was used in chapter 7 and 8, and reported in chapter 6. The results of the cross validation procedure for the level diagnosis are shown in table 9.3.

Iterations:	5000	10000	15000	20000	25000	30000	Regr.[1]
Model:							
Regr. (all)							0.221
Regr. (stepw.)							0.212
HID0	0.185	0.184	0.187	0.187	0.188	0.189	
HID2	0.251	0.162	0.160	0.160	0.161	0.160	
HID4	0.198	0.166	0.164	0.164	0.165	0.166	
HID6	0.191	0.164	0.162	0.163	0.161	0.164	
HID8	0.192	0.171	0.169	0.166	0.169	0.168	
HID10	0.183	0.174	0.168	0.170	0.170	0.169	
HID12	0.178	0.168	0.166	0.167	0.170	0.166	
HID14	0.174	0.167	0.169	0.163	0.168	0.169	

Table 9.3. Mean squared error (MSE) of the level diagnosis in a constrained model (N=75)

[1] Performance results of the regression models are placed in a separate column. These models were estimated using traditional OLS methods.

In table 9.3, cross validated average squared errors are generally lower than in chapter 8. For models having their minimum at 30000 iterations, tests were run to see if MSE continued to decrease after 30000 iterations. However, no model had minimum values for learning time beyond 30000 iterations[1]. The best performance was found for the connectionist model with two hidden units. As in chapter 8, the connectionist model with hidden units showed better fit than both the model without hidden units and the benchmarks. A t-test of the difference in cross validated average squared error between the best connectionist model and the best benchmark[2] was significant, and in favour of the connectionist model at α=0.01 (t=2.73, d.f.=74). Furthermore, the model without hidden units now also showed somewhat better results than the benchmarks.

The stepwise regression model had a lower MSE than the best regression model of chapter 8. Because of multicolinearity, the stepwise procedure used by SPSS in building the 75 regression models of chapter 8 did not give the best possible models. Much of this multicolinearity was reduced by constraining the number of inputs, and performance results of the regression models also improved. However, the improvement in the connectionist models was much greater than the improvement in the benchmark models.

To test the improved results, t-tests of the differences between the MSE values of the models in chapter 8 and the *constrained* models were performed. The results are shown in table 9.4.

Iterations:	5000	10000	15000	20000	25000	30000
HID2	-1.77	0.87	1.37	1.99*	2.07*	2.31*

Table 9.4. T-tests of the best level model of chapter 8 (HID2) vs. best constrained connectionist model ($\mu_{old} - \mu_{new}$) at increasing number of iterations (* indicates significant at α=0.05, d.f.=74)

For a comparable number of iterations, the constrained model performed significantly better. However, the best model of chapter 8 had its minimum MSE at 10000 iterations[3], and the difference in MSE between the two *best* models was not significant at α=0.05. Still the MSE was *generally* lower for the constrained models. The minimum was also generally found for a larger number of iterations than in chapter 8. This implied overfit did not occur so early during learning in the constrained models. This observation further supported our conclusion from the simulations of chapter 8, that stopping learning in an "oversized" model early did not give the same performance as a smaller connectionist model stopped at the minimum MSE point. Stopping the learning and reducing complexity of the models may be equivalent when

[1] This conclusion was valid for the separate level and trend models, and for the combined models.
[2] The stepwise regression model.
[3] Model with 2 hidden units after 10000 iterations

it comes to controlling the mapping on learning samples, but this rule did not seem valid for generalisation errors.

To test if the distribution of the errors was similar to the models of chapter 8, correlations of error with level target value and distance from target value were computed. The results are shown in tables 9.5 and 9.6.

Iterations:	5000	10000	15000	20000	25000	30000	Regr.
Model:							
Regr. (all)							0.169
Regr. (stepw.)							0.177
HID0	0.075	0.090	0.108	0.118	0.124	0.127	
HID2	-0.138	0.001	0.008	0.011	0.017	0.022	
HID4	-0.058	0.042	0.034	0.042	0.046	0.048	
HID6	0.005	0.020	0.036	0.040	0.051	0.047	
HID8	-0.043	0.025	0.047	0.052	0.057	0.064	
HID10	0.048	0.050	0.049	0.061	0.059	0.070	
HID12	0.003	0.031	0.054	0.057	0.068	0.070	
HID14	0.033	0.046	0.068	0.058	0.060	0.056	

Table 9.5. Correlations of SE and target for the level diagnosis in a constrained model (N=75)

Iterations:	5000	10000	15000	20000	25000	30000	Regr.
Model:							
Regr. (all)							0.149
Regr. (stepw.)							0.128
HID0	0.168	0.147	0.126	0.119	0.109	0.104	
HID2	0.634**	-0.017	-0.056	-0.083	-0.097	-0.104	
HID4	0.304**	-0.007	-0.030	-0.050	-0.067	-0.090	
HID6	0.196	0.038	0.006	-0.004	-0.019	-0.055	
HID8	0.173	0.038	0.009	-0.023	-0.045	-0.075	
HID10	0.147	0.051	0.030	0.012	-0.023	0.052	
HID12	0.178	0.056	0.042	0.019	-0.010	-0.038	
HID14	0.156	0.086	0.054	0.013	0.001	-0.037	

Table 9.6. Correlations of SE and distance from mean target for the level diagnosis in a constrained model (N=75)(* indicates significant at α=0.05)

The pattern of correlations in table 9.5 and 9.6, showed roughly the same pattern as the similar measures in chapter 8. The distribution of the connectionist model errors was favourable, with errors distributed uniformly over the range of the target value.

Separate simulations of trend diagnosis were run following the same procedures as reported above. The cross validated average squared errors of these models are shown in table 9.7.

Iterations:	5000	10000	15000	20000	25000	30000	Regr.
Model:							
Regr. (all)							0.321
Regr. (stepw.)							0.356
HID0	0.305	0.300	0.302	0.301	0.301	0.303	
HID2	0.563	0.308	0.306	0.289	0.285	0.279	
HID4	0.456	0.307	0.296	0.288	0.280	0.279	
HID6	0.414	0.302	0.302	0.298	0.299	0.300	
HID8	0.414	0.310	0.306	0.300	0.298	0.297	
HID10	0.380	0.307	0.301	0.300	0.288	0.289	
HID12	0.376	0.312	0.304	0.300	0.298	0.290	
HID14	0.342	0.302	0.308	0.296	0.291	0.294	

Table 9.7. Mean squared error (MSE) of the trend diagnosis in a constrained model (N=75)

Again, the simulations showed reduced MSEs when compared to the corresponding results of chapter 8. Rather surprising were the poor results of the stepwise procedure of SPSS, but this suggested multicolinearity still caused problems of variable selection for the stepwise procedure. However, the results for the full regression model were better than for the benchmarks of chapter 6. Consequently, the best traditional benchmark for evaluating the trend diagnosis of the connectionist models was the full regression model with the constrained set of variables.

A t-test of the difference in cross validated average squared error between the best connectionist model and the best benchmark now showed a significant difference in favour of the connectionist model at α=0.05, (t=2.24, d.f.=74). Thus, even stronger support of proposition P1 of chapter 4 was provided.

Strong support of proposition P2 had previously been found for the level diagnosis. Now, the difference between the best connectionist model without hidden units[1] and the best model with hidden units[2] was observable for trend diagnosis also. However, a test of the difference did not prove significant at α= 0.05 (t=1.68, d.f.=74) when a two sided test was used, but since proposition P2 was formulated in favour of the connectionist model only, significance at α= 0.05 was found when a one-sided test was used. Thus, some support for proposition P2 was provided for trend diagnosis also.

To test the improved results of the constrained connectionist models, t-tests of the difference between the MSE values of the models in chapter 8 and the constrained models were performed. The results are shown in table 9.8.

---

[1] At 10000 learning iterations.
[2] Model with 4 hidden units at 30000 iterations.

Iterations:	5000	10000	15000	20000	25000	30000
HID4	-1.56	1.25	1.76	1.74	2.00*	2.05*

Table 9.8. T-tests of the best trend model of chapter 8 (HID6) vs. best constrained connectionist model (μ_{old} - μ_{new}) at increasing number of iterations (* indicates significant at α=0.05, d.f.=74)

The differences tested in table 9.8 were significant for 25000 and 30000 iterations. Furthermore, the difference in MSE between the *best* trend diagnosis model of chapter 8[1] and the best constrained model was also significant at α= 0.05 (t=2.03, d.f.=74).

The trend diagnosis models of chapter 8 had biased distributions of errors over the range of target values. To test the distribution of errors for the constrained model, correlations with trend target and distance from mean trend target were calculated. The correlations are shown in tables 9.9 and 9.10.

Iterations:	5000	10000	15000	20000	25000	30000	Regr.
Model:							
Regr. (all)							-0.024
Regr. (stepw.)							0.146
HID0	-0.030	-0.058	-0.049	-0.044	-0.045	-0.048	
HID2	0.114	-0.011	-0.087	-0.066	-0.040	-0.039	
HID4	0.110	-0.072	-0.075	-0.075	-0.063	-0.073	
HID6	0.092	-0.080	-0.072	-0.059	-0.054	-0.061	
HID8	0.124	-0.082	-0.078	-0.064	-0.054	-0.072	
HID10	0.084	-0.079	-0.047	-0.062	-0.062	-0.076	
HID12	0.132	-0.082	-0.067	-0.039	-0.028	-0.038	
HID14	0.053	-0.072	-0.074	-0.071	-0.064	-0.060	

Table 9.9. Correlations of SE and target for the trend diagnosis in a constrained model (N=75)

Iterations:	5000	10000	15000	20000	25000	30000	Regr.
Model:							
Regr. (all)							0.264*
Regr. (stepw.)							0.256*
HID0	0.362**	0.325**	0.315**	0.304**	0.298**	0.284**	
HID2	0.896**	0.338**	0.269*	0.262*	0.275*	0.253*	
HID4	0.705**	0.309**	0.307**	0.285*	0.275*	0.251*	
HID6	0.581**	0.311**	0.293*	0.282*	0.275*	0.253*	
HID8	0.583**	0.313**	0.298**	0.276*	0.269*	0.246*	
HID10	0.506**	0.329**	0.295*	0.295*	0.284*	0.268*	
HID12	0.521**	0.310**	0.286*	0.287*	0.246*	0.248*	
HID14	0.455**	0.324**	0.306**	0.297**	0.263*	0.258*	

Table 9.10. Correlations of SE and distance from mean target for the trend diagnosis in a constrained model (N=75) (** and * indicates significant at α=0.01 and 0.05 respectively)

[1] Model with 6 hidden units after 20000 iterations.

Corresponding closely to the correlation pattern of the trend diagnosis models of chapter 8, the constrained model had low correlations with targets, but high positive correlations with distance from mean targets. In table 9.10, we found that all correlations were significantly different from 0 (α=0.05). This result was somewhat worse than the results in chapter 8 and further supported our conclusion that evaluating connectionist models by cross validated average squared errors only, was insufficient. Errors were largest for cases with extreme positive and negative values. If the constrained model was more regressive than the model of chapter 8, this result should not be surprising. One could assume that the reduction in the number of input units would make the model more regressive. Hypothetically, this should reduce errors on targets close to the mean target, but increase errors on the extreme targets. In summation, the effect on MSE was positive, but the effect on the distribution of errors was negative.

To test the proposition made by several authors (e.g. Bounds et al., 1990; Chakaborty et al., 1992), that a combined model with several output units performed better than separate models when these variables were correlated, a combined model simulation was run. The results for the level diagnosis of these *constrained and combined* models are shown in table 9.11.

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

Table 9.11. Mean squared error (MSE) of the level diagnosis in a constrained and combined model (N=75)

From table 9.11, we found that the combined model performed better than the separate level diagnosis model. This was rather surprising compared to the results of chapter 8, but corresponded well to the hypothesis of Chakaborty et al. (1992), referred to above. A t-test of the difference in MSE between the best separate and combined level diagnosis models, showed that the combined model performed *significantly* better at α=0.05 (t=2.31, d.f.=74). For the models of chapter 8, such a difference was not found. Consequently, with fewer input units, the performance improved considerably in a combined model. One is led to suggest that the model could use its improved representations of trend-relevant stimulus dimensions to improve its level diagnosis and vice versa.

A t-test of the difference in MSE between the best constrained and combined level diagnosis model and the corresponding combined model of chapter 8[1] is shown in table 9.12.

Iterations:	5000	10000	15000	20000	25000	30000
HID4	0.91	1.11	1.79	2.25*	2.48*	2.97**

Table 9.12. T-tests of the best combined model of chapter 8 (HID2) vs. the best constrained and combined connectionist model (level diagnosis)($\mu_{old} - \mu_{new}$) at increasing number of iterations (* and ** indicates significant at α=0.05 and 0.01 respectively, d.f.=74)

The tests of table 9.12 also showed that the combined model with fewer input units significantly outperformed the *corresponding* models of chapter 8.

The results of the constrained and combined model for trend diagnosis are shown in table 9.13.

Iterations:	5000	10000	15000	20000	25000	30000
Model:						
HID0	0.307	0.303	0.301	0.301	0.302	0.299
HID2	0.473	0.328	0.290	0.277	0.269	0.267
HID4	0.389	0.302	0.287	0.281	0.280	0.282
HID6	0.394	0.303	0.285	0.280	0.278	0.276
HID8	0.369	0.298	0.287	0.287	0.281	0.285
HID10	0.365	0.308	0.292	0.289	0.285	0.291
HID12	0.352	0.289	0.280	0.284	0.275	0.282
HID14	0.342	0.292	0.290	0.289	0.283	0.287

Table 9.13. Mean squared error (MSE) of the trend diagnosis in a constrained and combined model (N=75)

From table 9.13, we found that the performance of the constrained and combined models was somewhat better than that of the corresponding separate models. However, the difference between the two best models was not significant at α=0.05 (t=0.83, d.f.=74), even though the difference was in favour of the combined model.

However, the results of table 9.13 showed that the constrained and combined models performed better than the combined model of chapter 8. This finding was similar to the finding for the level diagnosis. Pairwise t-tests of the differences in cross validated average squared errors are shown in table 9.14.

[1] Model with 2 hidden units at 15000 iterations.

Iterations:	5000	10000	15000	20000	25000	30000
HID2	-2.28*	-0.01	1.33	2.12*	2.39*	2.51*

Table 9.14. T-tests of the best combined model of chapter 8 (HID10) vs. the best constrained and combined connectionist model (trend diagnosis) ($\mu_{old} - \mu_{new}$) at increasing number of iterations (* indicates significant at α=0.05, d.f.=74)

The tests showed that the constrained combined model *significantly* outperformed the combined model of chapter 8 for trend diagnosis also. The performance of the best constrained and combined model was significantly better than the best combined model of chapter 8[1] at α=0.05 (t=2.10, d.f.=74). This corresponded to similar findings for the constrained and combined connectionist models of the level diagnosis.

Even though we had previously established that the connectionist models with hidden units significantly outperformed the models without hidden units, and consequently found support for the proposition P2 of chapter 4, the strongest test would be a comparison of the two constrained and combined models developed here. For the level diagnosis, the difference in MSE between the best constrained and combined connectionist model without hidden units[2] and the best model with hidden units[3] was significant at α=0.05 (t=2.27, d.f.=74). For the trend diagnosis, the difference in MSE between the best constrained and combined connectionist model without hidden units[4] and the best model with hidden units[5] was not significant at α=0.05 (t=1.82, d.f.=74). However, the last conclusion was open for judgement. Since the proposition P2 was formulated in the favour of the connectionist models with hidden units only, a one-sided test may be applied. In this case, the t-test observator value 1.82 was significant at p=0.036. Since strong support had previously been found of proposition P2 for the level diagnosis, we concluded that proposition P2 was generally supported.

A problem with the models of chapter 8 and the separate trend diagnosis model of this section, was the unfavourable distribution of errors over the range of target values. The errors were correlated with distance from mean targets, indicating that the models' performance was poorer for the extreme targets. The same measures for the constrained and combined models are shown in table 9.15.

[1] Model with 10 hidden units at 10000 iterations.
[2] At 10000 learning iterations.
[3] The model with 4 hidden units at 30000 iterations.
[4] At 30000 learning iterations.
[5] The model with 2 hidden units at 30000 iterations.

Iterations:	5000	10000	15000	20000	25000	30000
Model:						
HID0	0.360**	0.325**	0.313**	0.307**	0.295*	0.292*
HID2	0.599**	0.259*	0.259*	0.251*	0.218	0.209
HID4	0.403**	0.263*	0.236*	0.217	0.204	0.197
HID6	0.422**	0.329**	0.288*	0.250*	0.224	0.199
HID8	0.421**	0.316**	0.283*	0.238*	0.240	0.219
HID10	0.371**	0.301**	0.274*	0.250*	0.230*	0.211
HID12	0.339**	0.334**	0.271*	0.239*	0.238*	0.214
HID14	0.382**	0.321**	0.268*	0.247*	0.232*	0.208

Table 9.15. Correlations of SE and distance from mean target for the trend diagnosis in a constrained and combined model (N=75) (** and * indicates significant at α=0.01 and 0.05 respectively)

The results in table 9.15 showed that the combined model had a somewhat better distribution of errors than the corresponding separate model, but the distribution was more unfavourable than in the combined model of chapter 8. It was difficult to get an impression of how important the bias of this distribution was. Regressing the SE on distances from mean targets gave an R^2 of 0.04. Thus, distance from mean targets explained 4 % of the variance in SE. Furthermore, the regression coefficient was not significantly different from 0 at α=0.05 (t=1.83, d.f.=73). Consequently, the bias of the errors was considered undesirable but acceptable. Unlike the trend diagnosis models of chapter 8, bias was unambiguously reduced by learning in the constrained and combined models. Despite this, the bias was still undesirable and the constrained models seemed more regressive than the model of chapter 8.

We can conclude that reduction of the number of free parameters did not improve the undesirable distribution of errors. Despite the similarities in the probability distributions of the level and trend diagnosis variables, learning did not eliminate the unfavourable distribution of errors for the trend diagnosis in the same way as it did for the level diagnosis. However, White (1989) has pointed out in general that a biased error distribution is explainable. When extreme values of the input cues are likely to give extreme diagnoses and the extreme values are less likely to occur than the moderate values, the errors should be largest for the less likely diagnostic values:

> *These weights give small errors (on average) for values of X that are very likely to occur at the cost of larger errors (on average) for values of X that are unlikely to occur. (White, 1989, p. 98)*

Thus, the more surprising result of our simulations was that the level diagnosis models did not have the same biased distribution of errors as the trend diagnosis models had.

In regression analysis, a study of the residuals is often used to further evaluate model fit. Analysis of the SEs is somewhat similar to analysis of the residuals in regression analysis, and an overview of the outliers[1] is shown in table 9.16.

Model/diagnosis	Outliers					Number
32 inputs/level	AB	CA	D	M	R	5
32 inputs/trend	AB	AU	D			3
12 inputs/level	AB	CA	D	AP	R	5
12 inputs/trend	AB	AU	D		R	4

Table 9.16 Cases classified as outliers in models with 32 and 12 inputs

When studying table 9.16, we found that many of the outliers were common to the constrained models and to the models of chapter 8. The outliers AB and D were common to all models and diagnoses. Both these firms were judged more favourably by the subjects than by the models. The first hypothesis was that these firms had only been diagnosed by subjects with extreme response styles. A response style indicator had previously been computed[2], and was compared among the subjects in each composite judge committee. We found no composite judge committees where all subjects had extreme response styles. Thus, the response style explanation of the outliers could be excluded.

Correlations of the models' SEs with all financial cues of each firm showed that most of the model outliers were also outliers in the distribution of a handful of financial cues. The cues in question were: OSTREC, OOPREV, CHOSTREC, ACCPAY and APT. This finding led us to formulate a second hypothesis explaining the model outliers. It was suggested that the subjects had found something in the financial cues not "discovered" by the model that could justify their diagnoses, and explain the model outliers. The findings relevant to this hypothesis are summarised in table 9.17.

[1] A squared error is defined as outlier when SE > MSE + 2 σ_{SE}

[2] The response style indicator was made by computing the average positive or negative difference on all variables and firms a subject had diagnosed. The differences were computed as distance from the average response in the composite judge committees. Positive and negative differences indicated positive and negative response styles respectively.

Outlier	Diagnosis	Judgement	Special conditions	Implication
AB	Level Trend	Extreme-positive	Other operating revenues are high. Accounts payable period large due to large accounts payable relative to inventory.	Opposite direction.
AP	Level Trend	Extreme-positive	Accounts payable period in year 1 large due to large accounts payable relative to inventory. Other short term receivables large in year 1.	Opposite direction Same direction
AU	Trend	Extreme-negative	Other operating revenues are high and increasing in year 2. Positive change in other short term receivables. Large increase in accounts payable period.	Same direction Opposite direction Same direction
CA	Level	Extreme-positive	None.	None
D	Level Trend	Extreme-positive	Other operating revenues are high and increasing in year 2.	Opposite direction
M	Level	Extreme-negative	Accounts payable period large due to large accounts payable relative to inventory.	Same direction
R	Level Trend	Regress-positive	Large other short term receivables. Decrease in other short term receivables in year 2.	Same dir. Opposite direction

Table 9.17 Analysis of model outliers and financial cues of special relevance to the errors.

Table 9.17 shows the seven outliers of the models, the diagnosis it was the outlier for, how the subjects judged the firm relative to the model, the special conditions of the relevant financial cues, and implications these financial cues should have if the special conditions were taken into consideration by the subjects. If the subjects correctly had taken the special conditions into consideration, this was marked in the "implications" column by the term "same direction". If not, the term "opposite direction" was used. If the hypothesis of "proficient" subjects was correct, most outliers in table 9.17 should have been marked "same direction". However, the two markings were equally frequent. Thus, the hypothesis of "proficient" subjects was also excluded.

To judge how important the outlier errors were, they where excluded in the computation of the performance results of the best connectionist model. When eliminated, correlation of the MSE with distance from target dropped from 0.197 to 0.140. MSE dropped from 0.267 to 0.200. Thus, only a small improvement in the error distribution measure was obtained, whereas the cross validated average squared error was notably reduced. Despite the improvements gained in model performance from excluding the outliers, it was not clear if the outliers indicated subject or model "proficiency" in financial diagnosis.

9.2 Analysis of model representations

The constrained model minimised error under more constraints than the models of chapter 8. Theoretically, fewer free parameters should reduce the number of possible local minima. Consequently, the functional organisation should be more similar across connectionist models.

To study the representations, a set of comparable models was developed using the same principles as in chapter 8. Weights were initialised by a random number generator using the ranges described in chapter 6. In the constrained models, optimal performances were found for 30000 iterations. This was higher than the corresponding optimal iterations in chapter 8. When generating the models for representational analysis, we followed the same procedure as in chapter 8. Since all stimulus-response pairs were in the learning sample, convergence was likely to be found somewhat earlier than 30000 iterations. Since minimal MSE was found after a number of iterations roughly twice the number of chapter 8, learning was stopped after twice the number of iterations[1], and weights were saved. This procedure was repeated for 10 versions of the connectionist models with two, three and four hidden units, respectively.

The performance results of these models are shown in table 9.18.

Model	Average level MSE	St dev. of level MSE	Average trend MSE	St. dev. of trend MSE	Correlation of level with diff. from target	Correlation of trend with diff. from target	Runs with common hidden unit
HID2	0.111	0.004	0.207	0.007	-0.053	0.233	1
HID3	0.112	0.003	0.203	0.008	-0.051	0.223	4
HID4	0.112	0.002	0.200	0.004	-0.053	0.233	6

Table 9.18. Results of 10 constrained and combined models with full learning sample and randomised initial weights (figures are averages of the 10 runs)

The performance results were very similar to the results of chapter 8, but the number of common hidden units was reduced. This can be confirmed by visual inspection of the Hinton diagram in figure 9.2. The additional hidden units in the larger models did not improve performance significantly, but the standard deviation of the errors dropped with more hidden units.

To study the distribution of the weights, summary statistics of the weights of each connectionist model were calculated. The statistics are shown in table 9.19.

[1] In chapter 8, learning was stopped after 10000 iterations when the full learning sample was used. In this section it was stopped after 20000 iterations.

	HID2 (N=32)		HID3 (N=47)		HID4 (N=62)	
Version::						
A	Mean:-0.35 Skew:-0.14	Stdev: 1.38 Kurt:-1.24	Mean:-0.12 Skew:-0.12	Stdev:1.10 Kurt:-0.38	Mean:-0.04 Skew:-0.13	Stdev:0.99 Kurt:0.37
B	Mean:0.18 Skew:0.37	Stdev:1.46 Kurt:-0.93	Mean:-0.01 Skew:-0.87	Stdev:1.10 Kurt:0.42	Mean:-0.04 Skew:-0.52	Stdev:0.99 Kurt:0.09
C	Mean:0.25 Skew:-0.75	Stdev:1.37 Kurt:-0.75	Mean:-0.02 Skew:-0.20	Stdev:1.16 Kurt:-0.51	Mean:0.09 Skew:-0.14	Stdev:1.01 Kurt:-0.74
D	Mean:-0.16 Skew:-0.61	Stdev:1.35 Kurt:-0.54	Mean:0.06 Skew:0.37	Stdev:1.14 Kurt:-0.31	Mean:0.09 Skew:0.33	Stdev:1.02 Kurt:-0.08
E	Mean:-0.01 Skew:-0.89	Stdev:1.35 Kurt:-0.24	Mean:-0.10 Skew:-0.34	Stdev:1.13 Kurt:-0.51	Mean:-0.05 Skew:-0.36	Stdev:1.06 Kurt:-0.83
F	Mean:0.18 Skew:-0.45	Stdev:1.37 Kurt:-0.66	Mean:-0.26 Skew:-0.20	Stdev:1.21 Kurt:-0.76	Mean:-0.12 Skew:-0.62	Stdev:1.00 Kurt:0.67
G	Mean:-0.33 Skew:-0.05	Stdev:1.42 Kurt:-1.30	Mean:0.22 Skew:0.11	Stdev:1.23 Kurt:-0.73	Mean:-0.25 Skew:-0.11	Stdev:1.02 Kurt:-0.69
H	Mean:-0.19 Skew:0.58	Stdev:1.36 Kurt:-0.32	Mean:0.31 Skew:-0.58	Stdev:1.09 Kurt:0.18	Mean:-0.01 Skew:-0.34	Stdev:1.00 Kurt:-0.49
I	Mean:-0.24 Skew:-0.40	Stdev:1.42 Kurt:-0.90	Mean:-0.05 Skew:-0.63	Stdev:1.11 Kurt:-0.29	Mean:-0.11 Skew:-0.29	Stdev:1.02 Kurt:-0.44
J	Mean:-0.34 Skew:-0.07	Stdev:1.43 Kurt:-1.27	Mean:-0.23 Skew:0.04	Stdev:1.15 Kurt:-0.51	Mean:0.14 Skew:0.31	Stdev:1.00 Kurt:0.49

Table 9.19. Statistics illustrating the differences in representations between the 10 versions of each constrained and combined model[1]

From table 9.19, the similarities of the different versions of the models could be studied. There were some variations in skewness and kurtosis, and in the versions with two hidden units, all models had a less peaked weight distribution than the normal distribution. The versions with three and four hidden units had weight distributions close to the normal distribution. We recall that the weight distributions of the models in chapter 8 were generally more peaked than the normal distribution. The difference between the two sets of weight distributions was a consequence of the reduction in the number of input units. The large amount of weights around zero was now eliminated, resulting in a larger proportion of the weights being "active" in the diagnosis. A preliminary conclusion was that the representations generally were more local than in chapter 8.

As in chapter 8, a cluster analysis of absolute values of the network weights coming into the hidden layer was used to detect local units and to investigate the clusters of functionally equivalent hidden units. A dendrogram from the cluster analysis of the weights is shown in figure 9.1.

---

[1] Skewness and kurtosis are standardised.

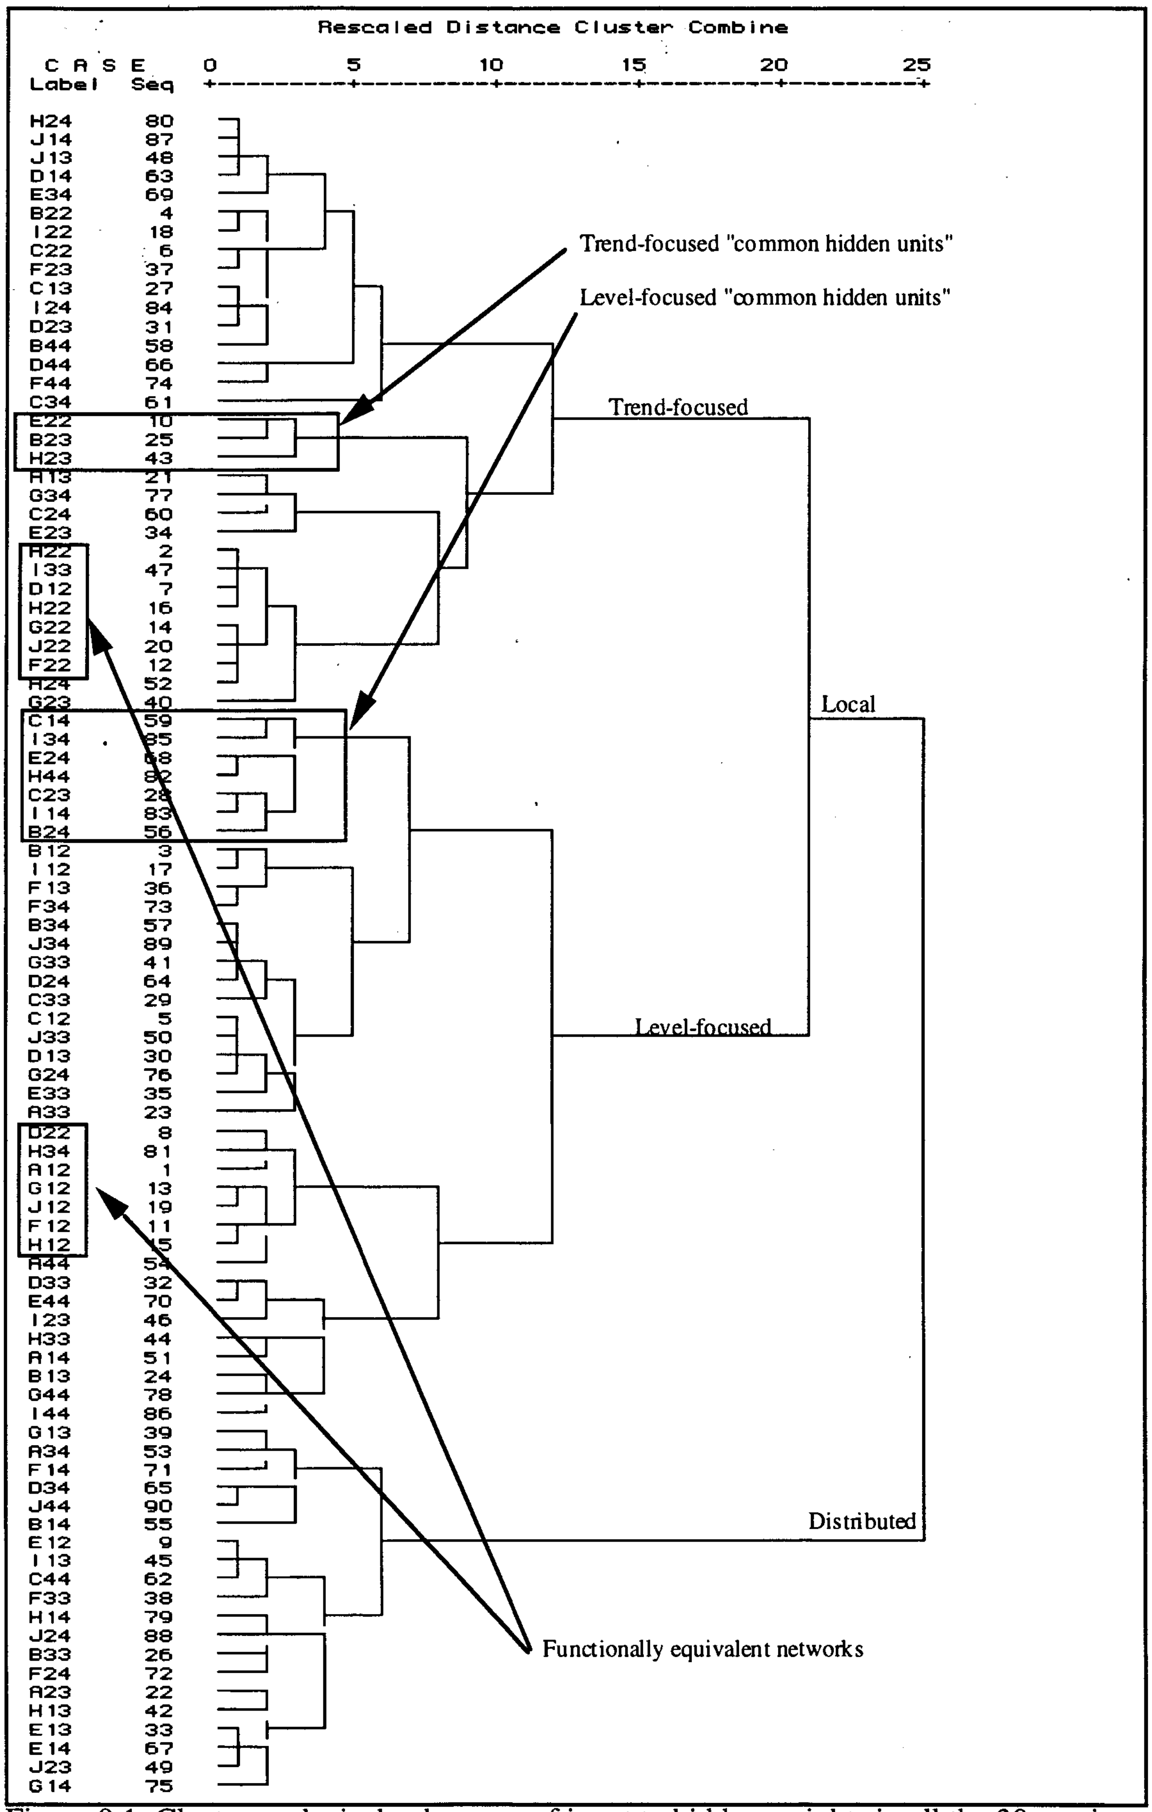

Figure 9.1. Cluster analysis dendrogram of input to hidden weights in all the 30 versions of the constrained and combined connectionist model with 2, 3 and 4 hidden units

In the models of chapter 8, we found that the *local* units were *common* hidden units. The first split of the dendrogram in figure 8.2 of chapter 8 was between these local common hidden units and the rest of the units. A similar split between local and distributed hidden units was found in figure 9.1. Here, the *largest* cluster consisted of the local hidden units.

Even though all common hidden units were found in the local cluster, they were placed in two different subclusters. This indicated that the common hidden units had a different organisation than in the models of chapter 8. They were either trend- or level-oriented. Consequently, the typical pattern of strong common hidden units in all models found in chapter 8 was broken for the constrained models. In all the constrained models, the functional organisation of the task was done in a more traditional manner. The hidden units seemed to create a task partitioning among themselves that was more output-directed or output-oriented than in chapter 8. The representation of a model with only two hidden units could be interpreted as if the hidden units had a functional organisation in which each of the units specialised on each of the output responses. Furthermore, the hidden units seemed to combine the inputs in a way that made their outputs more orthogonal to each other than the models of chapter 8. Testing this hypothesis on the small connectionist models, the average correlation between the output of the two hidden units was calculated. For the connectionist models in chapter 8, this correlation was 0.411, and in the constrained models, it was 0.312. Consequently, the two hidden units worked somewhat more independently. In addition, models with few weights now organised their representation to reach local solutions close to each other. An example of this can be found in figure 9.1, where a group of functionally equivalent models is shown.

However, these suggestions were only preliminary, and the procedure of analysis developed in chapter 8 was applied. In figure 9.2, a Hinton diagram of the weights between hidden units and output units is shown.

From the Hinton diagrams in figure 9.2, we see how the hidden units of models A2, D2, F2, G2, H2 and J2, specialised on each of the two responses and split the diagnostic task among the two hidden units. We also noticed the relatively small amount of units similar to the common hidden units of chapter 8. Such common hidden units were primarily found in the larger models with three and four hidden units.

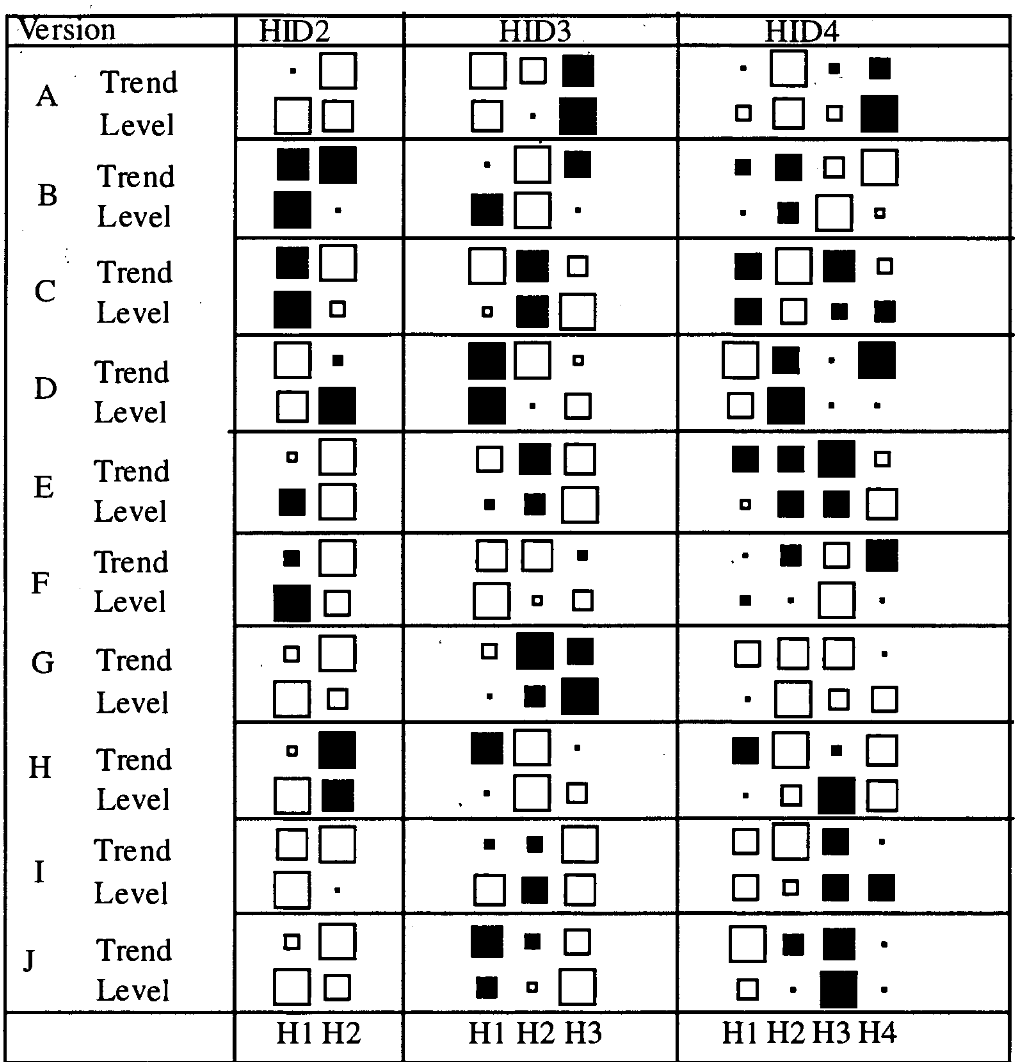

Figure 9.2 Hinton diagram of the weights between hidden and output units for the 10 versions of the constrained and combined models with 2, 3 and 4 hidden units respectively

In chapter 8, we showed how similarities in the weight patterns of the connections between the hidden and output units were reflected in similarities in the connections between hidden units and input units. The cluster analysis dendrogram in chapter 8 was also used to detect groups of *functionally equivalent* hidden units. In the cluster analysis dendrogram of figure 9.1 it was easy to recognise different clusters of functionally similar hidden units.

When compared to the Hinton diagram of figure 9.2, it was obvious that the first split of the local hidden units cluster in figure 9.1 was between units focusing on trend and units focusing on level diagnosis. In figure 9.1, all hidden units were analysed simultaneously. To clarify the analysis, each set of models could be analysed separately. The separate analysis also made it possible to combine dendrograms and Hinton diagrams for illustration. Thus, the 10 versions of the models with two, three and four hidden units were cluster analysed separately. A cluster analysis dendrogram of the input weights to the 20 hidden units in the 10 models with two hidden units is shown in figure 9.3. To further improve the analysis, we did not use absolute values, but developed a method that controlled for the sign differences and the biases in each model. The weights used in the cluster analysis of figure 9.3 have been adjusted for sign differences and biases[1], but the Hinton diagrams are shown with the original weight values.

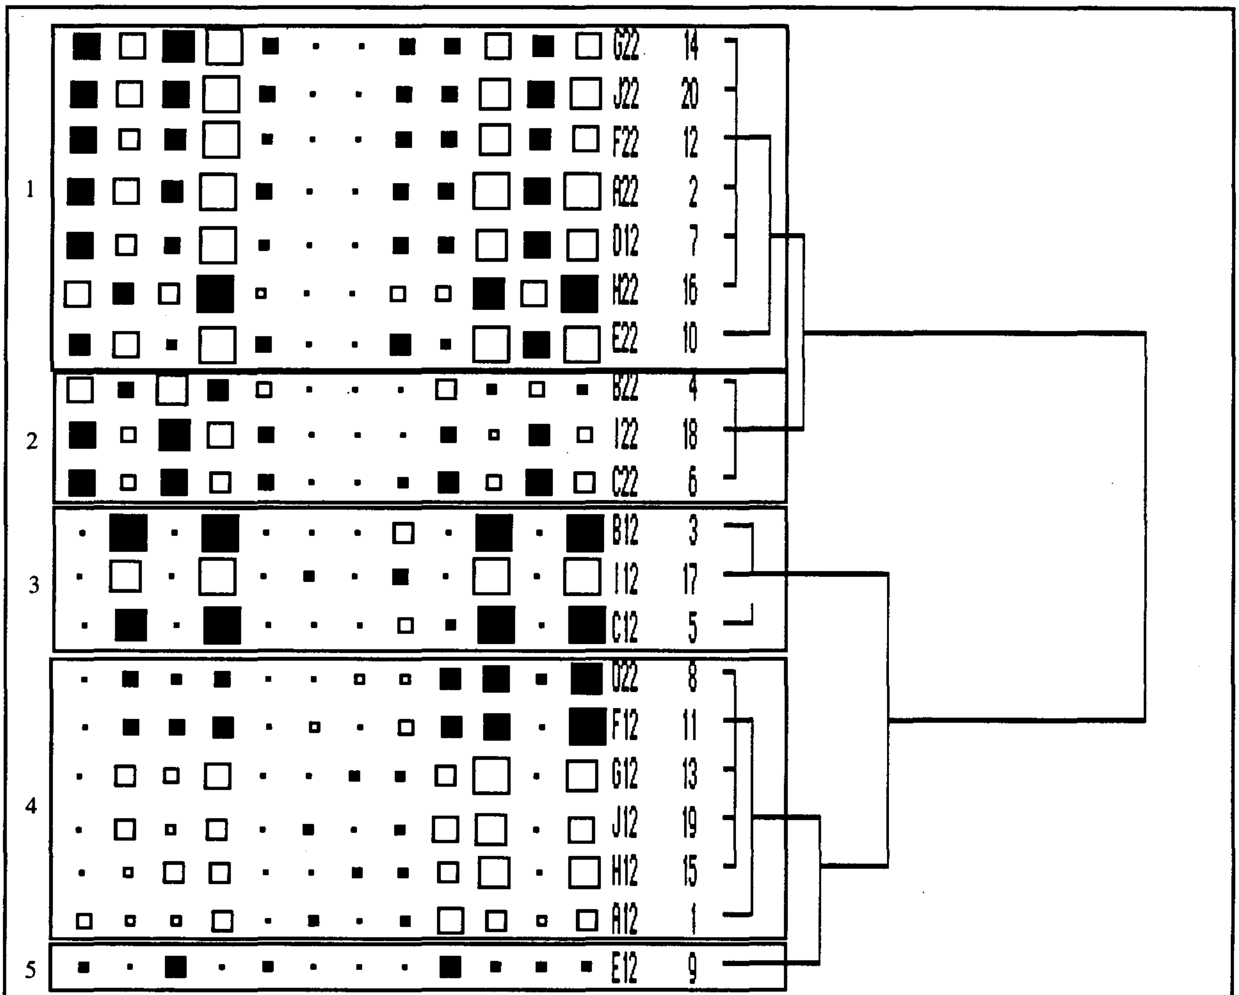

Figure 9.3 Cluster analysis dendrogram and Hinton diagram of the constrained and combined models with two hidden units

To compare the hidden units, the Hinton diagrams show the weight pattern of each connection between input and hidden units. The weights are placed in the following order: OPMA1, OPMA2, ROI1, ROI2, ROE1, ROE2, AIR1, AIR2, CURR1, CURR2, BER1 and BER2. When cluster analysis dendrograms and Hinton diagrams are combined in this way, visual inspection of the weight pattern of subclusters is eased. Five clusters are marked in figure 9.3.

---

[1] The adjustment was performed by computing the new weight as -[w_{ij} - (bias/12)] for inhibitory hidden units and w_{ij} - (bias/12) for excitatory hidden units.

We recognised the groups of functionally equivalent models described above and marked in figure 9.1. First we recognised the units of the functionally equivalent models A2, D2, F2, G2, H2 and J2 in clusters 1 and 4. Next, we found that the models B2, C2 and I2 that were functionally equivalent in figure 9.1, had units placed in clusters 2 and 3 only. Finally, we found the only model with the common hidden unit. The common hidden unit was placed in cluster 1, even though it differed from the other units in this cluster. The distributed hidden unit E12 was placed in the separate cluster 5.

We could now interpret the hidden units in each of the five clusters. In clusters 1 and 2, we found units most strongly activating the trend diagnosis unit. This can be confirmed by studying figure 9.2. The typical positive/negative weight pattern of a trend unit was also found in these units. Units of clusters 3 and 4 most strongly activated the level diagnosis units, and could be interpreted as level-oriented hidden units. The two subclusters in each of the two clusters differed in the way they specialised on the trend and level diagnosis respectively. Clusters 1 and 3 had more "common" trend- and level-oriented hidden units. The model with the most common trend-oriented hidden unit was E2, but even this model could be characterised similarly to the models belonging to cluster 1. Consequently, the models organised their hidden units in one of two functionally different ways. The first group organised their representations with a trend-oriented hidden unit, also taking care of the common parts of both diagnoses, and used the second hidden unit to *specialise* on level diagnosis adjustment. The opposite was done by the three resulting models. They had a level-oriented hidden unit also taking care of the common parts, and used the second hidden unit to specialise on trend adjustment. These two solutions could be characterised as *similar local solutions* to the mapping problem solved by the connectionist models, and no indication was found that one solution was better than the other.

In figure 9.4, a similar cluster analysis dendrogram of the input weights combined with a Hinton diagram is shown for the 10 versions of the model with three hidden units.

The first split of the dendrogram was now between the local and distributed hidden units. The models with two hidden units had only local units, but with three hidden units, distributed units seemed to develop. The next split was between the local hidden units focusing on trend and the units focusing on level diagnosis.

The next split of the clusters was in functionally different hidden units within the level- and trend-oriented hidden units clusters, and within the distributed hidden units cluster. The resulting clusters have been marked and are numbered in figure 9.4. Cluster 1 consisted of units with the typical positive/negative sign pattern of a trend-oriented hidden unit. However, these units differed from the units in cluster 2 by their relatively large weights from input units representing cues from the first of the two consecutive years. In figure 9.2, we see how weights from these units almost exclusively activated the trend diagnosis units. We characterised these units as *"exclusive trend"* oriented hidden units.

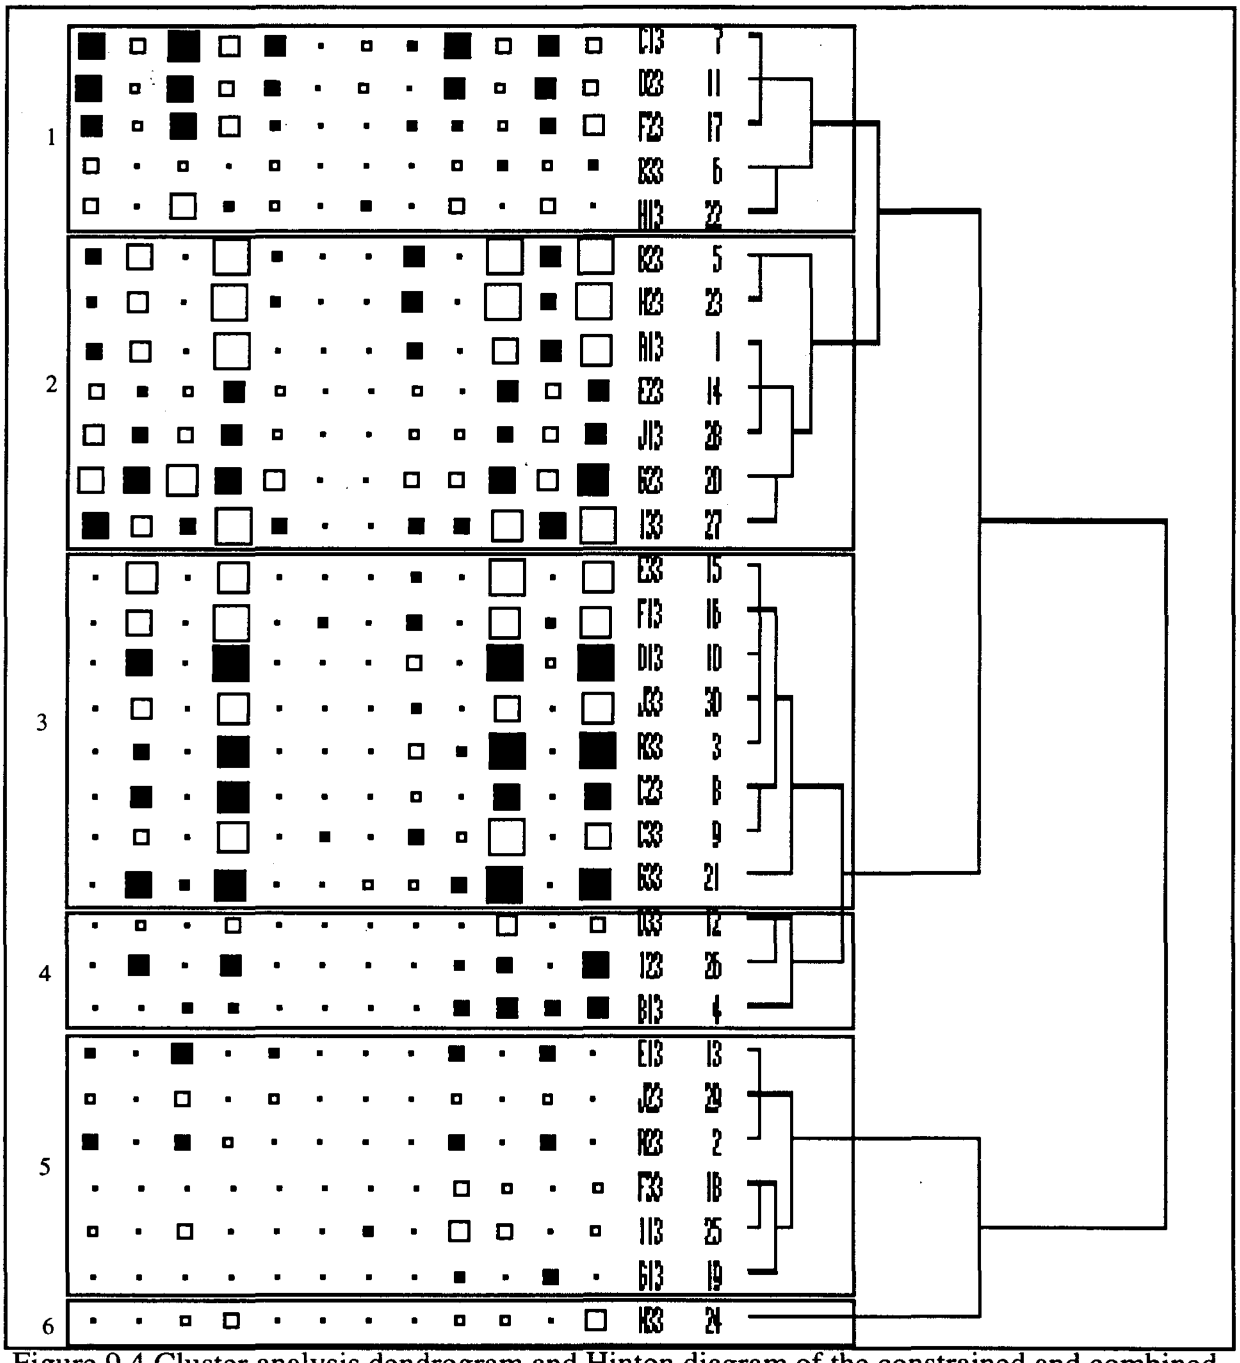

Figure 9.4 Cluster analysis dendrogram and Hinton diagram of the constrained and combined models with three hidden units

In cluster 2, we found the same sign pattern as in cluster 1, but weights coming from units representing cues from the second of the two consecutive years were much larger than those coming from units representing the first year cues. These units represented a more complex concept of common aspects of level and trend diagnoses, but their main focus was on trend. In the units of clusters 3 and 4, the positive/negative sign pattern of a trend-oriented hidden unit could not be found. These units were typical level-oriented units. The units in cluster 4 differed from the units in cluster 3 by their smaller absolute values. Small weight values can be compensated by larger weights in the layer between hidden and output units, and could not be used to discriminate the units of cluster 4 from the units in cluster 3 functionally.

Clusters 5 and 6 consisted of distributed hidden units. As opposed to the distributed hidden units in chapter 8, there was a typical weight pattern in some of these units. In cluster 5, we found a pattern close to the pattern of cluster 1, but the positive/negative sign pattern was weak. From the sign pattern of the connections between the hidden and the output units, found in figure 9.2, we preliminarily interpreted these units as "difference" hidden units. These units simultaneously inhibited the trend and excited the level diagnoses units, or the opposite. Because of their large values on weights coming from input units representing cues of the first consecutive year, we were tempted to characterise these units as "yesterday's situation" detectors. It seemed that the adjustments of a common diagnosis performed by these units, were done by focusing on the first of the two consecutive years. Thus, the values from the cues of the first year were used to represent the exception to the rule that level and trend were positively correlated. With the use of the combined cluster analysis and Hinton diagram, we could now interpret how the exceptions to this rule were detected. They were simply detected by focusing particularly on a set of cues representing several diagnostic concepts of the first of the two consecutive years. This corresponded well to our interpretation in chapter 8, that the trend diagnosis should be adjusted downwards from what was predicted by a common hidden unit when the cue values of the first of the two consecutive years were exceptionally large.

Cluster 6 consisted of only one unit with a weight pattern close to the units of cluster 4, and was interpreted as a "miniature" level-oriented unit. The only reason why this unit was not placed in cluster 4 was its generally lower absolute weight values.

Consequently, the functionality typical for the models with two hidden units was also found for models with three hidden units, but the use of the "yesterday's situation" detectors adjusting for the exceptions to the rule that level and trend diagnoses were positively correlated, was not found in the smaller models. The performance of the larger models was not significantly better than that of the models with two hidden units, suggesting that the additional hidden units of the larger models did not improve performance. Despite this, we showed how the combination of Hinton diagrams of the connections between hidden and output units and the diagrams of figure 9.4, could be used to identify units representing rule exceptions and how this exception was detected in the stimulus material.

A similar procedure was followed for the models with four hidden units. In figure 9.5, a cluster analysis dendrogram of the input weights of all hidden units in the models with four hidden units is combined with a Hinton diagram of the connection weights.

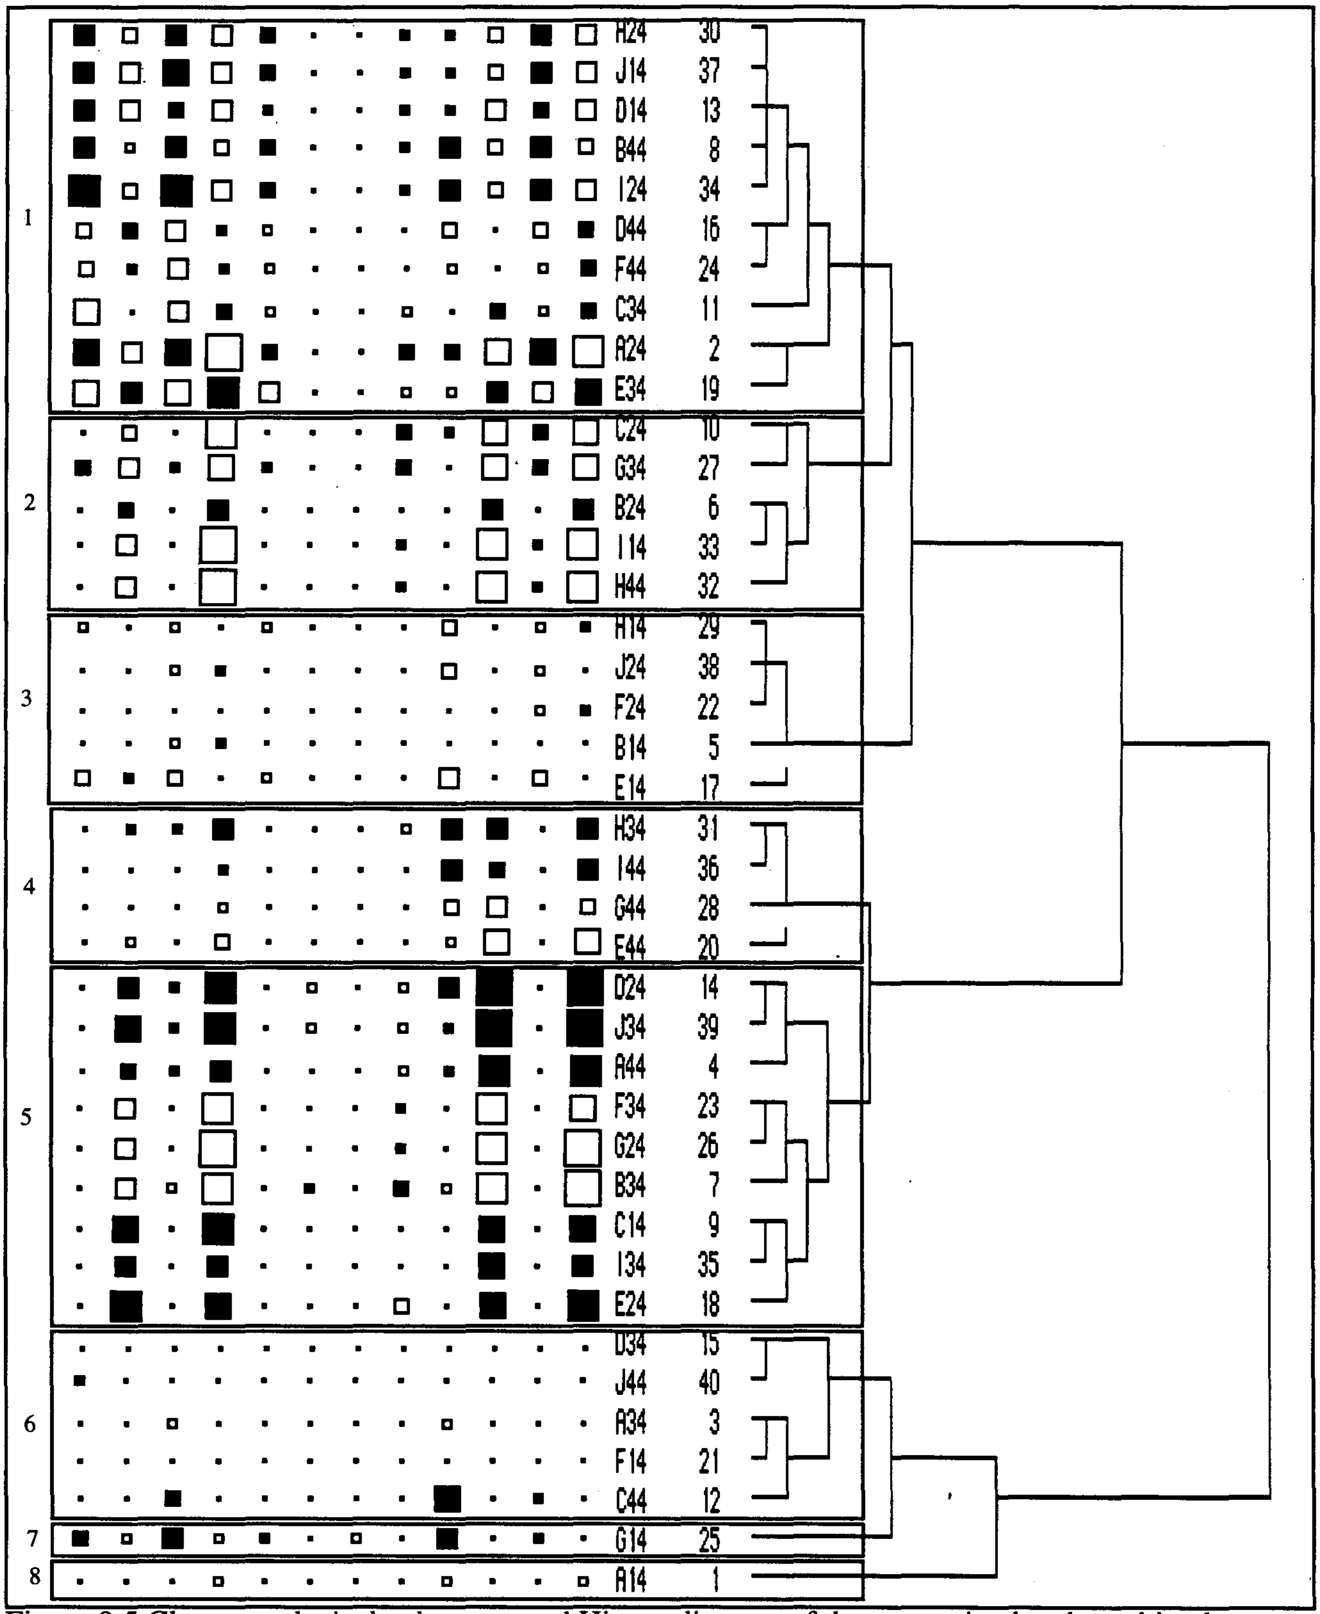

Figure 9.5 Cluster analysis dendrogram and Hinton diagram of the constrained and combined models with four hidden units

From figure 9.5, the functionality found in the models with four hidden units could be analysed. We recognised a functionality very similar to the one found in the models with two and three hidden units. The first split was found between local and distributed hidden units. The local hidden units were split between level- and trend-oriented hidden units. A subcluster in the trend-oriented hidden units cluster consisted of highly distributed hidden units. Eight clusters could be identified when setting the cluster distance at an interpretable level. All models except G4 had a unit in cluster 1 and all models except H4 had a unit in cluster 5. The two clusters 1 and 5 contained the trend- and level-oriented hidden units respectively. Model G4 had a unit in cluster 2 and H4 had a unit in cluster 4. These clusters were also trend- and level-oriented hidden unit clusters, respectively. The rest of the hidden units were distributed units found in clusters 3, 6, 7 and 8. Subclusters 6, 7 and 8 were placed in a separate cluster of distributed hidden units, whereas cluster 3 was placed among the trend-oriented units' clusters. The units of cluster 3 were "miniature" trend-oriented units, but nothing indicated that these units significantly improved the performance of the model. Some "yesterday's situation" indicators were found in the clusters 6, 7 and 8. Except for these, the units in these clusters seemed to be quite redundant.

Consequently, most models with three and four hidden units implemented a functionality close to the one implemented by the models with two hidden units. Additional hidden units did not improve the performance significantly. The *primary* functional organisation of the constrained connectionist models can thus be understood by investigating the smaller models.

If the connectionist models implemented a functionality as the one described above, one may ask why the performance was superior to the linear models. Traditionally, connectionist models perform in a more superior manner because they can detect and utilise nonlinearities in the mappings (Smith, 1993). Nonlinearities may be implemented by hidden units and by output units, or by combining the two groups of units. If both the mapping of inputs to hidden unit outputs and the hidden unit outputs to output unit outputs are linear, the mapping of inputs to outputs can be rewritten on linear form. Since our connectionist models with hidden units performed better than the linear benchmarks and the connectionist models without hidden units, nonlinear mappings were presumedly implemented in the models. By performing linear regressions using the hidden unit outputs as dependent variables, and the input cues as independent variables, the nonlinearities of the hidden unit mappings could be detected as deviations from predicted response. The same procedure was followed by using hidden unit outputs as independent variables, and output unit outputs as dependent variables.

Only the small models were analysed for nonlinearities. We have previously shown that the larger connectionist models *mainly* implemented the functionality of the models with only

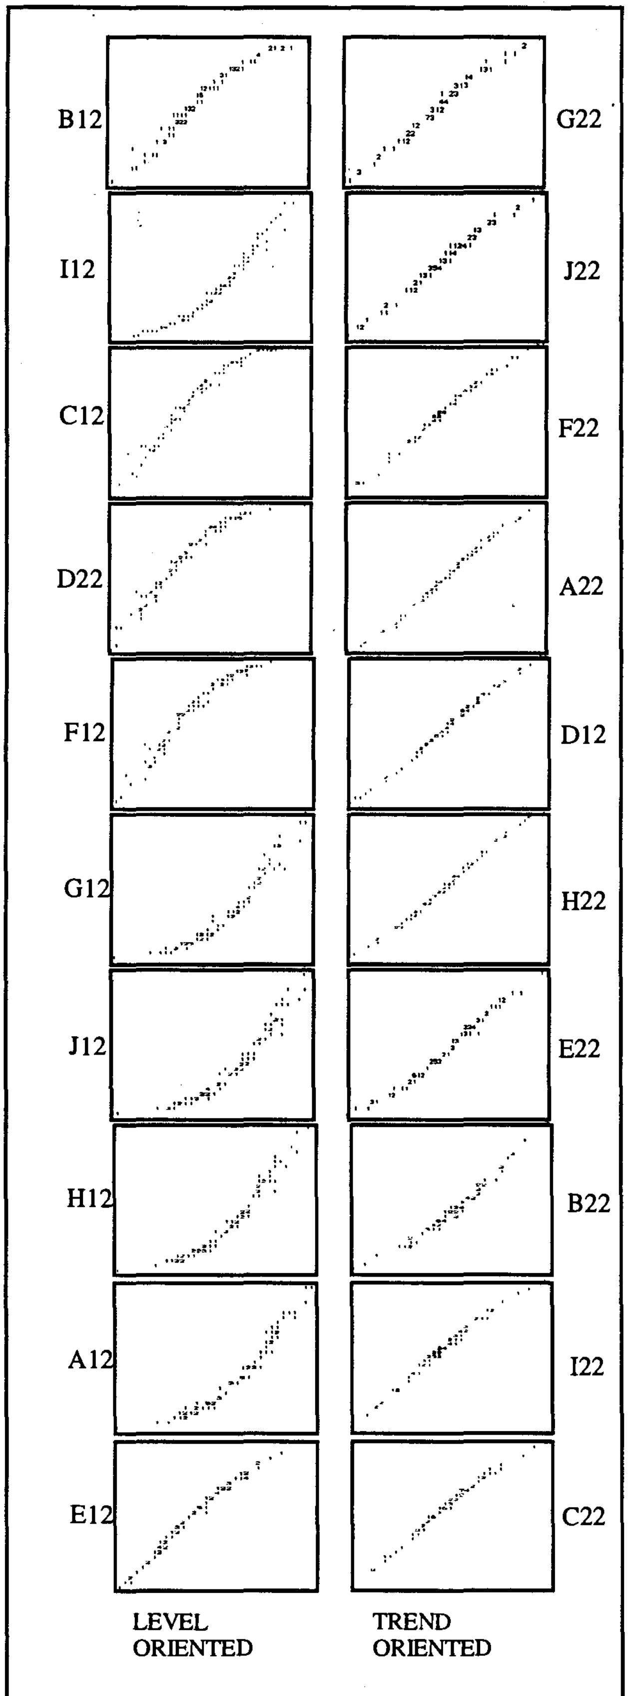

Figure 9.6 Plot of hidden unit outputs (y-axis) vs. linearly predicted values (x-axis)

two hidden units, and that the additional units of the larger models did not improve performance. A plot of actual output values of the hidden units versus the best linearly predicted value of the outputs using input cues as independent variables, is shown in figure 9.6 for each hidden unit.

In the figure, predicted hidden unit outputs are shown on the x-axis and actual hidden unit responses are shown on the y-axis. In the left column of figure 9.6, we have shown the units in the level-oriented cluster, and in the right column we have shown the trend-oriented hidden units from the dendrogram of figure 9.3. We found that the trend-oriented units had an output that could be approximated by a linear function of the inputs. However, the level-oriented units had an output that was highly nonlinear.

The functional form varied with excitatory and inhibitory units. Only model E2 deviated from this pattern. Actually, the response of both hidden units in E2 was somewhat nonlinear. The nonlinearities seemed to be present both in the common level-oriented units B2, I2 and C2, and in the more specialised level-oriented units in the lower part of the cluster. Since the trend-oriented hidden units had approximately linear outputs both for common trend-oriented hidden units and for specialised trend-oriented hidden units, we can conclude that it was not the common part of the level diagnosis that was modelled by nonlinearity in the connectionist models. Rather, it seemed to be the parts of the representation that were *uniquely* associated with the level diagnoses that was modelled by the nonlinearities of the hidden units. The nonlinearity was used to *moderate* the level diagnosis of the cases with high cue values and to *amplify* the level diagnosis of stimuli having cues in the medium high range.

Nonlinearities may also be introduced in the layer between hidden and output units. A similar procedure as above was followed to investigate the nonlinearities introduced in this layer. The results showed only approximately linear mappings between the hidden unit and output unit response. Consequently, the hidden unit layer was used to form a nonlinear indicator of the unique aspects of the level diagnosis and an approximately linear indicator of the common and trend-oriented aspects. These indicators were linearly combined to form the level and trend *diagnoses.*

Hidden unit outputs resemble factor scores in traditional models[1]. In the connectionist models one of these "factors" were nonlinear while the other was approximately linear. The two "factors" were combined in two classes of models. In the models with the common part of the level and trend diagnoses treated by the level-oriented hidden unit, the trend diagnosis was formed by approximately linearly combining the common part and the uniquely trend-oriented hidden unit output. Level diagnosis was modelled solely by the level-oriented hidden unit, representing both the common part and the uniquely level-oriented part of the level diagnosis. In the models with the common part treated by the trend-oriented hidden units, a level diagnosis was formed by combing the common part and the uniquely level-oriented part. Trend diagnosis was modelled solely by the trend-oriented hidden unit, representing both the common part and the uniquely trend-oriented part of the trend diagnosis.

A second way to explain the common and unique parts of the different hidden unit representations, is by considering the level and trend diagnoses as being composed of two parts each. One part is common, while the unique part can be considered as the level *independent* part of the trend concept. Similarly, we can think of a trend *independent* part of the level concept. These common and independent parts of the concepts existed because the two diagnoses level and trend, were highly *correlated* in empirical cases. As such, the *empirical* measures of the level and trend diagnoses were correlated, while they may be more "independent" or separated as *theoretical* concepts. The hidden units were used to identify the relevant inputs to such theoretical concepts of "level" and "trend".

---

[1] Such as in the benchmarks A and B of chapter 6.

If the connectionist model was expanded by adding a third hidden unit, one may expect the task to be partitioned in the uniquely level-oriented, the uniquely trend-oriented and the common parts. Generally, this did not seem to be the case. The functional organisation of the task performed by the models with two hidden units was retained, and the additional hidden unit did not improve performance significantly. Some of the additional hidden units were very distributed, and the rest of them seemed to be "miniatures" of similar, but more local, versions. However, some exceptions corresponding more to the "three dimensional" organisation of the representation were found. Upon visual inspection of figure 9.2, the closest functional organisation to the one proposed above was found in model C3[1]. The first hidden unit of this model seemed to be an inhibitory, specialised trend-oriented hidden unit. The second seemed to be a hidden unit representing the common aspects of the level and trend diagnoses, and the third seemed to be an inhibitory, specialised level-oriented hidden unit. A procedure similar to the one used in figure 9.6, was applied to the hidden unit outputs in model C3. The result is shown in figure 9.7.

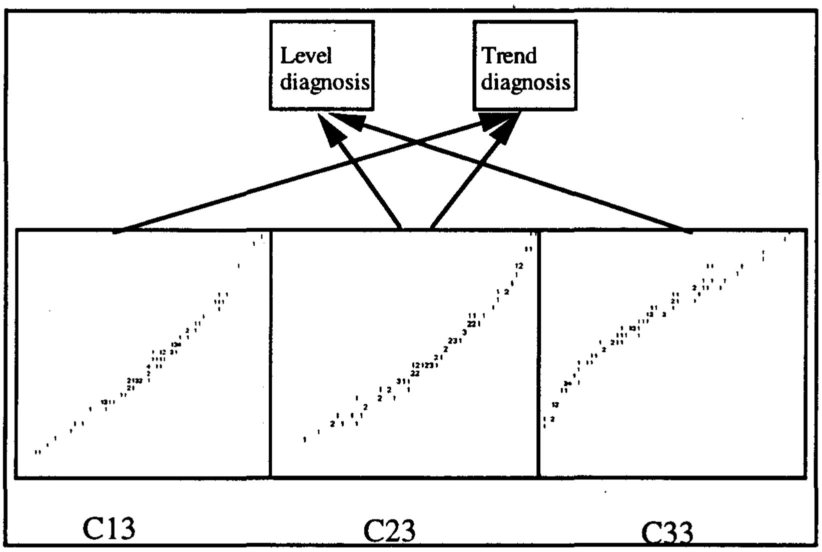

Figure 9.7 Linearities and nonlinearities in model C3

Hidden units C23 and C33 had a very similar incoming weight pattern. However, the outputs of the two units were combined very differently to form the two diagnoses. While hidden unit C33 almost only connected to the level diagnosis unit, hidden unit C23 also connected to the trend diagnosis unit. Despite their similar incoming weight pattern, the output function of the two units transformed the incoming signals differently. The specialised level-oriented hidden unit C33 was highly nonlinear. Consequently, hidden units with similar incoming weight patterns may be used differently to form the responses of a connectionist model.

To compare the "factors" formed by the hidden units, traditional principal components analysis[2] of the input variables revealed four factors with an eigenvalue greater than 1.00.

[1] The models B3, D3 and H3 also had a similar functional organisation.

[2] The principal components analysis was done on the 12 independent variables of the constrained connectionist models.

These factors represented 81.2 % of the variance in the original twelve-variable set. The rotated factor loadings of these factors using varimax rotation are shown in table 9.20.

In table 9.20, the highest factor loadings have been marked. We interpreted the factors as the first representing a combined indication of "liquidity" and "leverage". The second factor represented the "profitability" level of the second of the two consecutive years. The third factor represented "financing", and the last factor represented the "profitability" level of the first of the two consecutive years.

Variable	FACTOR 1	FACTOR 2	FACTOR 3	FACTOR 4
OPMA1	0.068	0.225	-0.189	0.903
OPMA2	0.192	0.806	-0.166	0.353
ROI1	0.004	0.139	0.083	0.939
ROI2	0.115	0.864	0.176	0.238
ROE1	-0.260	0.422	0.213	0.602
ROE2	-0.086	0.836	-0.026	0.055
AIR1	-0.056	0.066	0.894	0.060
AIR2	0.038	-0.017	0.877	-0.001
CURR1	0.921	-0.012	0.048	-0.033
CURR2	0.847	0.059	0.136	-0.122
BER1	0.716	-0.055	-0.548	0.161
BER2	0.711	0.199	-0.527	0.102

Table 9.20 Factor loadings from principal components analysis of the input variables with varimax rotation

The factors were recognised as belonging to different diagnostic areas, and no factor loadings corresponding to the level- and trend-oriented hidden units of the connectionist models were found. The pattern of the rotated factors roughly corresponded to the *hypothesised* functional task partitioning. Regressing trend and level diagnosis on the factor scores of the four factors gave an MSE of 0.214 for the level diagnosis and 0.380 for the trend diagnosis, when the cross validation procedure was applied. This was a significant loss in performance when compared to the connectionist models, but it represented no considerable loss in performance when compared to the best traditional linear regression or benchmark models. Consequently, the difference between the "factors" developed by the connectionist models and the traditional factors was that the "factors" of the connectionist models were *functionally related to the task the model performed.* This implied that the "factors" of the connectionist models were not necessarily functional in separating the input variables into orthogonal dimensions, but they certainly represented a functional separation of the input variables that had vital importance to the task the connectionist model was set up to minimise the error of. This separation of the input variables represented the abstracted stimulus dimensions used by the hidden units of the connectionist models.

To summarise in terms used in classification research, three functionally different *derived stimulus dimensions* were represented by the hidden units of the constrained models. The "common" dimension was common to both level and trend diagnosis, whereas the "specialised level" dimension was only relevant to level diagnosis. The "specialised trend" dimension was only relevant to trend diagnosis. The "common" dimension was an approximately linear function of the values of the *original* stimulus dimensions of the task. The "specialised level" dimension was a nonlinear function of the same values of the original stimulus dimensions. The "specialised trend" dimension was an approximately linear function of the *change* in the original stimulus dimension values. For all these abstracted stimulus dimensions, no obvious selective attention (Kruschke, 1992, 1993a, 1993b) was paid to stimulus dimensions presumed to represent different diagnostic concepts, such as "profitability", "liquidity", "financing", or "leverage". In some of the larger connectionist models, one hidden unit was allocated to represent each of the three abstracted stimulus dimensions. In the smaller connectionist models, a combination of the "common" stimulus dimension and one of the "specialised trend" or "specialised level" dimensions, was implemented by one hidden unit alone.

To further interpret the behaviour of the models in classification research terminology, a cluster analysis of hidden unit outputs during processing, so called "in vivo" clustering (Hanson & Burr, 1990), was used. The outputs of the C3 model was recorded while all stimulus-response patterns were presented to the model. Next, the recorded outputs were cluster analysed to investigate how different stimulus-response patterns were placed in different clusters of the hidden unit output space. A dendrogram illustrating the clusters is shown in figure 9.8.

Figure 9.8 illustrates the clusters formed at an interpretable cluster distance. The diagnosis values shown on the left are the average level and trend diagnoses of each cluster, whereas the case indicator shows the number of firms found in each cluster. The cluster distance is illustrated in an ordinary dendrogram.

As can be seen in figure 9.8, the first split was between low and moderate to high level diagnosed firms. The low level diagnosis cluster was divided into three subclusters. One cluster consisted of extremely low valued firms, while the other clusters were split between promising and less promising firms. The average trend diagnosis of the promising firms was as high as 2.6.

The moderate and high level diagnosed firms were divided into two main subclusters depending on their trend diagnosis value. In the first main subcluster, average trend diagnosis value was only 2.8, whereas it was 3.9 in the second. In the high trend diagnosis subcluster,

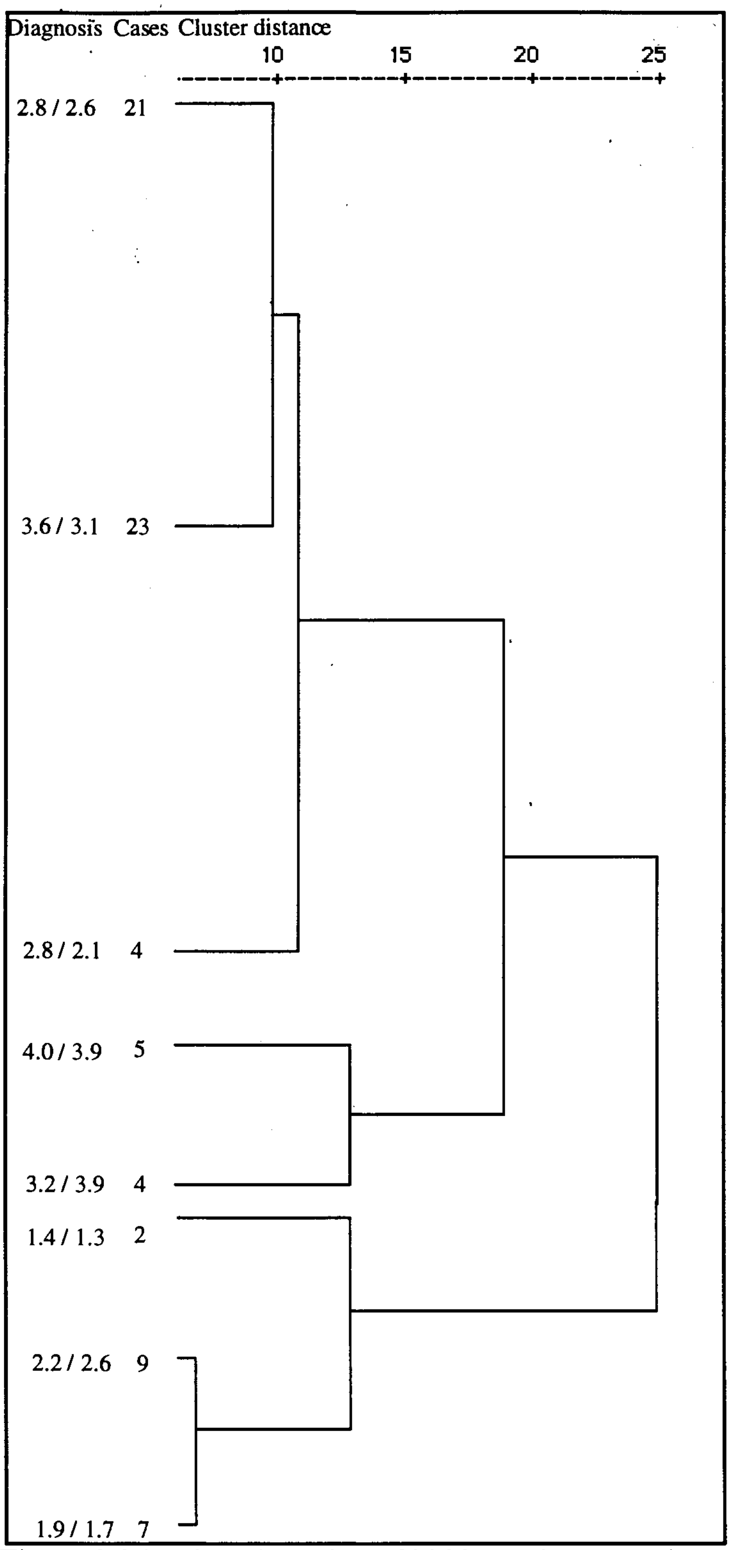

Figure 9.8 In vivo clustering of hidden unit outputs in the C3 model

new subclusters were formed depending on the level diagnosis value. In the low trend diagnosis subcluster, a further split of both level and trend diagnosis values was found.

The most important conclusions that could be drawn from this analysis was that the hidden unit output space was diagnosis oriented. If, for example, similar average diagnosis values had been found in each subcluster, we could infer that the hidden units detected stimulus dimensions or features in the financial cues that *later* could be used to form final diagnoses. Then, firms with similar diagnoses would be found in different subclusters. However, they would have been placed in different subclusters for different *reasons* found in the structure of their financial cues. This was, however, not the case. Consequently, it was very unlikely that a "subclass detector" or a diagnosis independent "feature detector" representation of the kind illustrated in figure 4.2 was used[1].

[1] Average values were computed for each cluster of figure 9.8 for the subjects' judgements of the diagnostic areas "profitability", "financing", "liquidity" and "leverage" also, showing pattern very similar to the average values illustrated in figure 9.8.

Furthermore, there seemed to be no explicit prototype representations *directly* implemented by hidden units in the model. However, looking at the cluster diagram of figure 9.8, cluster centres represented prototypical combinations of level and trend diagnoses. The "prototypes" did not consist of diagnosis independent subclasses, but rather of level and trend diagnosis-related prototypes, such as "bankruptcy", "bad, but promising", or "good, but with alarming trend", firms. Thus, prototype *interpretation* of the functional organisation of the task was possible, even though no *explicit* prototype representation took place in the hidden units. In this way, prototype-based explanation of the functional organisation was similar to a heuristics- or rule-based interpretation. Even though no rules were directly implemented in single hidden units, model behaviour could be described as if it followed rules.

Correspondingly, even though no hidden units directly implemented prototypes or measured similarities to prototypes, a functional organisation of the task could be found in the dendrogram of figure 9.8 that could be interpreted as if similarity to prototypes was used in the model. Thus, both prototype and rule-based *interpretations* of connectionist model behaviour seemed relevant.

9.3 Constraining complexity by cue importance

In the introduction to this chapter, we pointed out that the most successful constrained model was developed by cue elimination when some sensitivity based measure of cue importance was used. This procedure was followed in section 9.1. Subjects' cue importance indicators can also be used for cue selection or elimination. In chapter 5, we used this approach to focus on the ratio section of the stimulus material in order to develop the independent variables of the models. Using the cue importance measures, we could either eliminate cues from the ratio cue list of chapter 5, or we could use the importance indicator to select more freely from the total list of cues. While the first approach led to elimination of cues also eliminated by the sensitivity based method of section 9.1[1], the second approach was used in this section. In table 9.22, a summary of the cues with the highest importance indicator[2] of each diagnostic area is shown[3].

---

[1] The only exception to this rule was the selection of AIR as an important indicator. This cue had a low importance indicator, but it had high sensitivity and was included in the models of section 9.1 and 9.2.
[2] See chapter 5 for explanation of the importance indicator.
[3] The table is an exhibit of table 5.3.

Diagnostic area	Priority 1 (cumm.prop.)	Priority 2 (cumm.prop.)	Priority 3 (cumm.prop.)	Priority 4 (cumm.prop.)	Priority 5 (cumm.prop.)
Profitability	ROI (19.6)	PROMARG (33.3)	ROE (42.7)	OPMARG (50.4)	OPROF (58.0)
Financing	LTINV (19.0)	LTL (29.6)	BER (38.8)	STL (46.8)	ICOV (54.3)
Liquidity	CURR (19.6)	ACID (37.7)	CHLIK (46.7)	CASH (52.9)	APT (58.7)
Leverage	BER (50.1)	EQUITY (61.6)	ICOV (66.2)	ROE (70.7)	ROI (73.5)

Table 9.22 Five most important cues of each diagnostic area (cumulative proportion of cues indicated in parentheses)

From table 9.22, we found that the five selected indicators in each diagnostic area represented over 50 % of all indicated cues. For "leverage", BER completely dominated the other cues. Most of the cues were from the ratio section of the stimulus material. For LTL and STL, most subjects had indicated a relationship between the two indicators. This was interpreted as if the subjects had missed a ratio computed as LTL/STL or STL/LTL. Some of the cues were multiple cues or cues indicating closely related concepts. The high correlation of such cues was undesirable in both connectionist models and linear models. An example of two such cues was CURR and ACID. Only one of these cues was selected.

To make the model comparable to the models of section 9.1, the number of cues should be roughly similar to the number of cues used in these models. For "profitability", PROMA and OPMA were closely related indicators. Selecting ROI, PROMA and ROE represented a selection of the 42.7 % most indicated "profitability" cues. For "financing", a ratio of STL/LTL was computed. BER was selected as the only indicator of the "leverage" concept, and the selection of LTINV, LTL/STL and BER represented 46.8 % of the most indicated "financing" cues. For "liquidity", CURR and ACID were closely related, and selecting one should be sufficient. Selecting CHLIK in addition, implied selecting an indicator not of ratio scale, but a total of 46.7 % of the indicators of "liquidity" was then represented. For "leverage", selecting BER represented 50.1 % of the indicators used, and should be sufficient. Consequently, a model with 15 input units was set up. 14 of these were from 7 ratios of the two consecutive years, the last was a change indicator; CHLIK.

The standard cross validation procedure described above was followed with similar parameters as in section 8.1 and 9.1. The performance results of the simulations are shown in tables 9.23 and 9.24.

Iterations:	5000	10000	15000	20000	25000	30000
Model:						
HID0	0.219	0.220	0.226	0.229	0.234	0.235
HID2	0.241	0.209	0.200	0.197	0.201	0.203
HID4	0.202	0.203	0.205	0.211	0.214	0.217
HID6	0.189	0.202	0.203	0.206	0.209	0.215
HID8	0.191	0.196	0.197	0.201	0.203	0.208
HID10	0.201	0.193	0.196	0.200	0.205	0.206
HID12	0.196	0.201	0.200	0.206	0.208	0.208
HID14	0.199	0.197	0.199	0.207	0.208	0.212

Table 9.23. Mean squared error (MSE) of the level diagnosis in a combined subject selected model (N=75)

Iterations:	5000	10000	15000	20000	25000	30000
Model:						
HID0	0.357	0.350	0.351	0.355	0.359	0.363
HID2	0.476	0.371	0.347	0.342	0.341	0.340
HID4	0.422	0.345	0.328	0.332	0.342	0.340
HID6	0.400	0.337	0.342	0.343	0.344	0.350
HID8	0.390	0.344	0.337	0.334	0.333	0.336
HID10	0.395	0.342	0.338	0.344	0.341	0.349
HID12	0.364	0.343	0.336	0.347	0.348	0.349
HID14	0.369	0.338	0.340	0.354	0.350	0.352

Table 9.24. Mean squared error (MSE) of the trend diagnosis in a combined subject selected model (N=75)

The model generally showed results comparable to the results of chapter 8 with a minimal MSEs at 0.189 for level diagnosis and 0.332 for trend diagnosis. Tests of the correlation of SE with targets and distance from mean targets showed a pattern similar to the models of chapter 8. Consequently, selection by cue importance did not give an improvement in performance results when compared to the results of chapter 8. In chapter 8, the models "determined" cue importance by adjusting their weights to relevant input units. The selection by subjects' cue importance indicators did not give a model with better generalisation results than the "self selection" of cues performed by the models of chapter 8.

9.4 Conclusions

A main conclusion of chapter 8 was that the number of weights was too large to be properly set by the relatively small learning sample of this study. To overcome some of these problems and to improve the generalisation properties of the models, weights were reduced by elimination of input units. Two procedures were explained, of which the quantitative approach was the most successful. A procedure based upon selection by using subjects' cue importance indicators did not prove significantly better than the original models of chapter 8. However, the *sensitivity* analysis based procedure gave models with significantly better performance results.

The procedure of chapter 8 was followed, and no other parameters than the number of input units were changed. The *performance* was significantly better for the constrained models than for the models of chapter 8. Generalisation was improved by weight reduction in the model. In our models, weight reduction was done by sensitivity analysis, but automated weight reduction during learning may have given similar results (Weigend et al., 1991).

Furthermore, it was shown that the fit of the connectionist models was significantly better than any of the 12 benchmarks provided in chapter 6 or in sections 8.1 and 9.1, when judged by the recommended measure, cross validated average squared error (White, 1990), for models of both level and trend diagnosis. These findings strongly supported the proposition P1 of chapter 4.

Furthermore, it was found that the connectionist models with hidden units performed significantly better than the models without hidden units. The comparisons between the performance of the two sets of models were done using the same estimation methods[1]. Thus, the improvement in performance was attributed to the internal representations formed in the hidden units. These findings generally supported proposition P2 of chapter 4.

From the simulations, *combined* models were found superior to the separate models. It seemed that level diagnosis was used as a "hint" to form relevant representations for trend diagnosis and vice versa within the same model.

In addition to these findings, some observations of relevance to connectionist modelling research were made. As expected, minimum average squared errors were found for a higher number of learning iterations in the constrained models than in the models of chapter 8. Thus, overfit did not occur so early. Furthermore, the reduced complexity gave better performance, supporting our suggestion in chapter 8 that stopping learning and reducing the number of free parameters were not equivalent methods for increasing model performance. Finally, the sensitivity based method of input unit elimination applied here, was found more successful in increasing performance than using subjects' indications of important input variables.

The *representations* of the best constrained and combined models were analysed following a procedure similar to the one used in chapter 8. However, in this section we took the analysis further by using separate cluster analyses for each set of models, and combined these analyses with Hinton diagrams. Furthermore, we selected some representative models, and analysed these using cluster analysis of the hidden unit output space and plots of nonlinearity.

---

[1] Learning rule and procedure.

The final representations found in the constrained model were different from the representations in the models of chapter 8. As expected, more *local* hidden units were developed. Among the local units, we did not find that every model used common hidden units as the models in chapter 8 did. In the smaller models, all local hidden units now had either a level or a trend orientation. Cluster analysis showed that the two groups of units had very different weight patterns. Two types of level-oriented and trend-oriented hidden units were found, reflecting the two functionally different ways of organising the small models. Some connectionist models developed level-oriented units also representing common aspects of level and trend diagnosis, whereas the opposite situation was found for the other type of small models.

In models with more than 2 hidden units, the additional units did not improve the performance of the model, but their functional organisation differed somewhat from the smaller models. Generally, the additional units were somewhat more distributed, and some "miniatures" of the local hidden units were found. In some models, a special type of unit presumed to detect "yesterday's situation" was found. This unit had a weight pattern to the output units similar to the "difference" detectors of chapter 8. However, the most interesting organisation of the larger models was illustrated by the analysis of the model C3. This model had three hidden units representing the common aspects of the level and trend diagnoses by one unit, the unique aspects of the level diagnosis in one hidden unit, and the unique aspects of the trend diagnosis in the last hidden unit. The common hidden unit's output was an approximately linear function of a selected set of financial cues from the last of the two consecutive years. The specialised level-oriented hidden unit's output was a highly nonlinear function of the same cues, whereas the output of the specialised trend-oriented hidden unit was an approximately linear function of the change in a selected set of financial cues. Using these representations, a level diagnosis was formed by linearly combining the common and specialised level-oriented hidden units' outputs. Trend diagnosis was formed by linearly combining the common and the specialised trend-oriented hidden units' outputs. Consequently, the abstracted stimulus dimensions represented by the hidden units were complex.

Because the connectionist models with two hidden units compressed the representation of two of the hidden units in the model described above into one hidden unit, their representation was even harder to interpret. When investigating the nonlinear outputs of these hidden units, we found that the units showing nonlinear outputs, were always level-oriented. However, some of these units also activated the trend diagnosis unit. The approximately linear outputs of the trend-oriented hidden units were also used either to activate the level and trend diagnosis, or only to activate the trend diagnosis units. Thus, we concluded that two functional organisations were used. In one organisation, level diagnosis was indirectly[1] a nonlinear function of the value of a selected set of financial cues from the last of the two consecutive years *and* an approximately linear function of the change in the value of a selected set of financial cues. Trend diagnosis was indirectly an approximately linear function of the change in the same set of financial cues. In the other organisation, trend diagnosis was indirectly an approximately linear function of change in a selected set of cues *and* a nonlinear function of the value of a selected set of financial cues from the last of the two consecutive years. Level diagnosis was now only a nonlinear function of the value of the same selected cues of the last of the two consecutive years.

As in chapter 8, we concluded that the hidden units did not specialise on diagnostic concepts or diagnostic areas in the way expected by theory. To further elaborate on this conclusion, we tested the performance of a model with the connections restricted to impose the *expected* functional organisation of diagnostic task. The performance results of this model are shown in table 9.25.

Iterations:	5000	10000	15000	20000	25000	30000
Model:						
Restricted level	0.466	0.182	0.197	0.198	0.197	0.198
Restricted trend	0.651	0.377	0.345	0.349	0.346	0.347

Table 9.25. Mean squared error (MSE) of the level and trend diagnoses in a constrained and restricted model (N=75)

From table 9.25, we found that the performance of the restricted model was quite comparable to the corresponding restricted model of chapter 8. The correlation of errors and targets followed the general results found in both chapter 8 and section 9.1. To illustrate the restricted model, a Hinton diagram of the weights of one version of the models is shown in figure 9.9. The weights of figure 9.9 were quite different from what should be expected from theory. Hidden unit 1 worked as a trend-oriented common hidden unit. Hidden unit 2 was so far considered to be a

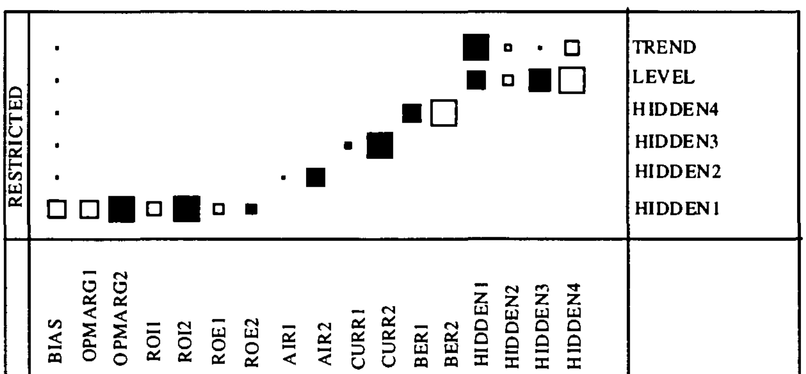

Figure 9.9 Hinton diagram of weights in a constrained and restricted connectionist model

"financing" indicator. Hidden unit 3 was a level-oriented unit, possibly detecting "liquidity

[1] By the term "indirectly", we mean that the diagnoses are functions of the functions specified in the text.

level", and hidden unit 4 computed a reversed difference in BER1 and BER2, and inhibited both the level and trend diagnosis units. In chapter 8 we tested the relationship between the hidden unit outputs and the responses collected on the particular diagnostic variables from the subjects. A similar test was performed here, and the correlations between hidden unit outputs and subjects' diagnostic area responses are shown in table 9.26.

Diagnostic area	HIDDEN1	HIDDEN2	HIDDEN3	HIDDEN4
PROFLEVEL	0.742**	-0.091	0.105	-0.464**
PROFTREND	0.826**	0.017	0.152	-0.298**
FINLEVEL	0.411**	-0.254*	0.561**	-0.614**
FINTREND	0.591**	-0.204	0.379**	-0.529**
LIQLEVEL	0.223	-0.241*	0.719**	-0.562**
LIQTREND	0.380**	-0.225	0.431**	-0.506**
LEVLEVEL	0.420**	-0.376**	0.475**	-0.898**
LEVTREND	0.662**	-0.163	0.297**	-0.659**

Table 9.26. The constrained and restricted model's correlations of hidden unit outputs with composite judge diagnosis of the four diagnostic areas "profitability", "financing", "liquidity" and "leverage" for level and trend respectively (** and * indicates significantly different from 0 at α=0.01 and 0.05 respectively) (N=75).

From table 9.26, we found that the correlation pattern did not correspond to the pattern expected from theory. Both hidden unit 1 and 4 operated as detectors of a "common factor". Hidden units 2 and 3 detected all but "profitability", and despite the restricted connections in the model, the correlation with the corresponding diagnostic area was not particularly high when compared across diagnostic areas. As in chapter 8, it seemed as if the restricted model tried to implement the functionality also found in the unrestricted connectionist models of section 9.1. The functional organisation of the task developed in the connectionist models by learning still seemed superior to the one imposed by theory.

Consequently, the representations of the constrained connectionist models did not correspond well to theoretical concepts in financial diagnosis. Rather, they corresponded more to complex and coarse detectors of the common and unique aspects of the diagnostic concepts "level" and "trend". The representations in the constrained models were more complex than the representations of the models in chapter 8. Thus, analysis and interpretation was more difficult. The complex concepts represented in the constrained models were implemented in *variables* representations. Even though these variables were both linear and nonlinear, they had response patterns more similar to a variable than to a feature detector.

Judging the cognitive relevance of these representations was difficult, but the rule-plus-exception interpretation of the models in chapter 8 previously found to have cognitive relevance did not give an accurate description of the behaviour of the models in this chapter. As in chapter 8, further interpretation of the representations should be made with reference to the molar behaviour of the models. Two alternatives were found useful in chapters 8 and 9. A rule-based interpretation was used for the models of chapter 8, whereas both rule-based and prototype-based interpretations were used in this chapter.

Using a rule-based interpretation, we could describe very roughly what the *simplest* of the constrained *small* connectionist models did. By simplifying, the functional organisation can be described in the following way: One hidden unit detected a general change in most of the financial cues over the two consecutive years. This change was combined linearly to form the diagnosis of both level and trend. Another hidden unit focused on the cue values of the second of the two consecutive years for most of the financial cues. The response of this hidden unit was nonlinear. It was roughly linear in the low to moderate range, but highly nonlinear for high and extremely high cue values. The response of this hidden unit was used to reduce and enhance the value of the level diagnosis response unit, but extremely high cue values did not add linearly to the level diagnosis. Thus, the simplified heuristic implemented by these models were: Test if there is a positive change in financial cue values over the two consecutive years. If so give a correspondingly good preliminary level and trend diagnosis. Check the values of a broad range of financial cues in the second of the two consecutive years. If they are generally low, adjust the level diagnosis down as a function of the value of the financial cues, and retain the trend diagnosis. If they are moderate, retain both preliminary diagnoses. If they are high or extremely high, retain trend diagnosis and adjust the level diagnosis as a nonlinear function of the cue values somewhat upward.

Whether this heuristic had cognitive relevance was not definite, but if viewed in the light of a bounded rational (Newell & Simon, 1972) diagnostician, it certainly seemed to reduce the demands on cognitive information processing capacity.

A prototype-based explanation of the connectionist models could also be made. When cluster analysing the hidden unit outputs of a model with three hidden units, the different stimulus patterns were placed in different clusters depending on their hidden unit output patterns. Eight different clusters were identified in figure 9.8. The output of the hidden units could, when they were combined, be interpreted as if they computed the similarity of the stimulus patterns to the centre of each of these clusters. Thus, the cluster centres could be interpreted as prototypes. Some interpretable prototypes could be identified, such as the "bankruptcy firm", the "bad, but promising firm", or the "good, but with alarming trend firm". It seemed from this analysis that using prototype theory terminology was a second fruitful way to interpret the behaviour of the connectionist models.

Consequently, the representations of the connectionist models could be explained by studying the internal representations formed in the hidden units of the models. However, these representations were often difficult to interpret *directly* in terms like strategies, heuristics, rules, variables, concepts, prototypes or classes, previously shown to have cognitive relevance in either competence or behavioural theory of financial diagnosis, or in general classification theory. However, the molar behaviour of the connectionist models could be interpreted using several of these terminologies. Two were applied here; the rule-based and the prototype-based, and these interpretations resembled quite closely similar explanations given in production system terminology and in prototype theory terminology.

In addition to these findings on the representations of the constrained connectionist models, some other observations of relevance to connectionist modelling research were made. The weight pattern of the trained constrained connectionist models varied, but compared to the models of chapter 8, weight initialisation did not have the same effect. For most constrained connectionist models, a small number of local solutions to the minimisation problem was found. These local solutions were represented by a similarly small set of local hidden units. The use of distributed hidden units was considerably reduced in the constrained models, but despite this reduction, the complexity of the local hidden units was very high.

PART V - DISCUSSION

Chapter 10. Discussion

In this study, attempts were made to follow many of the principles traditionally applied to studies in the "context of justification". To give some examples, propositions were developed, an experimental research design was set up to provide valid measures of stimulus-response data, a comparative perspective was applied to evaluate connectionist model success, and statistical measures and cross validation procedures were used to evaluate propositions. Despite these attempts, connectionist modelling is research still in the "context of discovery" (see Seidenberg, 1993). In the evaluation and discussion of the conclusions of this study, the exploratory nature of connectionist modelling should be kept in mind.

The Cook and Campbell (1979) lists of threats to validity are often used as a basis for a discussion of the conclusions in administrative science studies. There are several difficulties with this approach in the present study. First, the framework of Cook and Campbell (1979) was primarily developed for evaluation of causal studies, and this study is exploratory in nature. Second, many of the methods used to establish that Cook and Campbell's (1979) validity threats are insignificant, are based upon the application of linear methods. To take a few examples, reliability is often documented using correlational measures or other linear methods such as principal components analysis. Construct validity is often tested using linear measures of convergence and discriminant validity. In fact, when these methods are used to develop measures, linear models will benefit over alternatives even by the procedure of measurement development. Despite these difficulties, the framework of Cook and Campbell (1979) is carefully applied here to evaluate the empirical parts of the thesis. However, modification of the framework is made to organise the discussion of simulations and results.

Before turning to the discussion of conclusions, section 10.1 summarises the research efforts and conclusions documented in previous chapters. In section 10.2, we discuss and evaluate our main conclusions. Some theoretical, methodological and practical implications of our research effort are presented in section 10.3, and in section 10.4 suggestions for further research on connectionist models of financial diagnosis are made.

10.1 Summary and conclusions

In chapter 2, we showed how financial diagnosis has been studied from three different perspectives; the judgement modelling, the cognitive, and the predictive perspective. Despite their different foci and modelling methods, our understanding of the financial diagnosis task benefited from this theoretical triangulation. Somewhat surprising was the absence of applications of cognitive classification and categorisation theory to financial diagnosis.

Classification theory has been applied to other diagnostic tasks with considerable success (e.g. Weber, Böckenholt, Hilton and Wallace, 1993). Consequently, traditional and contemporary theories of classification were reviewed in chapter 3, and a model of financial diagnosis based upon a connectionist model of classification, was suggested in chapter 4. Furthermore, propositions on the relationship between this model and traditional models of the three approaches to financial diagnosis were made. The first proposition, P1, stated that connectionist models of financial diagnosis should show better fit than benchmarks of traditional models. The second proposition, P2, stated that the improved fit could primarily be explained by the ability of the connectionist models to build internal representations, and the third proposition, P3, stated that these internal representations should have cognitive relevance.

To evaluate these propositions, a financial diagnosis experiment was set up. 108 subjects participated in the diagnosis of 75 randomly selected small and medium sized firms. Full financial statements and selected ratios of two consecutive years were used as stimulus material, and several measures of diagnostic responses were collected. The treatment plan resulted in 324 diagnoses of the 75 firms, averaging 4.32 diagnoses per firm. To create the stimulus-response pairs representing learning and test samples of the connectionist model, composite judge diagnoses were computed.

A simulation design was developed that accommodated resampling methods and cross validated measures to evaluate the performance of the connectionist model. Furthermore, a number of benchmarks were developed using traditional methods of the judgement modelling approach to financial diagnosis.

In chapters 7, 8 and 9, the propositions of chapter 4 were evaluated using three simulation experiments with varying stimulus and response representations. The first simulation used a stimulus representation consisting of 17 selected financial ratios, of which 15 were provided with values from two consecutive financial statements[1]. Diagnostic response was measured by a bankruptcy classification variable. The second simulation used the same stimulus representation as the first, but diagnostic responses were measured by composite judge assessments of level and trend diagnoses of the financial situation. In the third simulation, the diagnostic response representations of the second simulation were used, but sensitivity based measures were used to select a constrained set of stimuli. Generally, model fit was improved from simulation one through simulation three. The main conclusions that could be drawn from the simulations should be related to the propositions made above, but a set of

[1] Thus, 32 independent variables were used.

conclusions with relevance to connectionist modelling in general could also be drawn, and have been summarised in the concluding sections of chapters 7, 8 and 9.

Strong support was found for proposition P1. The connectionist models showed significantly better fit than traditional benchmarks when evaluated by cross validated average squared error. Furthermore, analysis of error distributions revealed a smaller standard deviation of errors and explainable outliers for the connectionist models. In particular, the model with constrained stimulus representations and composite judge diagnoses showed favourable fit to the financial diagnostic data.

For the connectionist models showing significantly better fit than the benchmarks, tests were made to evaluate proposition P2. In these tests, significantly better fit was found for the connectionist models with hidden units than for the models without hidden units. Equivalent initial weights and parameter settings were used for both types of models, and the same learning rules were applied. Furthermore, overfit was controlled by the optimal stopping point rule for both types of models. Consequently, it could be concluded that the difference in performance was explained by the internal representations of the hidden units. These findings supported proposition P2.

Evaluation of proposition P3 was much more difficult. The representations built by the hidden units of the connectionist models were expected to consist of derived stimulus dimensions reflecting different diagnostic concepts, such as "profitability", "financing", "liquidity" and "leverage". The theoretical importance of these concepts documented in chapter 2 and discovered in factor analysis studies of financial statement data (Gombola and Ketz, 1983; Pinches et al., 1973), founded such expectations. However, completely different and much more complex representations were built by the hidden units. Different internal representations were developed when stimulus and response representations were changed. At first sight, the internal representations also seemed sensitive to initial weights. These findings complicated the analysis and interpretation of the representations. However, a method was developed that made analysis of internal representations possible across models with different initial weights, but differences in internal representations resulting from different stimulus and response representations had to be interpreted separately for each model.

The most typical representation developed by the models with the large stimulus set consisted of a local "common" hidden unit and one or several more distributed hidden units. Both these representations were in the "variables" form. The way these hidden units formed the financial diagnosis could be explained by using a rule plus exception heuristic. Empirically, level and trend diagnoses were strongly correlated, and a rule stating that level and trend diagnosis values were similar and depended partly upon the values of financial statement cues from the most recent year and partly upon the change in financial statement cue values, was implemented by the "common" hidden unit. Exceptions to this rule were implemented by the distributed hidden units. The rule plus exception heuristic has been proposed by several authors (e.g. Nosofsky, Palmieri & McKinley, 1994) as a cognitive model of classification.

The representations developed by the models with the constrained stimulus set consisted of two or three local, and no, one, or some distributed hidden units. The distributed hidden units were not shown to improve model performance. For networks with two hidden units, the "common" hidden units previously found in the larger networks had developed an increased focus on either level or trend diagnosis. The remaining unit was also local, and focused exclusively on either trend or level diagnosis respectively. Of particular interest were the networks with three hidden units with one "common", one trend independent level-oriented, and one level independent trend-oriented hidden unit. The representations of the local hidden units were still of the "variables" form, even though all level-oriented "common" and unique hidden units, now were nonlinear and highly nonlinear, respectively.

A rule-based interpretation of the molar behaviour of the small models suggested that one hidden unit still implemented the rule that level and trend diagnoses were correlated. If this unit was trend-oriented, the remaining hidden unit was used to adjust the level diagnosis nonlinearly. Thus, a linear rule and a nonlinear regulator were used. A rule-based interpretation of the networks with three hidden units was simpler, suggesting one unit implemented the rule that level and trend diagnoses were correlated. The two resulting hidden units were used to adjust the diagnosis given by this rule linearly for the trend-oriented hidden unit, and nonlinearly for the level-oriented hidden unit. An elaboration of the rule-based interpretation was given in section 9.4. This representation constituted a more complex rule-and-adjustment heuristic than the rule-plus-exception heuristic found in the connectionist models of chapter 8. The cognitive relevance of this representation was open for discussion, but compared to a functional division of the diagnosis into diagnostic areas later to be combined into a final diagnosis, it represented an interesting simplification. For subjects of limited cognitive capacity (Newell & Simon, 1972), a heuristic detecting the change in the values of two financial cues and an adjustment based on the value of the most recent cues, certainly reduces the strain put on cognitive resources.

Furthermore, the improved performance of the combined and constrained connectionist model could be explained by the "hints" given to each diagnostic response variable from the other. The representations developed in a model of level diagnosis were more favourable when "hints" of the trend diagnosis were given to the same model, and vice versa. This finding was exclusive to the constrained models, and illustrated that model representations showed sensitivity to both correlated stimuli and correlated responses.

In addition to a rule-based interpretation of the molar behaviour of the connectionist models, a constrained connectionist model's behaviour could be interpreted using prototype terminology. By using cluster analysis on hidden unit output space of a selected model with three hidden units, we showed that the behaviour of the model could be interpreted as if similarity to a set of prototypes was computed. Examples of the different prototypes were; the "bankruptcy firm", the "bad, but promising firm", or the "good, but with alarming trend firm" prototypes. Consequently, interpretation of the behaviour of connectionist models could be made in terms used in several cognitive theories of classification, but the *direct* interpretation of the representations formed by hidden units in such terms, was difficult.

10.2 Discussion

A major threat to the validity of our conclusions is lack of validity in the collected stimulus-response data pairs. The *empirical* validity of the data can be evaluated using the framework of Cook and Campbell (1979). This evaluation is organised in a discussion of the statistical conclusion validity, the internal validity, the construct validity and the external validity of the empirical parts of the study. However, more threats to the validity of our conclusions are related to the *simulation* design and to the models used to operationalise the proposed theory. These threats are mostly threats to statistical conclusion, construct and external validity, and are discussed separately. Finally, we discuss the validity of our conclusions in *general*.

Empirical validity in this study means that the stimulus-response pairs used in the simulations were valid financial diagnoses. Of the traditional threats to *statistical conclusion validity*, lack of reliability and small sample sizes were most relevant. To secure reliability, the instrument was pretested, the stimulus situation was concentrated and treatments were given simultaneously to all subjects. In general, small sample sizes also threaten validity, but as shown in chapter 5, there was a close relationship between stimulus sample size and subject sample size in this study. More financial diagnoses could only be obtained at the cost of reduced control over the individual variation among diagnosticians. Thus, we refer to the considerations made on this subject in chapter 5.

Internal validity threats, such as history and maturation, constituted relevant threats in our study. However, most of these threats were controlled by randomisation. For example, maturation effects occurring as a consequence of repeated stimuli, were controlled by randomising the order of the stimuli so that some subjects diagnosed a firm early, while other diagnosed the same firm late during the experiment. The similar principle was used to control most threats to internal validity, such as test effects, instrument variation, mortality, imitation and compensation. This procedure may have introduced random errors in our data, but the principle of avoiding bias at the cost of increased random errors was followed since these errors would only strengthen significant findings.

Similar to other threats to internal validity, the interaction between treatment and selection was controlled by randomisation. However, careful considerations were made of the relationship between internal and external validity before the homogeneity of the subject sample was accepted. A less homogeneous sample could weaken the internal validity, but strengthen the external validity. Since the propositions of this study were related to modelling properties of connectionist models and not to the generality of the final models, internal validity considerations were regarded most important. Thus, the homogeneity of the subject sample serves to strengthen internal validity in this study. Furthermore, there were no findings in financial diagnosis research indicating that the homogeneity of the chosen subject sample was in obvious favour of our propositions. Except for these considerations, the internal validity obtained by the experimental design strengthened the assumption that the stimulus-response pairs produced, were valid.

Whether or not these valid stimulus-response pairs were valid financial diagnoses further depends on the *construct validity* of the study. Construct validity is three-fold. To secure construct validity, theoretical concepts must be clarified properly, they must be operationalised properly, and the relationship between them must be operationalised properly. Of these presumptions, the first two are of primary relevance here[1].

Proper clarification of theoretical constructs was attempted by the reviews of chapters 2, 3 and 4. To simplify, clarification of relevant stimuli and responses was attempted by the review of financial diagnosis theory in chapter 2. Different theories clarifying the relationship between stimuli and responses in classification tasks in general and in financial diagnosis tasks in particular were reviewed in chapters 3 and 4. Thus, lacking theoretical clarification of concepts was hopefully avoided by the weight put on theory reviews.

Using the reviews as theoretical bases, stimuli and responses of the financial diagnosis task were operationalised in the stimulus material and response measures presented in chapter 5. However, several threats to construct validity are relevant. Mono-operational bias was reduced by multiple operations of both treatments and responses. However, mono-method biases represented by the form of stimulus presentation and response measures are relevant threats. The format, order and content of the stimulus material were carefully developed to match a realistic set of information given in financial diagnosis tasks, and written form of this material is traditional. Actually, giving the stimulus set in other forms could itself have

[1] Operationalisations of the relationship between stimuli and responses were done by model and are evaluated separately below.

threatened construct validity (see Moriarty, 1979). To obtain realism in the stimulus manipulations, real financial statement information was used. To avoid mono-method bias in the response measures, both predefined ordinal response scales and the subjects' own linguistic terms were used as measures. Both measures were used in the simulations of part IV. To further avoid effects resulting from different response styles and individual measurement bias, composite judge diagnoses were used. These were designed by applying composite judge rules, or by computing composite judge averages of the response measures as explained in chapter 5.

A danger in using real financial statement stimuli and random allocation of this stimuli to subjects, was the danger of insufficient manipulation. However, manipulation checks were performed by analysis of variance showing sufficient variation on both the level and trend diagnoses across firms. Hypothesis guessing would most likely give more extreme responses, but no information that this was a problem was found, and prior probabilities of diagnostic classes were not given. Furthermore, no indication was found that hypothesis guessing, if present, would consistently favour connectionist models.

The homogeneity of subjects represents a major threat to *external validity*. However, several actions were taken to investigate and control for this threat. First, previous research (see Bonner & Pennington, 1991) had shown that financial diagnosis was likely to be a task where small differences in task behaviour could be found between experienced and less experienced subjects. Second, our subjects were graduate students with financial diagnosis experience, the majority through professional experience, and a small fraction through prior education. Third, a homogeneous sample of more experienced subjects may also have represented a threat to external validity. To secure internal validity, some homogeneity of the subject sample was desirable, and it was beyond the purpose of this thesis to generalise the stimulus-response relationships obtained to more heterogeneous populations. However, it was a purpose of this thesis to generalise the conclusions on propositions P1, P2 and P3. Of these, no prior assumptions could be made that connectionist models with hidden units were favourably sensitive to the homogeneity of our sample. However, the generality of the connectionist model representations is in principle threatened by the homogeneity of our sample, and it is not obvious that the heuristics implemented by the connectionist model representations generalise to, for example, expert subjects. To establish whether this is the case or not, further research is required.

The second way homogeneity threatens external validity in this study is caused by the small and medium sized firms used in the sampling frame for the randomisation of stimuli. Again, after careful consideration of the relationship between internal and external validity, we concluded that internal validity would be strengthened by the homogeneity, while external validity would be weakened. As an example, subjects' firm recognition, a threat to internal validity[1], was considered reduced by using small and medium sized firms. The gain from reducing internal validity threats due to homogeneity of firm size, was considered more important than the small loss in external validity. Besides, with manipulation checks we showed that stimulus variation was sufficient. In general, the homogeneity of subject and stimulus samples in this study reduces the generality of our conclusions somewhat, but since external validity presupposes internal validity, the *relevant* generality of our conclusions must be considered high.

Interactions between situation or time and treatments represented inescapable, but minor threats to external validity in this study. Actually, obtaining realism in the diagnostic situation was a major goal in constructing the stimulus material, and time will always represent an economic and historical context for financial diagnosis, reducing generality of the judgements performed by the diagnosticians.

The methodological principle of using operationalisation by model, requires a special discussion of relevant validity issues. Most of the validity issues of the *simulation* design applied in this study are statistical conclusion, construct and external validity issues.

Statistical conclusion validity is threatened by inappropriate establishment of model performance differences. Care must be taken not to violate statistical assumptions of benchmark model methods and in the selection of performance measures. Since, for example, multinormality and multicolinearity assumptions could easily have been violated in this study, recommended methods for independent variable construction in benchmark models were followed. In addition, a cross validation procedure was adapted and applied, making statistical inference based upon learning sample data unnecessary. Similar procedures have been recommended for modern nonparametric methods (Efron & Tibshirani, 1993). Statistical inference was necessary in the evaluation of performance differences. A simple t-test of the differences between means was applied, based upon the applicability of the central limit theorem, even though cross validated squared errors were not normally distributed.

Traditional construct validity was evaluated above for the links of figure 3.10, going from theoretical constructs to operations at the observational level. Three operationalisations illustrated by short arrows in figure 3.10 have not been discussed; the operationalisations of representational assumptions and the two implementations of stimulus and response operations at the model level. These are all related to *construct validity*, but require special discussion when operationalisation by model is used. To secure construct validity of the

---

[1] And consequently also to external validity.

representational assumptions of connectionist theory, we stayed as close as possible to the original model implementation suggested by Rumelhart, Hinton and Williams (1986). Only simulation parameters deviated somewhat, but these must always to be set relative to the data set and size of the model. In addition, parameter settings were kept constant during all simulations. Stimulus and response measures were only rescaled linearly to fit the transfer functions of the connectionist models. No other change in the representational form of the data was made to fit model operations.

Benchmark model operationalisations of the relationship between stimulus and response were of the simplest form. These operationalisations took form either as linear transformations of stimulus data weighted to form responses, or as direct linear weightings of stimuli to form responses. Our proposition was that the mapping was more complex, and required internal representations. Thus, the operationalisation of the connectionist model was an operationalisation of the "relationship between putative stimuli and responses"[1]. Since this relationship was cognitive, it was not directly observational. Consequently, evaluation of model operationalisation and interpretation of developed representations are closely interrelated. However, in the alternative proposed by information processing theory, operationalisation is completely theory driven, and no posterior interpretation and evaluations are allowed. As previously concluded, an advantage of connectionist models is that they are open to validity evaluation based upon both responses and the internal representations developed.

The *external validity* of the model simulations refers to whether or not there was something unique with our subject data, simulation situation and time, that prevented similar simulation results to be obtained in general. Again, the uniqueness of the data due to homogeneity of subjects and stimuli seems the most relevant threat to such generalisations. Certainly, other representations may develop if subjects using different representations and information processes had performed the diagnoses, or if all diagnoses were of firms requiring different diagnostic behaviour. The question is if it was something special with our subjects or stimuli that prevents generalisation. If the intermediate abstractions are interpretable in the "variables" form, for instance if they consist of variables representing traditional diagnostic concepts, equally superior results should be obtained with similar connectionist models. If exemplar based representations are used by the subjects, other models, such as ALCOVE, may prove superior. The alternative model is also connectionist, and our model may be modified to handle similar representations (see Kruschke, 1993b). There was no indication available, a priori, that the homogeneity of subject or stimulus samples were consistently in favour of our simulation results, or suggested different representations should be developed

[1] The terms are deliberately chosen to illustrate similarity to the "construct validity of putative causes of effects" term used by Cook and Campbell (1979).

that our model could not implement[1]. In addition, the purpose of this study made us give priority to internal validity of *simulation* design as well as to experimental design. However, further investigation to increase the generality of simulation results and model representations should be performed with other subject and stimulus samples.

Of the *conclusions* on the propositions in chapter 4, the conclusions on P1 and P2 were relatively decisive, and their validity has been justified above. The conclusions to be drawn on proposition P3 seemed less obvious, and require further discussion. We generally applied the principle that the internal representations were considered to have cognitive relevance if they resembled intermediate abstractions, such as strategies, heuristics, rules, variables, concepts, prototypes or classes, previously documented of relevance in either competence or behavioural theory of financial diagnosis, or in general classification theory. The hypothesised task partitioning into concepts or variable representations of different diagnostic areas was not found. Neither were *direct* representations of prototypes or rules found. Rather, the representations were complex variable representations of abstract concepts. These concepts had not previously been identified in financial diagnosis research, but upon investigation they seemed meaningful as concepts in simplified financial diagnosis. Interpreted as heuristics, rules or prototypes, the operations on these representations could be described at the molar level. As heuristics, they *simplified* the financial diagnosis task. The question remains, however, whether the subjects actually applied the heuristics, or if the heuristic only approximately described their behaviour. Unfortunately, this question remains principally unsolved. Some investigation into the cognitive relevance of the heuristics may be obtained by interviews and subjects' posterior introspection, but this is an issue for further research[2]. However, there is a fundamental way in which the representations must be considered cognitively relevant: As long as the representations are functional in the description of subject behaviour, and they are based on a set of operating assumptions used to describe similar cognitive phenomena, they must be considered cognitively relevant at a molar level. As mentioned in chapter 3, the representations of any cognitive model[3] should give functional explanations at a psychological level, but we can not necessarily prove their cognitive relevance at a lower level.

---

[1] A considerable number of cross validated and non-cross validated initial experiments were performed in our project that are not reported in this study. Examples of these simulation experiments are tests of RBF-models (Moody & Darken, 1988), tests of modular versions of the models (see Haykin, 1994, p. 473-478), auto associative pre-processing of inputs (Chalmers, 1990), variations in binary and continuous stimulus and response representations, models of individual general diagnoses and of diagnostic area response, models including response error and measurement error representations, and models using modified learning rules (Jacobs, 1988), just to mention a few. However, the most consistent results were obtained following the fairly simple principles and formulations reported here.

[2] As mentioned in chapters 1, 2 and 3, the methodological problems with such research were some of the reasons why a connectionist approach was taken in this study.

[3] And other models including latent variables or processes as well. For a similar discussion of models in economics, see Cyert and Grunberg (1963).

10.3 Implications

The main findings in this study have several theoretical and methodological, and some practical implications. Before turning to these, we summarise some of the contributions made by the research effort in this study.

In addition to the main conclusions drawn from the evaluation of propositions P1, P2 and P3, summarised and discussed above, the research documented in this thesis contributes in several ways. First, a theoretical framework was provided in chapter 2, integrating research on the financial diagnosis task from several perspectives. Second, the application of classification theory in general and connectionist classification theory in particular to financial diagnosis, represents a new approach in cognitive accounting. The application of connectionist classification theory to financial diagnosis further represents a new method for studying cognitive phenomena in accounting, using the well known methodological principles of judgement modelling research, while accepting the need for intermediate abstractions in complex cognitive tasks.

The methodological contributions are represented by the experimental design used to provide the stimulus-response data set for the connectionist models, and the simulation design developed to evaluate these models. First, the controlled conditions of the experimental setting represent a validity securing strategy, not traditionally applied in connectionist research. Second, the realistic task context and stimuli used here are similarly important validity securing techniques, not traditionally applied in classification research, whether connectionist or not. Third, the randomisation of natural stimuli was an important premise in creating task realism, but is also an important technique in eliminating internal validity threats.

Among the most important contributions of the simulation design were the use of cross validated performance measures and the method developed to provide these measures. Furthermore, a method for simultaneous control of performance measures and complexity in connectionist modelling was developed independently for our purpose, even though we acknowledge that a similar procedure was developed simultaneously by Moody (Moody & Utans, 1995). A sensitivity based method was combined with subjects' measures of cue importance to further reduce connectionist model complexity. This method proved successful, and has, to our knowledge, not previously been used in connectionist modelling.

In part IV, we showed how these measures and procedures could be used to develop and test connectionist model fit against traditional models, and how different connectionist models could be compared. With cross validated measures, traditional statistical tests could be applied to evaluate model performance differences formulated as propositions. To analyse the representations of connectionist models, several known methods were applied (Hanson and Burr, 1990), and new methods were developed to analyse several versions of a connectionist model across different initial conditions. The use of cluster analysis techniques to investigate common aspects of several versions of a connectionist model represents a new and promising method.

10.3.1 Theoretical and methodological implications

The main findings and the other contributions referred to above, give implications of three types; theoretical, methodological and practical. The theoretical implications relate both to the application area of the study; financial diagnosis theory, and to the theory-supplying area of the study; cognitive classification theory.

This study has shown that classification theory in general and connectionist classification theory in particular, offer cognitive models of basic tasks in behavioural accounting, such as financial diagnosis. However, the main implication for financial diagnosis research is that these models can be used to unify judgement modelling research and methodology with cognitive process research. With reference to the lens model of figure 2.1, connectionist models can be used to build cognitive models of the right hand side of the lens model with the use of methodological principles similar to judgement modelling methodology. Furthermore, these cognitive models are open for posterior analysis and evaluation to determine the cognitive relevance of their internal representations.

To cognitive classification research, this study has shown that practical task contexts in accounting offer application areas of classification theory. We suggest this implies that cognitive classification theory could fruitfully be applied to practical tasks with richer task contexts than the "synthetic" classification tasks traditionally used. The discovery and evaluation of internal representations in connectionist models of financial diagnosis in this study illustrate how practical application areas require careful analysis and interpretation when diagnostic relationships are not predefined by the experimenter.

The methodological implications of this study relate both to traditional experimental methodology and to connectionist simulation methodology. The method for developing internal representations in connectionist models is an indirect method of investigating intermediate abstractions. This method can be applied to avoid the methodological problems of the direct methods of information processing theory. Two methodological implications stemming more directly from this thesis are the applicability of natural and of randomised stimuli. Particularly, when the stimulus material is complex and artificial construction of the material represents a threat to validity, natural stimuli can be used. To secure realism in this material, the selection can be controlled by a randomisation procedure. Here, we have shown that this procedure is applicable in practice, and consequently, it may be applied in similar experimental settings.

The method often used in connectionist simulations to test generalisation of model performance is to hold a few preselected examples out of the learning sample[1] (Gluck & Bower, 1988a, 1988b). Full cross validation is an extension of the hold-out principle, traditionally considered to be too computationally demanding for practical applications. Here, we have shown that the method is applicable in connectionist simulations of realistic tasks when the sample size is not too large. Fortunately, cross validation has greater relevance when the sample size is small. Thus, better measures of model fit can be obtained in simulation studies in experimental psychology following the cross validation procedure. The last methodological implication that can be drawn from this study is that connectionist models are analysable across initial conditions using the cluster analysis techniques shown in part IV. This represents an extension of previous cluster analysis techniques that have exclusively focused on one version of a connectionist model (Hanson and Burr, 1990).

10.3.2 Practical implications

The support for the propositions P1, P2 and P3 of chapter 4 found in this study, gives several implications to practising financial diagnosticians, to financial diagnosis teaching, to financial diagnosis expert systems applications, and to other practical task areas in accounting similar to financial diagnosis.

The simple heuristics implemented by internal representations in our connectionist models[2] stand in contrast to the rather complex diagnostic areas claimed to be functional by practising financial diagnosticians. Accepting the claim that intermediate abstractions representing these diagnostic areas are functional in diagnosis, the simplified heuristics applied by our models may explain some of the diagnostic errors performed by practising diagnosticians. It may well be the case that practising diagnosticians claim the functionality of complex intermediate abstractions representing diagnostic areas, but actually apply simple heuristics in their own diagnostic behaviour. The indirect methods applied in this study may have been an important factor in revealing this discrepancy. As mentioned above, however, whether or not these heuristics are generally applied by financial diagnosticians remains an unresolved question requiring further research.

---

[1] So called "transfer" tests in experimental psychology (see Gluck & Bower, 1988a, 1988b; Shanks, 1992)
[2] Even though the connectionist implementations are complex, the heuristics are simple.

Since the subjects in this study were graduate students participating in an advanced course in financial analysis, the findings have some implications for accounting teaching. First, the implications above for practising financial diagnosticians also apply to students. Despite the effort put on teaching analysis of diagnostic areas, simplified heuristics may still be applied by diagnosticians. Second, the methods applied in this study for indirect cognitive mapping may be applied during a financial analysis course to investigate how internal representations change. Consequently, connectionist methodology complements other measures of knowledge, and can be used in experimental settings that realistically simulate the tasks financial analysis knowledge should be applied to.

A large area of practical knowledge modelling is the area applying the principles of cognitive modelling, artificial intelligence and expert systems development to practical tasks. Connectionist models have been proposed by several authors as providing the means to indirectly elicit human knowledge (Hawley, Johnson & Raina, 1990; Liang, Moskowitz & Yih, 1992). In this study, we have applied connectionist methodology to a task where expert system applications are numerous[1]. Our research has two important implications for the expert systems application area. First, connectionist models showed significantly better fit to human financial diagnostic data than other well established traditional benchmarks. Thus, connectionist models can competitively be applied to model human behaviour in financial diagnosis tasks. However, this study has not established if artificial neural network methodology is similarly superior to predictive benchmarks, even though research has been reviewed that weakly supports this suggestion[2]. Consequently, connectionist methodology should be applied in task contexts where human subjects have shown significantly better performance than predictive models.

Second, this study has shown how careful analysis of the internal representations developed in connectionist models can be performed, and that this evaluation should be an integral and important part of the evaluation of model validity. Evaluation of model fit by performance measures alone is not sufficient to establish the cognitive relevance of the internal representations built by connectionist models. Thus, if connectionist models are applied by expert systems practitioners, the principles developed in this study should be used to evaluate the relevance of the proposed elicited knowledge. In this study, cross validation principles were shown applicable in practice, and should be applied to establish if an observed model fit is valid in general, or if it is only caused by favourable learning and test sample splits. When applying these methods with care, connectionist methodology certainly represents an alternative and indirect method of knowledge elicitation.

[1] See Klein and Methlie (1990) and chapter 2.
[2] See chapter 4.

Finally, the findings in this study suggest that connectionist methodology can be successfully applied to practical tasks where human subjects have shown superior performance to formal models, and to tasks where traditional knowledge elicitation methods threaten validity or are costly.

10.4 Suggestions for further research

Two strategies can be used to suggest further research. One starts with the weaknesses and limitations of this study, and suggests new research improving upon these. The other strategy starts with the prospects of connectionist models demonstrated in this study and extends and elaborates on these.

To increase and test the external validity of our findings, replications of the principal methodology applied in this study with other, similarly homogeneous or heterogeneous subject samples, other firms in the stimulus sampling frame, and in a different historical context are suggested. Furthermore, replications of the principal methodology using other or more specified task contexts, are suggested.

A second suggestion for further research is replicating the principles in this study on a larger scale. First of all, increasing the subject and stimulus sample sizes is suggested. With a larger subject sample, the validity of the composite judge diagnoses used in this study can be evaluated. With a larger stimulus sample, more stimulus-response pairs can be obtained, and the general validity of the cross validated measures of generalisation error used in this study can be evaluated.

Finally, the generality of the internal representations of connectionist models discovered in this study can be further investigated using methodological triangulation, such as connectionist modelling, verbal protocol analysis and interviews, simultaneously.

These suggestions are all improvements in experimental and simulation methodology aimed at answering the same research questions as those investigated here. Because of the exploratory nature of the project, however, several new and interesting questions have been generated by our research. Suggestions for further research based on these questions should receive more attention.

A general question that can be raised from the research in our study is how formal and mathematical analysis can be used to explain similarities and differences between the mapping functions implemented by connectionist models and by the traditional benchmark models. As an example, large differences were detected between "factors" represented by our connectionist models and those extracted by principal components analysis. Similar topics are currently investigated in the artificial neural network community (Cheng & Titterington, 1994), but should receive similar attention in cognitive connectionist modelling[1].

The randomised stimuli used in this study can be applied to classification research, but the control principles used in classification research can also be applied to accounting tasks. We suggest that classification research should pay attention to realistic and practical task contexts, and that research on such tasks can benefit both classification research and research in the application area.

To behavioural accounting research, the success of connectionist models demonstrated in this study suggests that similar models should be applied to other accounting tasks than financial diagnosis. One question that is particularly interesting is if the favourable properties of connectionist models demonstrated in this study generalise to other cognitive accounting tasks. Another interesting suggestion stemming from our finding that connectionist models and artificial neural network models can be combined within the lens model framework of figure 2.1, is how this combination applies to accounting tasks. Judgement modelling studies have a long tradition in behavioural accounting, and the lens model can both regain its relevance and unify judgement modelling and cognitive studies with the application of connectionist theory.

Several unanswered questions on financial diagnosis were also generated by our research. Of special interest is how connectionist models' representations change as a consequence of richer stimulus representations. One example is if the trend- and level-orientations shown to be functional in some connectionist models in this study also extend to, for example, "stability" representations if longer time series of financial data are presented to subjects and models. Of similar interest is the question whether or not connectionist models show superiority to traditional benchmark models when stimulus-response data on expert financial diagnosticians are used. The preliminary suggestion stemming from our research is that this will be the case.

No comparisons of model fit between information processing theory based models and connectionist models have been made in this study, and differences in the two traditions make comparisons difficult. However, connectionist methodology and traditional information processing methodology can be combined. An interesting question is how the knowledge elicited with information processing methodology relates to the internal representations

[1] For an example, see the research on general recognition theory performed by Ashby and colleagues (Ashby, 1992).

developed in connectionist models. Such a combination would represent a dual route knowledge elicitation methodology.

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.

Ackley, D.H., Hinton, G.E., & Sejnowski. T.J. (1985). A learning algorithm for Boltzmann machines. *Cognitive Science*, 9, 147-169.

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. (1983). *Corporate financial distress*. New York: John Wiley.

Altman, E.I., & McGough, T.P. (1974). Evaluation of a company as a going concern. *The Journal of Accountancy*, 43(1), 113-122.

Altman, E.I. & Katz, S. (1976). Statistical bond rating classification using financial and accounting data. In Schiff, M. & Sorter, G. (Eds.), *Topical research in accounting*. New York: NYU Press.

Altman, E.I., Haldeman, R.G,. & Narayanan, P. (1977). ZETA analysis: A new model to identify bankruptcy risk of corporations. *Journal of Banking and Finance*, 1, 29-54.

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.

Altman, E.I., Marco, G., & Varetto, F. (1994). Corporate distress diagnosis: Comparisons using linear discriminant analysis and neural networks. *Journal of Banking and Finance*, 18, 505-529.

Anderson, J.A. (1977). Neural models with cognitive implications. In LaBerse, D., & Samuels, S.J. (Eds.), *Basic processes in reading: Perception and comprehension*. Hillsdale, NJ.: Erlbaum.

Anderson, M.J. (1985). Some evidence on the effect of verbalization on process: A methodological note. *Journal of Accounting Research*, 23(2), 843-852.

Anderson, M.J. (1988). A comparative analysis of information search and evaluation behavior of professional and non-professional financial analysts. *Accounting, Organization and Society*, 13,431-446.

Argenti, J. (1976). *Corporate collapse. The causes and symptoms*. London: McGraw-Hill.

Ashby, F.G. (1992). Multidimensional models of categorization. In Ashby, F.G. (Ed.), *Multidimensional models of perception and cognition* (pp. 449-483). Hillsdale, NJ: Erlbaum.

Ashton, R.H. (1982). *Human information processing in accounting*. Sarasota, FL: American Accounting Association.

Baestaens, D.E., & van den Bergh, W.M. (1995). Tracking the Amsterdam Stock Index using neural networks. In Refenes, A.P. (Ed.), *Neural networks in the capital markets* (pp. 148-161). New York: Wiley.

Barker, D., (1990). Analyzing financial health: Integrating neural networks and expert systems. In Trippi, R.R., & Turban, E. (Eds.), *Neural Networks in Finance and Investing* (pp. 85-102). Chicago, IL: Probus Publishing.

Barnes, P. & Huan, H.D. (1993). The auditor's going concern decision: Some UK evidence concerning independence and competence. *Journal of Business, Finance and Accounting*, 20(2), 213-228.

Barto, A.G. & Anandan, P. (1985). Pattern recognizing stochastic learning automata. *IEEE Transactions on Man, Systems and Cybernetics*, 15, 360-375.

Bassøe, C.F. (1995). Automated diagnoses from clinical narratives: A medical system based on computerized medical records, natural language processing and neural network technology. *Neural Networks*, 8(2), 313-319.

Baum, E.B., & Haussler, D. (1989). What size net gives valid generalization ? *Neural Computation*, 1, 151-160.

Beaver, W.H., (1966). Financial ratios as predictors of failure. *Journal of Accounting Research*, supplement 1966a, 71-111.

Bechtel, W. (1993). Currents in connectionism. *Minds and Machines*, 3, 125-153.

Bedard, J. (1989). Expertise in auditing: Myth or reality. *Accounting, Organizations and Society*, 14(1/2), 113-131.

Bedard, J. & Biggs, S.F. (1991). Pattern recognition, hypotheses generation, and auditor performance in an analytical task. *The Accounting Review*, 66(3), 622-642.

Bedard, J. & Chi, T.H. (1993). Expertise in auditing. *Auditing: A Journal of Practice & Theory*, 12(supplement), 21-45.

Belkaoui, A. (1983). *Industrial bonds and the rating process*. Westport, CT: Quorum Books.

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.

Bhandari, S., Soldofsky, R., & Boe, J. (1979). Discriminant analysis model for predicting bond rating changes of industrial bonds. *Financial Management*, 8, 74-81.

Biggs, S.F., & Mock, T.J. (1983). An investigation of auditor decision processes in the evaluation of internal controls and audit scope decisions. *Journal of Accounting Research*, 21(1), 234-255.

Biggs, S.F. (1984). Financial analysts information search in the assessment of corporate earnings power. *Accounting, Organizations and Society*, 9(3/4), 313-323.

Biggs, S.F., & Wild, J.J. (1985). An investigation of auditor judgment in analytical review. *The Accounting Review*, 60(4), 607-633.

Biggs, S.F., Bedard, J., Gaber, B.G., & Linsmeier, T.J. (1985). The effects of task size and similarity on the decision behavior of bank loan officers. *Management Science*, 31, 970-987.

Biggs, S.F., Messier, W.F., & Hansen, J.V. (1987). A descriptive analysis of computer audit specialists decision-making behavior in advanced computer environments. *Auditing: A Journal of Practice and Theory*, 6(2), 1-21.

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.

Binks, D.L., & Allinson, N.M. (1991). Financial data recognition and prediction using neural networks. *Artificial Neural Networks: Proceedings of the ICANN-91,* (pp. 1709-1712). Amsterdam: North-Holland.

Blocher, E., & Cooper, J.C. (1988). A study of auditors analytical review performance. *Auditing: A Journal of Practice and Theory*, 7(2), 1-28.

Bonner, S.E. (1990). Experience effects in auditing: The role of task-specific knowledge. *The Accounting Review*, 65, 72-92.

Bonner, S.E. (1991). Is experience necessary in cue measurement? The case of auditing tasks. *Contemporary Accounting Research,* 8, 253-269.

Bonner, S. E., & Pennington, N. (1991). Cognitive processes and knowledge as determinants of auditor expertise. *Journal of Accounting Literature,* 10, 1-50.

Bounds, D.G., Lloyd, P.J., & Mathew, B.G. (1990). A comparison of neural network and othet pattern recognition approaches to the diagnosis of low back disorders. *Neural Networks*, 3, 583-591.

Bouwman, M.J. (1982). The use of accounting information: Expert versus novice behavior. In Ungson, G.R., & Braunstein, D.N. (Eds.), *Decision making: An interdisciplinary inquiry* (pp. 134-167). Boston, MA: Kent.

Bouwman, M.J. (1983). Human diagnostic reasoning by computer: An illustration from financial analysis. *Management Science*, 29, 653-670.

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.

Bremner, F.J., Gotts, S.J., & Denham, D.L. (1994). Hinton diagrams: Viewing connection strengths in neural networks. *Behavior Research Methods, Instruments and Computers*, 26, 215-218.

Brooks, L.R. (1978) Non-analytic concept formation and memory for instances. In Rosch, E. & Lloyd, B. (Eds.), *Cognition and categorization* (pp. 169-211). Hillsdale, NJ: Erlbaum.

Brooks, L.R, Norman, G.R., & Allen, S.W. (1991). Role of specific similarity in a medical diagnosis task. *Journal of Experimental Psychology: General*, 120, 278-287.

Brown, C.E., & Salomon, I. (1991). Configural information processing: The role of domain-specific knowledge. *The Accounting Review*, 66, 100-119.

Brown, T.J. (1992). Schemata in consumer research: A connectionist approach. *Advances in Consumer Research*, 19, 787-794.

Bruner, J., Goodnow, J.J., & Austin, G.A. (1956). *A study of thinking*. New York: Wiley.

Brunswik, E. (1952). *The conceptual framework of psychology.* Chicago, IL: University of Chicago Press.

Buta, P. (1994). Mining for financial knowledge with case based reasoning. *AI Expert,* February 1994, 34-41.

Campbell, J. E. (1984). An application of protocol analysis to the "little GAAP" controversy. *Accounting, Organizations and Society,* 9, 329-342.

Carpenter, G.A. (1989). Neural network models for pattern recognition and associative memory. *Neural Networks*, 2, 243-257.

Carpenter, G.A., & Grossberg, S. (1987). ART2: Self-organization of stable category recognition codes for analog input patterns. *Applied Optics*, 26, 4919-4930.

Carpenter, G.A., & Grossberg, S. (1990). ART3: Self-organization of distributed pattern recognition codes in neural network hierarchies. *Proceedings of the INNC-90* (pp. 801-804), Dordrecht: Kluwer.

Casey, C. (1980a). The usefulness of accounting ratios for subjects' predictions of corporate failure: Replication and extensions. *Journal of Accounting Research*, 18, 603-613.

Casey, C. (1980b). Variation in accounting information load: The effect on loan officers predictions of bankruptcy. *The Accounting Review*, 55, 36-49.

Chakaborty, K., Mehrotra, K., Mohan, C.K., & Ranka, S. (1992). Forecasting the behavior of multivariate time series using neural networks. *Neural Networks*, 5, 961-970.

Chalmers, J. (1990). Syntactic transformations on distributed representations. *Connection Science*, 2, 53-62.

Chalos, P., & Pickard, S. (1985). Information choice and cue use: An experiment in group information processing. *Journal of Applied Psychology*, 70, 634-641.

Chalos, P. (1985). Financial distress: A comparative study of individual, model and committee assessments. *Journal of Accounting Research*, 23, 527-543.

Chandrasekaran, B., & Goel, A. (1988). From numbers to symbols to knowledge structures: Artificial intelligence perspectives on the classification task. *IEEE Transactions on Systems, Man and Cybernetics*, 18, 415-424.

Chang, C.F., Sheu B.J., & Thomas, J. (1993). Multi layered back propagation neural networks for finance analysis. *Proceedings of the World Congress on Neural Networks* (445-450). New York: IEEE.

Cheng, B., & Titterington, D.M. (1994). Neural networks: A review from a statistical perspective. *Statistical Science*, 9(1), 2-54.

Chesser, D. (1974). Predicting loan noncompliance. *Journal of Commercial Bank Lending*, 56(8), 28-38.

Chewning, E.G., & Harrell, A.M. (1990). The effect of information load on decision makers' cue utilization levels and decision quality in a financial distress decision task. *Accounting, Organizations and Society*, 15, 527-542.

Choo, F. (1989). Expert-novice differences in judgement/decision making research. *Journal of Accounting Literature,* 8, 106-136.

Choo, F., & Trotman, K.T. (1991). The relationship between knowledge structure and judgements for experienced and inexperienced auditors. *The Accounting Review*, 66, 464-485.

Clark, A. (1989) *Microcognition.* Cambridge, MA: MIT Press.

Clark, A. (1993). *Associative engines. Connectionism, concepts and representational change*. Cambridge, MA: MIT Press.

Coats, P.K., & Fant, L.F. (1993). Recognizing financial distress patterns using a neural network tool. *Financial Management,* 22, 142-156.

Collins, A.M., & Lofthus, E.F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82, 407-428.

Collins, E., Gosh, S., & Scofield, C. (1988). An application of a multiple neural network learning system to emulation of mortgage underwriting judgements. *Proceedings of the IEEE International Conference on Neural Networks* (II459-II466).

Cook, T.D., & Campbell, D.T. (1979). *Quasi-experimentation: Design and analysis issues for field settings*. Boston, MA: Houghton Mifflin.

Cowan, J.D., & Sharp, D.H. (1988). Neural nets and artificial intelligence. *Dædalus*, 117(1), 85-121.

Cronbach, L.J. (1951) Coefficient alpha and the internal structure of tests. *Psychometrica*, 16, 297-334.

Crutcher, R.J. (1994). Telling what we know: The use of verbal report methodologies in psychological research. *Psychological Science*, 5(5), 241-244.

Cummins, R. (1989). *Meaning and Mental Representation.* Cambridge, MA: MIT Press.

Cybenko, G. (1989). Approximation by superpositions of a sigmoid function, *Mathematics of Control, Signals and Systems*, 2, 303-314.

Cyert, R.M., & Grunberg, E. (1963). Assumption, prediction and explanation in economics. In Cyert, R.M. & March, J.G. (Eds.), *A Behavioral Theory of the Firm.* (pp. 298-311), Englewood Cliffs, NJ: Prentice-Hall.

Dagli, C.H., Burke, L.I., & Shin, Y.C. (1992). Intelligent engineering systems through artificial neural networks : *Proceedings of the Artificial Neural Networks in Engineering (ANNIE '92) conference*. New York: ASME Press.

Danos, P., & Imhoff, E.A. (1983). Factors affecting auditors evaluation of forecasts. *Journal of Accounting Research*, 21, 473-494.

Danos, P., Holt, D.L., & Imhoff jr., E.A. (1984). Bond raters use of management financial forecasts: Experiment in expert judgement. *The Accounting Review*, 59, 547-573.

Danos, P., Holt, D.L., & Imhoff jr., E.A. (1989). The use of accounting information in bank lending decisions. *Accounting, Organizations and Society*, 14, 235-246.

Deng, P.S. (1993). Automating knowledge acquisition and refinement for decision support: A connectionist inductive inference model. *Decision Sciences*, 24, 371-393.

Dennett, D.C. (1978). *Brainstorms, philosophical essays on mind and psychology*. Sussex, England: Harvester Press.

Diamond, C., Shadbolt, J., Barac, M.A., & Refenes, A.P. (1993). Neural network system for tactical asset allocation in the global bonds markets. *Proceedings of the 3rd International Conference on Artificial Neural Networks* (pp. 118-122), Stevenage, England: IEEE

Dietrich, J.R,. & Kaplan, R.S. (1982). Empirical analysis of the commercial loan classification decision. *The Accounting Review*, 57, 18-38.

Doukas, J. (1986). Bankers versus bankruptcy prediction models: An empirical investigation, 1979-82. *Applied Economics*, 18, 479-493.

Dreyfus, H.L. (1972). *What computers can't do: A critique of artificial reason*. New York: Harper & Row.

Dretske, F. (1981). *Knowledge and the flow of information*. Oxford: Basil Blackwell.

Dun & Bradstreet-Soliditet. (1995). *Regnskapsfakta: Hovedtall for alle registrerte regnskap*. Oslo, Norway: Dun & Bradstreet - Soliditet A/S.

Dutta, S., & Shekhar, S. (1988). Bond rating: A non-conservative application of neural networks. *Proceedings of the IEEE International Conference on Neural Networks* (pp. II443- II450), San Diego, CA: IEEE

Dutta, S., Shekhar, S., & Wong, W.Y. (1994). Decision support in non-conservative domains: Generalizations with neural networks. *Decision Support Systems*, 11, 527-544.

Eberhart, R.C., & Dobbins, R.W. (Eds.). (1990). *Neural network PC tools: A practical guide*. San Diego, CA: Academic Press.

Eklund, T., & Knutsen, K. (1994). *Regnskapsanalyse: Aktiv bruk av regnskapet*. Oslo, Norway: Universitetsforlaget.

Edmister, R.O. (1971). Financial ratios and credit scoring for small business loans. *Journal of Commercial Bank Lending*, 53, 10-23.

Efron, B., & Gong, G. (1983). A leisurely look at the bootstrap, the jackknife, and cross-validation. *The American Statistician*, 37, 36-48.

Efron, B., & Tibshirani, R.J. (1993). *An introduction to the bootstrap*. New York: Chapmann and Hall.

Ellis, H.C., & Hunt, R.R. (1993). *Fundamentals of cognitive psychology* (5. ed.). Madison, WI: Brown & Benchmark.

Elman, J.L. (1990). Finding structure in time. *Cognitive Science*, 14, 179-211.

Elstein, A.S., Shulman, L.S., & Sprafka, S.A. (1990). Medical problem solving: A ten-year perspective. *Evaluation & The Health Professions*, 13, 5-36.

Enis, C.R. (1988). The impact of current-valued data on the predictive judgements of investors. *Accounting, Organizations and Society*, 13, 123-145.

Ericsson, K.A., & Simon, H.A. (1984). *Protocol analysis: Verbal reports as data.* Cambridge, MA: MIT Press.

Erxleben, K., Baetge, J., Feidicker, M., Koch, H., Krause, C., & Mertens, P. (1992). Klassifikation von Unternehmen: Ein Vergleich von Neuronalen Netzen und Diskriminanzanalyse. *Zeitschrift für Betriebswirtschaft*, 62, 1237-1262.

Estes, W.K. (1950). Towards a statistical theory of learning. *Psychological Review*, 57, 94-107.

Estes, W.K. (1986). Array models for category learning, *Cognitive Psychology*, 18, 500-549.

Estes, W.K. (1993). Models of categorization and category learning. *The Psychology of Learning and Motivation*, 29, 15-56.

Estes, W.K. (1994). *Classification and cognition.* Oxford: Oxford Univ. Press.

Estes, W.K., Campbell, J.A., Hastapoulos, N., & Hurwitz, J.B. (1989). Base-rate effects in category learning: A comparison of parallel networks and memory storage-retrieval models. *Journal of Experimental Psychology, Learning, Memory and Cognition*, 15, 556-571.

Fahlman, S. (1989). Faster-learning variations of backpropagation: An empirical study. In Touretzky, D., Hinton, G., & Sejnowski, T. (Eds.), *Proceeding of the 1988 Connectionist Summer School* (pp. 38-51). San Mateo, CA: M. Kaufmann.

Fahlman, S., & Lebiere, C. (1990). The cascade correlation learning architecture. In, Touretzky, D. (Ed.), *Advances in Neural Information Processing Systems* , 2 (pp. 524-532). San Mateo, CA: M. Kaufmann.

Falbo, P. (1991). Credit scoring by enlarged discriminant models. *OMEGA International Journal of Management Science*, 19, 275-289.

Farelly, G.E., Ferris, K.R., & Reichenstein, W.R. (1985). Perceived risk, market risk, and accounting determined risk measures. *The Accounting Review*, 60, 278-288.

Fausett, L. (1994). *Fundamentals of neural networks: Architectures, algorithms and applications*. Englewood Cliffs, NJ.: Prentice-Hall.

Feldman, J.A., & Ballard, D.H. (1982). Connectionist models and their properties. *Cognitive Science* , 6, 205-254.

Fisher, L. (1959). The determinants of risk premiums on corporate bonds. *Journal of Political Economy*, June 1959, 217-237.

Fodor, J. (1980). Methodological solipsism considered as a research strategy in cognitive psychology. *The Behavioral and Brain Sciences*, 3, 63-73.

Fodor, J. (1983). *Modularity of mind*. Cambridge, MA: MIT Press.

Fodor, J., & Pylyshyn, Z. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28, 3-71.

Foster, G. (1986). *Financial statement analysis* (2. ed..). Englewood Cliffs, NJ: Prentice Hall.

Frederick, D.M. (1991). Auditors' representation and retrieval of internal control knowledge. *The Accounting Review*, 66, 240-258.

Friedman, J.H. (1991). Multivariate adaptive regression splines, *Ann. Statist.*, 19, 1-141.

Friedman, J.H., & Stuetzle, W. (1981). Projection pursuit regression. *Journal of the American Statistical Association*, 76, 817-823.

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.

Fukushima, K. (1975). Cognitron: A self-organizing multilayered neural network. *Biological Cybernetics*, 20, 121-136.

Gallant, S. (1993). *Neural network learning and expert systems.* Cambridge MA.: MIT Press.

Garner, W.R. (1978). Aspects of a stimulus: Features, dimensions and configurations. In Rosch, E., & Lloyd, B.B. (Eds.), *Cognition and categorization* (pp. 99-133). Hillsdale, NJ: L. Erlbaum.

Garson, G.D. (1991). Interpreting neural network connection weights. *AI Expert*, april 1991, 47-51.

Geman, S., Bienenstock, E., & Doursat, R. (1992). Neural networks and the bias/variance dilemma. *Neural Computation*, 4, 1-58.

Gentry, J., Newbold, P., & Whitford, D. (1985). Classifying bankrupt firms with funds flow components. *Journal of Accounting Research*, spring 1985, 146-159.

Gibbins, M., & Jamal, K. (1993). Problem-centred research and knowledge-based theory in the professional accounting setting. *Accounting, Organizations and Society*, 18, 451-466.

Gilbert, L.R., Menon, K., & Schwartz, K.B. (1990). Predicting bankruptcy for firms in financial distress. *Journal of Business Finance & Accounting*, 17, 161-171.

Gluck, M.A. (1991). Stimulus generalization and representation in adaptive network models of category learning. *Psychological Science*, 2, 50-55.

Gluck, M.A. (1992). Stimulus sampling and distributed representations in adaptive network theories of learning. In Healey, A.F., Kosslyn, S.M., & Shiffrin, R.M. (Eds.) *From learning processes to cognitive processes* 2 (pp. 169-199). Hillsdale, NJ: L. Erlbaum.

Gluck, M.A., & Bower, G.H. (1988a). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, 117, 227-247.

Gluck, M.A., & Bower, G.H. (1988b). Evaluating an adaptive network model of human learning. *Journal of Memory and Language*, 27,166-195.

Gluck, M.A., & Rumelhart, D.E. (1990). *Neuroscience and connectionist theory*. Hillsdale, NJ: L. Erlbaum.

Goldberg, D.E. (1989). *Genetic algorithms in search, optimization and machine learning.* Reading, MA: Addison-Wesley.

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.

Govindarajan, V. (1980). The objectives of financial statements: An empirical study of the use of cash flow and earnings by security analysts. *Accounting, Organizations and Society*, 5, 383-392.

Graesser, A.C. (1981). *Prose comprehension beyond the word.* New York: Springer.

Grossberg, S. (1982). *Studies of mind and brain; Neural principles of learning, perception, cognition and motor control.* Dordrecht: Reidel.

Grossberg, S., & Gutowski, A.E. (1987). Neural dynamics of decision making under risk: Affective balance and cognitive-emotional interactions. *Psychological Review*, 94, 300-318.

Grudnitski, G., & Osburn, L. (1993). Forecasting S&P and gold futures prices: An application of neural networks. *The Journal of Futures Markets*, 13, 631-543.

Grudnitski, G., & Do, A.Q. (1995). Important factors in neural networks' forecasts of gold futures prices. In Refenes, A.P. (Ed.), *Neural networks in the capital markets* (pp. 163-173). New York: Wiley.

Hampson, S. (1990). *Connectionist problem solving.* Boston, MA.: Birkhauser.

Hansen, J.V., Koehler, G.J., Messier, W.F., & Mutchler, J.F. (1993). Developing knowledge structures: A comparison of a qualitative-response model and two machine-learning algorithms. *Decision Support Systems*, 10, 235-243.

Hanson, S.J., & Burr, D.J. (1990). What connectionist models learn: Learning and representation in connectionist networks. *Behavioral and Brain Sciences*, 13, 471-518.

Haugeland, J. (1985). *Artificial intelligence, the very idea.* Cambridge, MA: MIT Press.

Hawkins, D.F. (1986). *Corporate financial reporting and analysis: Text and cases.* Homewood, IL.: Irvin.

Hawley, D.D., Johnson, J.D., & Raina, D. (1990). Artificial neural systems: A new tool for financial decision making. *Financial Analyst Journal*, November/December 1990, 63-72.

Haykin, S. (1994). *Neural networks: A comprehensive foundation* . New York: Macmillian.

Hebb, D.O. (1949). *The organization of behavior.* New York: Wiley.

Hecht-Nielsen, R. (1987). Counterpropagation networks. *Applied Optics*, 26, 4979 - 4984.

Hertz, J., Krogh, A., & Palmer, R.G. (1991). *Introduction to the theory of neural computation.* Reading, MA: Addison-Wesley.

Hinton, G.E. (1989). Connectionist learning procedures. *Artificial Intelligence*, 40, 185-234.

Hinton, G.E., & Anderson, J.A. (Eds.). (1981). *Parallel models of associative memory.* Hillsdale, NJ: Erlbaum.

Hinton, G.E., & Sejnowski, T.J. (1986). Learning and relearning in Boltzmann machines. In Rumelhart, D.E. & McClelland, J.L. (Eds.), *Parallel distributed processing; Explorations in the microstructure of cognition.* Cambridge, MA.: MIT Press.

Hinton, G., McClelland, J.L., & Rumelhart, D.E. (1986). Distributed representations. In Rumelhart, D.E. & McClelland, J.L. (Eds.), *Parallel distributed processing; Explorations in the microstructure of cognition* (pp. 77-109). Cambridge, MA.: MIT Press.

Hintzman, D.L. (1988). Judgements of frequency and recognition memory in a multiple-trace memory model. *Psychological Review*, 95, 528-551.

Ho, J., & Rodgers, W. (1993). A review of accounting research on cognitive characteristics. *Journal of Accounting Literature.* 12, 101-130.

Hogarth, R.M. (1987). *Judgement and choice: The psychology of decision* (2. ed.). Chichester: Wiley.

Holland, J.H., Holyoak, K.J., Nisbett, R.E., & Thagard, P.R. (1986). *Induction: Processes of inference, learning and discovery.* Cambridge, MA.: MIT Press.

Holt, R.N., & Carroll, R.J. (1980). Classification of commercial bank loans through policy capturing. *Accounting, Organizations and Society*, 5, 285-296.

Holt, D.L., & Morrow, P.C. (1992). Risk assessment judgements of auditors and bank lenders: A comparative analysis of conformance to Bayes theorem. *Accounting, Organizations and Society*, 17, 549-559.

Holyoak, K.J. (1990). Problem solving. In Osherson, D.N., & Smith, E.E. (Eds.), *Thinking: An invitation to cognitive science* (pp. 117-146). Cambridge, MA: MIT Press

Hopfield, J.J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proc. of the National Academy of Sciences*, 79, 2554-2558.

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.

Hornik, K., Stinchcombe, M., & White, H. (1989). Multi-layer feedforward networks are universal approximators. *Neural Networks*, 2, 359-366.

Horrigan, J. (1966). The determinants of long term credit standing with financial ratios. *Journal of Accounting Research*, 4 (supplement), 44-62.

Houghton, K.A., & Sengupta, R. (1984). The effect of prior probability disclosure and the information set construction on bankers' ability to predict failure. *Journal of Accounting Research*, 22, 768-775.

Houghton, K.A., & Woodliff, D.R. (1987). Financial ratios: The prediction of corporate "success" and failure. *Journal of Business Finance & Accounting*, 14, 537-554.

Houghton, K.A. (1984). Accounting data and the prediction of business failure: The setting of priors and the age of data. *Journal of Accounting Research*, 22, 361-368.

Hruschka, H. (1993). Determining market response functions by neural network modelling: A comparison to econometric techniques. *European Journal of Operations Research*, 66, 27-35.

Hsu, W., Hsu, L.S., & Tenorio, M.F. (1994). Parameter significance estimation and financial prediction. *Neural Computing & Applications*, 1, 280-286.

Hurwitz, J.B. (1990). *A hidden-pattern unit network model of category learning*. Unpublished doctoral dissertation, Harvard University.

Hurwitz, J.B. (1994). Retrieval of exemplar and feature information in category learning. *Journal of Experimental Psychology, Learning, Memory and Cognition*, 20, 887-903.

Hutchinson, J.M., Lo, A.W., & Poggio, T. (1994). A nonparametric approach to pricing and hedging derivative securities via learning networks, *The Journal of Finance*, 49, 851-889.

Iselin, E.R. (1991). Individual versus group decision-making performance: A further investigation of two theories in a bankruptcy prediction task. *Journal of Business, Finance & Accounting*, 18,191-208.

Iselin, E.R. (1993). The effects of the information and data properties of financial ratios and statements on managerial decision quality. *Journal of Business Finance & Accounting*, 20, 249-266.

Jacobs, R.A. (1988). Increased rates of convergence through learning rate adaption. *Neural Networks*, 1, 295-308.

James, H. (1994). Software for studying and developing applications of artificial neural networks. *Economic Journal*, 104, 181-196.

Jensen, H.L. (1992). Using neural networks for credit scoring. *Managerial Finance*, 18, 15-26.

Johansson, E.M., Dowla, F.U., & Goodman, D.M. (1992). Backpropagation learning for multilayer feed-forward neural networks using the conjugate gradient method. *International Journal of Neural Systems*, 2, 291-301.

Jones, M.C., & Sibson, R. (1987). What is projection pursuit ? *Journal of the Royal Statistical Society*, 150 (1), 1-36.

Jones, F.L. (1987). Current techniques in bankruptcy prediction. *Journal of Accounting Literature*, 6, 131-164.

Jung, D., & Burns, J.R. (1993). Connectionist approaches to inexact reasoning and learning systems for executive and decision support. *Decision Support Systems*, 10, 37-66.

Jøreskog, K.G., & Sørbom, D. (1988). *LISREL 7: User's guide*. Chicago: National Educational Resources.

Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgement under uncertainty: Heuristics and biases*. Cambridge: Cambridge Univ. Press.

Kamijo, K., & Tanigawa, T. (1990). Stock price pattern recognition: A recurrent neural network approach. Proceedings of the IEEE International Joint Conference on Neural

Networks. Reprinted in Trippi R.R., & Turban E. (Eds.), *Neural Networks in Finance and Investing* (pp. 357-369).Chicago, IL.: Probus Publishing.

Kaplan, R.S. & Urwitz, G. (1979). Statistical models of bond ratings: A methodological inquiry, *Journal of Business*, 52, 231-261.

Karels, G.V., & Prakash, A.J. (1987). Multivariate normality and forecasting of business bankruptcy. *Journal of Business Finance & Accounting*, 14, 573-593.

Karnin, E.D. (1990). A simple procedure for pruning back-propagation trained neural networks. *IEEE Transactions on Neural Networks*, 1, 239-242.

Kida, T. (1980). An investigation into auditors' continuity and related qualification judgements, *Journal of Accounting Research*, 18, 506-523.

Kida, T. (1984). The impact of hypothesis-testing strategies on auditors use of judgement data, *Journal of Accounting Research*, 22, 332-340.

Kim, J.W., Weistroffer, H.R., & Redmond, R.T. (1993). Expert systems for bond rating: A comparative analysis of statistical; rule-based and neural network systems. *Expert Systems*, 10, 167-171.

Kimoto, T., Asakawa, K., Yoda, M., & Takeoka, M. (1990). Stock market prediction system with modular neural networks, Proceedings of the IEEE International Joint Conference on Neural Networks. Reprinted in Trippi, R.R., & Turban, E. (1993). *Neural networks in finance and investing* (pp. 343- 356). Chicago, IL: Probus Publishing.

Kinney, W.R. (1987). Attention-directing analytical review using accounting ratios: A case study. *Auditing: A Journal of Practice & Theory*, 6, 59-73.

Kinserdal, A. (1992). *Regnskap med analyse*. Oslo, Norway: Bedriftsøkonomens forlag.

Kirkebøen, G. (1993). *Psykologi, informasjonsteknologi og ekspertise: En studie av informasjonsbehandlingsteoriens innflytelse i psykologien - og vice versa.* Unpublished doctoral dissertation, Universitetet i Oslo.

Klein, M., & Methlie, L.B. (1990). *Expert systems: A decision support approach.* Reading, MA: Addison-Wesley.

Klersey, G.F., & Mock, T.J. (1989). Verbal protocol research in auditing. *Accounting, Organizations and Society*, 14, 133-151.

Koh, H.C., & Killough, L.N. (1990). The use of multiple discriminant analysis in the assessment of the going-concern status of an audit client. *Journal of Business Finance & Accounting*, 17, 179-192.

Kohonen, T. (1977). *Associative memory: A system theoretic approach.* New York: Springer.

Kohonen, T. (1995). *Self-organizing maps*. Berlin: Springer.

Kosko, B. (1987). Competitive bidirectional associative memories, *IEEE First International Conference on Neural Networks* (pp. II759-766). San Diego, CA: SOS Printing.

Kosslyn, S.M. (1980). *Image and mind.* Cambridge MA.: Harvard University Press.

Kruschke, J.K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, 99, 22-44.

Kruschke, J.K. (1993a). Three principles for models of category learning. *The Psychology of Learning and Motivation*, 29, 57-90.

Kruschke, J.K. (1993b). Human category learning: Implications for backpropagation models. *Connection Science*, 5, 3-36.

Kryzanowski, L., Galler, M., & Wright, D.W. (1993). Using artificial neural networks to pick stocks. *Financial Analyst Journal*, July-August 1993, 21-27.

Kuan, C.M., & White, H. (1994). Artificial neural networks: An econometric perspective. *Econometric Reviews*, 13(1), 1-91.

Kvadsheim, R. (1992). *The intelligent imitator: Towards an exemplar theory of behavioral choice*. Amsterdam: North-Holland.

Lachenbruch, P., & Mickey, M. (1968). Estimation of error rates in discriminant analysis. *Technometrics*, 10, 1-11.

Lachenbruch, P. (1967). An almost unbiased method of obtaining confidence intervals for the probability of misclassification in discriminant analysis. *Biometrics*, 23, 639-645.

Laitinen, E.K. (1991). Financial ratios and different failure processes. *Journal of Business Finance & Accounting*, 18, 649-673.

Lakoff, G. (1987). *Women, fire and dangerous things, What categories reveal about the mind.* Chicago: University of Chicago Press.

Le Cun, Y. (1985). A learning scheme for asymmetric threshold networks. *Proceedings Cognitiva* (pp. 599-604) Paris, France.

Le Cun, Y. (1989). Generalization and network design strategies. In Pfeifer, R., Schreter, Z., Fogelman-Soulie, F., & Steels, L., (Eds.), *Connectionism in perspective* (pp. 143-155). Amsterdam: North-Holland.

Lev, B., & Thiagarajan, S.R. (1993). Fundamental information analysis. *Journal of Accounting Research*, 31, 190-215.

Levine, D.S., & Aparicio, M. (Eds.). (1994). *Neural networks for knowledge representation and inference*. Hillsdale, NJ: L. Erlbaum.

Lewis, B.L., Patton, J.M., & Green, S.L. (1988). The effects of information choice and information use on analysts' predictions of municipal bond rating changes. *The Accounting Review*, 63, 270-282.

Liang, T.P., Chandler, J.S., Han, I., & Roan, J. (1992). An empirical investigation of some data effects on the classification accuracy of probit, ID3, and neural networks. *Contemporary Accounting Research*, 9, 306-328.

Liang, T.P., Moskowitz, H., & Yih, Y. (1992) Integrating neural networks and semi-Markov processes for automated knowledge acquisition: An application to real-time scheduling. *Decision Sciences*, 23, 1297-1314.

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.

Libby, R., & Blashfield, R.K. (1978). Performance of a composite as a function of the number of judges. *Organizational Behavior and Human Performance*, April 1978, 121-129.

Libby, R., Trotman, K.T., & Zimmer, I. (1987). Member variation, recognition of expertise, and group performance. *Journal of Applied Psychology*, 72(1), 81-87.

Libby, R., & Frederick, D.M. (1990). Experience and the ability to explain audit findings. *Journal of Accounting Research*, 28, 348-367.

Libby, R., & Luft, J. (1993). Determinants of judgement performance in accounting settings: Ability, knowledge, motivation and environment. *Accounting, Organizations and Society*, 18, 435-450.

Lippmann, R.P. (1993). Neural networks, Bayesian a posteriori probabilities, and pattern classification. In Cherkassy, V, Friedman, J.H., & Wechsler, H. (Eds.), *From statistics to neural networks, theory and pattern recognition applications.* Berlin: Springer.

Luce, R.D. (1963). Detection and recognition. In Luce, R.D., Bush, R.B., & Galanter, E. (Eds.), *Handbook of mathematical psychology* (pp. 103-189). New York: Wiley.

Lyngstad, P.B. (1987). *Kunnskapsmodellering av ekspertise i finansiell diagnostisering.* Unpublished masters thesis, Norwegian School of Economics and Business Administration, Bergen.

MacKay, D.J.C. (1992). A practical Bayesian framework for backpropagation networks. *Neural Computation*, 4, 448-472.

Mallaris, M., & Salchenberger, L. (1993). A neural network model for estimating option prices. *Journal of Applied Intelligence*, 3, 193-206.

Mar Molinero, C., & Ezzamel, M. (1991). Multidimensional scaling applied to corporate failure. *OMEGA International Journal of Management Science*, 19, 259-274.

March, J.G. (1978). Bounded rationality, ambiguity and the engineering of choice. *Bell Journal of Economics*, 9, 587-608.

Martin-del-Brio, B., & Serrano-Cinca, C. (1993). Self-organizing neural networks for the analysis and representation of data: Some financial cases. *Neural Computing & Applications*, 1, 193-206.

Massaro, D.W., & Friedman, D. (1990). Models of integration given multiple sources of information. *Psychological Review*, 97, 225-252.

McClelland, J.L. & Rumelhart, D.E. (Eds.). (1986). *Parallel distributed processing; Explorations in the microstructure of cognition.* Cambridge, MA: MIT Press.

McClelland, J.L., & Elman, J.L. (1986). Interactive processes in speech perception: The TRACE model. In McClelland, J.L. & Rumelhart, D.E. (Eds.), *Parallel distributed* *processing; Explorations in the microstructure of cognition*. Cambridge, MA: MIT Press.

McClelland, J.L., & Kawamoto, A.H. (1986). Mechanisms of sentence processing: Assigning roles to constituents. In McClelland, J.L. & Rumelhart, D.E. (Eds.), *Parallel distributed processing; Explorations in the microstructure of cognition*. Cambridge, MA: MIT Press.

McCloskey, M., & Cohen, N.J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. in Bower G. (Ed.), *The psychology of learning and motivation*, 24. New York: Academic Press.

McGee, W., Shields, M.D., & Birnberg, J.G. (1978). The effects of personality on a subject's information processing. *The Accounting Review*, 53, 681-697.

Mear, R., & Firth, M. (1987a). Cue usage and self-insight of financial analysts. *The Accounting Review*, 63, 176-182.

Mear, R., & Firth, M. (1987b). Assessing the accuracy of financial analyst security return predictions. *Accounting, Organizations and Society*, 12, 331-340.

Mear, R., & Firth, M. (1988). Risk perceptions of financial analysis and the use of market and accounting data. *Accounting and Business Research*, 18, 335-340.

Medin, D., & Schaffer, M.M. (1978). Context theory of classification learning. *Psychological Review*, 85, 207-238.

Medin, D., & Schwanenflugel, B.J. (1981). Linear separability in classification learning. *Journal of Experimental Psychology, Learning, Memory and Cognition*, 7, 355-368.

Medin, D.L., Altom, M.W., Edelson, S.M., & Freko, D. (1982). Correlated symptoms and simulated medical classification. *Journal of Experimental Psychology, Learning, Memory and Cognition*, 8, 37-50.

Medin, D.L., & Edelson, S.M. (1988). Problem structure and the use of base rate information from experience. *Journal of Experimental Psychology: General*, 117, 68-85.

Medin, D.L., & Florian, J.E. (1992). Abstraction and selective coding in exemplar-based models of categorization. In Healey, A.F., Kosslyn, S.M., & Shiffrin, R.M. (Eds.), *From learning proceesses to cognitive processes* vol 2 (pp. 207-234). Hillsdale, NJ: L. Erlbaum.

Mehta, M. (1995). Foreign exchange markets. In Refenes A.P. (Ed.), *Neural networks in the capital markets* (pp.177-198). New York: Wiley.

Mensah, Y. (1983). The differential bankruptcy predictive ability of specific price level adjustments: Some empirical evidence. *The Accounting Review*, April 1983, 228-246.

Meservy, R.D., Bailey, A.D., & Johnson, P.E. (1986). Internal control evaluation: A computational model of the review process. *Auditing: A Journal of Practice and Theory*, 6, 44-74.

Messier, W.F., & Hansen, J.V. (1988). Inducing rules for expert system development: An example using default and bankruptcy data. *Management Science*, 34, 1403- 1415.

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.

Methlie, L.B. (1993). Performance modelling: A cognitive approach to building knowledge-based systems. In Watkins, P.R., & Eliot, L.B. (Eds.), *Expert systems in business and finance: Issues and applications*. New York: Wiley.

Methlie, L.B. (1994). Regnskapsanalyse i et kognitivt perspektiv. In *Regnskap og ledelse* (pp. 335-355). Oslo, Norway: Bedriftsøkonomens forlag.

Minsky, M. (1975). A framework for representing knowledge. In Winston, P. (Ed.), *The psychology of computer vision*. New York: Mc Graw Hill.

Minsky, M., & Papert, S. (1969). *Perceptrons*. Cambridge, MA: MIT Press.

Mintzberg, H., Raisinghani, H., & Theoret, A. (1976). The structure of unstructured decision processes. *Administrative Science Quarterly*, 21, 246-275.

Moody, J.E. (1993). Prediction risk and architecture selection for neural networks. In Cherkassy, V., Friedman, J.H., & Wechsler, H. (Eds.), *From statistics to neural networks, theory and pattern recognition applications*. Berlin: Springer.

Moody, J.E., & Darken, C. (1988). Learning with localized receptive fields. In Touretzky, D., Hinton, G., & Sejnowski, T. (Eds.), *Proceedings of the 1988 Connectionist Summer School* (pp. 174-185). San Mateo, CA: Morgan Kaufmann.

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.

Moriarty, S. (1979). Communicating financial information through multidimensional graphics. *Journal of Accounting Research*, spring 1979, 205-224.

Mumma, G.H. (1993). Categorization and rule induction in clinical diagnosis and assessment. *The Psychology of Learning and Motivation*, 29, 283-326.

Mutchler, J.F. (1985). A multivariate analysis of auditor's going-concern opinion decision. *Journal of Accounting Research*, 23, 668-682.

Myers, J.L., Lohmeier, J.H., & Well, A.D. (1994). Modelling probabilistic categorization data: Exemplar memory and connectionist nets. *Psychological Science*, 5, 83-89.

Nelson, M.W. (1993). The effects of error frequency and accounting knowledge on error diagnosis in analytical review. *The Accounting Review*, 68, 804-824.

Nesvik, G.O. (1993). *An empirical study of selected learning algorithms for feed-forward neural networks*. Unpublished doctoral dissertation, NTH: Univeristy of Trondheim.

Neural Ware (1993a). *Neural computing: A technology handbook for Professional II/Plus*. Pittsburg, PA.: Neural Ware Inc.

Neural Ware (1993b). *Reference guide: Software reference for Professional II/Plus*. Pittsburg, PA.: Neural Ware Inc.

NeuralWare (1993c). *Advanced reference guide: Software reference for Professional II/Plus.* Pittsburg, PA: Neural Ware Inc.

Newell, A. (1980). Physical symbol systems, *Cognitive Science*, 4, 135-183.

Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.

Newell, A., & Simon, H.A. (1972) *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.

Nisbett, R., & Wilson, T. de C. (1977). Telling more than we can know: Verbal reports on mental processes, *Psychological Review*, 21, 227-236.

Nisbett, R., & Ross, L. (1980). *Human inference: Strategies and shortcomings of social life*. Englewood Cliffs, NJ: Prentice-Hall.

Nosofsky, R.M. (1984). Choice, similarity and the context theory of classification. *Journal of Experimental Psychology, Learning, Memory and Cognition*, 10, 104-114.

Nosofsky, R.M. (1986). Attention, similarity and the identification-categorization relationship. *Journal of Experimental Psychology, General*, 115, 39-57.

Nosofsky, R.M. (1992a). Similarity scaling and cognitive process models. *Annual Review of Psychology*, 42, 25-53.

Nosofsky, R.M. (1992b). Exemplars, prototypes and similarity rules. In Healey, A.F., Kosslyn, S.M., & Shiffrin, R.M. (Eds.), *From learning processes to cognitive processes* 2 (pp. 149-167). Hillsdale, NJ: L. Erlbaum.

Nosofsky, R.M., & Kruschke, J.K. (1992). Investigations of an exemplar-based connectionist model of category learning. *The Psychology of Learning and Motivation*, 28, 207-250.

Nosofsky, R.M., Palmieri, T.J., & McKinley, S.C. (1994). Rule-plus-exception model of classification learning, *Psychological Review*, 10, 53-79.

Nosofsky, R.M., Gluck, M.A., Palmieri, T.J., McKinley, S.C., & Glauthier, P. (1994). Comparing models of rule-based classification learning: A replication and extension of Shepard, Hovland and Jenkins (1961), *Memory & Cognition*, 22, 352-369.

Nottola, C., Condamin, L., & Naim, P. (1991). On the use of hard neural networks for symbolic learning : Application to company evaluation. *Proceedings of the IEEE International Joint Conference on Neural Networks* (pp. 509-522), New York: IEEE.

Nunnally, J.C. (1978). *Psychometric Theory*. New York: McGraw-Hill.

Odom, M.D., & Sharda, R. (1990). A neural network model for bankruptcy prediction. Proceedings of the IEEE International Conference on Neural Networks (pp. II163-II168). Reprinted in Trippi, R.R., & Turban, E. (Eds.), *Neural networks in finance and investing*. Chicago, IL: Probus Publishing.

Ohlson, J.A. (1980). Financial ratios and the probabilistic prediction of bankruptcy. *Journal of Accounting Research*, 18, 109-131.

Oja, E. (1989). Neural networks, principal components, and subspaces. *International Journal of Neural Systems*, 1, 61-68.

Orgler, Y.E. (1970). A credit scoring model for commercial loans. *Journal of Money Credit and Banking*, 2, 434-445.

Osherson, D.N. (1990). Judgment. In Osherson, D.L., & Smith, E.E. (Eds.), *Thinking: An invitation to cognitive science* 3 (pp. 55-87). Cambridge, MA: MIT Press.

Osherson, D.L., & Smith, E.E. (1990). *Thinking: An invitation to cognitive science* 3. Cambridge, MA: MIT Press.

Papert, S. (1988). One AI ore many. *Dædalus*, 117, (1), 1-14.

Park, J., & Sandberg, I.W. (1991). Universal approximation using radial-basis function networks. *Neural Computation*, 3, 246-257.

Parker, D.B. (1985). *Learning logic* (Technical Report TR-47). Cambridge, MA: MIT, Center for Computational Research in Economics and Management Science.

Pedersen, P.E. (1988). *To perspektiver i kognisjonsvitenskap*. Unpublished masters thesis, Norwegian school of economics and business administration, Bergen, Norway.

Pedersen, P.E. (1989). Validitetseffekter av modifikasjoner i det klassiske eksperiment, *Beta*, 3(1), 53-65.

Peng, Y., & Reggia, J.A. (1989). A connectionist model for diagnostic problem solving. *IEEE Transactions on Systems, Man and Cybernetics*, 19, 285-298.

Peters, J.M. (1990). A cognitive computational model of risk hypothesis generation. *Journal of Accounting Research*, 28 (supplement), 83-103.

Peters, J.M. (1993). Decision making, cognitive science and accounting: An overview of the intersection. *Accounting, Organizations and Society*, 18, 383-405.

Pinches, G. & Mingo, K. (1973). A multivariate analysis of industrial bond ratings. *Journal of Finance*, 28, 1-18.

Pinches, G.E., Mingo, K.A., & Caruthers, J.K. (1973). The stability of financial patterns in industrial organizations. *Journal of Finance*, 28, 389-396.

Piramuthu, S., Shaw, M.J., & Gentry, J.A. (1994). A classification approach using multi-layered neural networks. *Decisison Support Systems*, 11, 509-525.

Poddig, T. (1995). Bankruptcy prediction: A comparison with discriminant analysis. In Refenes, A.P. (Ed.), *Neural networks in the capital markets* (pp. 311-323). New York: Wiley.

Poggio, T., & Girosi, F. (1990). Regularization algorithms for learning that are equivalent to multilayer networks. *Science*, 247, 978-982.

Pylyshyn, Z. (1984). *Computation and cognition: Toward a foundation for cognitive science*. Cambridge, MA: MIT Press.

Quinlan, J.R. (1986). Induction of decision trees. *Machine Learning*, 1, 81-106.

Quinlan, P.T. (1991). *Connectionism and psychology: A psychological perspective on new connectionist research*. New York: Harvester.

Raghupathi, W., Schkade, L.L., & Raju, B.S. (1991). A neural network approach to bankruptcy prediction. Proceedings of the IEEE 24th Annual Hawai International Conference on Systems Sciences. Reprinted in Trippi, R.R., & Turban, E. (Eds.), *Neural networks in finance and investing* (pp. 141-158). Chicago, IL: Probus Publishing.

Rahimian, E., Singh, S., Thammachote, T., & Virmani, R. (1993). Bankruptcy Prediction by Neural Networks. In Trippi, R.R., & Turban, E. (Eds.), *Neural networks in finance and investing* (pp. 159-171), Chicago, IL: Probus Publishing.

Raman, K.K. (1981). Financial reporting and municipal bond rating changes. *The Accounting Review*, 56, 910-926.

Reed, S.K. (1972). Pattern recognition and categorization. *Cognitive Psychology*, 3, 382-407.

Reeke, G.N., & Edelman, G.E. (1988). Real brains and artificial intelligence, *Dædalus*, 117 (1), 143-174.

Refenes, A.P. (1993). Constructive learning and its application to currency exchange rate forecasting. In Trippi, R.R., & Turban, E. (Eds.), *Neural networks in finance and investing* (pp. 465-493). Chicago, IL: Probus Publishing.

Refenes, A.P. (Ed.). (1995). *Neural networks in the capital markets.* New York: Wiley.

Refenes, A.P., Azema-Barac, M., Chen, L., & Karoussos, S.A. (1993). Currency exchange rate prediction and neural network design strategies. *Neural Computing & Applications*, 1, 46-58.

Refenes, A.P., Zapranis, A.D., & Francis, G. (1995). Modelling stock returns in the framework of APT: A comparative study with regression models. In Refenes A.P. (Ed.), *Neural networks in the capital markets* (pp. 101-125). New York: Wiley.

Rehkugler, H., & Poddig, T. (1991). Künstliche Neuronale Netze in der Finanzanalyse: Eine neue Ära der Kursprognosen ? *Wirtschaftsinformatik*, 33, 365-374.

Reiter, S.A., & Emery, D.R. (1991). Estimation issues in bond-rating models. *Advances in Quantitative Analysis of Finance and Accounting*, 1 (part A), 147-163.

Rescorla, R.A., & Wagner, A.R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In Black, A.H., & Prokasey, W.F. (Eds.), *Classical conditioning: II. Current research and theory.* New York: Appleton-Century-Crofts.

Ricchiute, D.N. (1984). An empirical assessment of the impact of alternative task presentation modes on decision-making research in auditing, *Journal of Accounting Research*, 22, 341-350.

Ripley, B.D. (1993). Statistical aspects of neural networks. In Barndorff-Nielsen, O.E., Jensen, J.L., & Kendall, W.S. (Eds.), *Networks and chaos- Statistical and probabilistic aspects* (pp. 40-123). London: Chapman and Hall.

Rips, L.J., Shoben, E.J., & Smith, E.E. (1973). Semantic distance and the verification of semantic relations. *Journal of Verbal Learning and Verbal Behavior*, 12, 1-20.

Robins, A.V. (1992). Incorporating supervised learning in the domains account of categorization. *Connection Science*, 4, 45-56.

Rodgers, W. (1991a). How do loan officers make their decisions about credit risks? A study of parallel distributed processing. *Journal of Economic Psychology*, 12, 243-265.

Rodgers, W. (1991b). Evaluating accounting information with causal models: Classification of methods and implications for accounting research. *Journal of Accounting Literature*, 10, 151-180.

Rodgers, W. (1992). The effects of accounting information on individuals' perceptual processes. *Journal of Accounting, Auditing and Finance*, 7, 67-95.

Rodgers, W., & Housel, T.J. (1987). The effects of information and cognitive processes on decision making. *Accounting and Business Research*, 18 (69), 67-74.

Rodgers, W., & Johnson, L.W. (1988). Integrating credit models using accounting information with loan officers' decision processes. *Accounting and Finance*, November 1988, 1-22.

Romaniuk, S.G., & Hall, L.O. (1992). Decision making on creditworthiness, using a fuzzy connectionist model. *Fuzzy Sets and Systems*, 48, 15-22.

Rosch, E. (1978). Principles of categorization. In Rosch, E., & Lloyd, B.B. (Eds.), *Cognition and categorization* (pp. 27-48). Hillsdale, NJ: Erlbaum.

Rosch, E., & Mervis, C.B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7, 573-605.

Rosenberg, C.R. (1988). *Learning the connection between spelling and sound: A network model of oral reading*. Unpublished doctoral dissertation, Princeton University.

Rosenberg, E., & Gleit, A. (1994). Quantitative methods in credit management: A survey. *Operations Research*, 42, 589-613.

Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65, 386-408.

Rosenblatt, F. (1962). *Principles of neurodynamics*. New York: Spartan.

Ross, S. A. (1976). The arbitrage price theory of capital asset pricing. *Journal of Economic Theory*, 13, 341-360.

Rouse, W.B., Hammer, J.M., & Lewis, C.M. (1989). On capturing human skills and knowledge: Algorithmic approaches to model identification. *IEEE Transactions on Man, Systems and Cybernetics*, 19, 558-573.

Rumelhart, D.E., & Norman, D.A. (1985). Representation of knowledge. In Aitkenhead, A.M., & Slack, J.M. (Eds.), *Issues in cognitive modelling* (pp. 15-62). Hillsdale, NJ: Erlbaum.

Rumelhart, D.E., Hinton, G.E. & McClelland, J.L. (1986). A general framework for parallel distributed processing. In Rumelhart, D.E., & McClelland, J.L. (Eds.), *Parallel distributed processing; Explorations in the microstructure of cognition*. Cambridge, MA: MIT Press.

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.

Rumelhart, D.E., & McClelland, J.L. (Eds.). (1986). *Parallel distributed processing; Explorations in the microstructure of cognition*. Cambridge, MA: MIT Press.

Rumelhart, D.E., Smolensky, P., McClelland, J.L., & Hinton, G.E. (1986). Schemata and sequential thought processes in PDP models. In Rumelhart, D.E., & McClelland, J.L. (Eds.), *Parallel distributed processing; Explorations in the microstructure of cognition*. Cambridge, MA: MIT Press.

Rumelhart, D.E., & Zipser, D. (1986). Feature discovery by competitive learning. In Rumelhart, D.E., & McClelland, J.L. (Eds.), *Parallel distributed processing; Explorations in the microstructure of cognition*. Cambridge, MA: MIT Press.

Rumelhart, D.E., & Todd, P.M. (1993). Learning and connectionist representations. *Attention and Performance*, 14, 3-30.

Salchenberger, L.M., Cinar, E.M., & Lash, N.A. (1992). Neural networks: A new tool for predicting thrift failures. *Decision Sciences*, 23, 899-916.

Sanger, D. (1989). Contribution analysis: A technique for assigning responsibilities to hidden units in connectionist networks. *Connection Science*, 1, 115-138.

Schepanski, A. (1983). Tests of theories of information processing behavior in credit judgement. *The Accounting Review*, 58, 581-599.

Schöneburg, E. (1990). Stock price prediction using neural networks: A project report. *Neurocomputing*, 2, 17-27.

Searle, J.R. (1980). Minds, brains and programs. *The Behavioral and Brain Sciences*, 3, 417-424.

Seidenberg, M.S. (1993). Connectionist models and cognitive theory. *Psychological Science*, 4, 228-235.

Seidenberg, M.S., & McClelland, J.L. (1989). A distributed developmental model of word recognition and naming. *Psychological Review*, 96, 523-568.

Sejnowski, T.J., & Rosenberg, C.R. (1986). *NETtalk: A parallel network that learns to read aloud* (Technical report JHU-EECS-86-01). John Hopkins University.

Selling, T., & Shank, J. (1989). Linear versus process tracing approaches to judgement modelling: A new perspective on cue importance. *Accounting, Organizations and Society*, 14, 65-77.

Selfridge, M., Biggs, S.F., & Krupka, G.R. (1992). A cognitive model of auditor's going-concern judgement. *International Journal of Intelligent Systems*, 7, 393-417.

Sen, T.K., Oliver, R., & Sen, N. (1995). Predicting corporate mergers. In Refenes, A.P. (Ed.), *Neural networks in the capital markets* (pp. 325-340). New York: Wiley.

Shaw, M.J., & Gentry, J.A. (1990). Inductive learning for risk classification. *IEEE Expert*, February 1990, 47-53.

Shanks, D.R. (1990). Connectionism and the learning of probabilistic concepts. *The Quarterly Journal of Experimental Psychology*, 42(A), 209-237.

Shanks, D.R. (1991). Categorization by a connectionist network. *Journal of Experimental Psychology, Learning, Memory and Cognition*, 17, 433-443.

Shanks, D.R. (1992). Connectionist accounts of the inverse base-rate effect in categorization. *Connection Science*, 4, 3-18.

Shanon, B. (1992). Are connectionist models cognitive? *Philosophical Psychology*, 5, 235-255.

Sharda, R., & Patil, R.B. (1993). A connectionist approach to time series prediction: An empirical test. In Trippi, R.R., & Turban, E. (Eds.), *Neural networks in finance and investing* (pp. 451-464). Chicago, IL: Probus Publishing.

Sharkey, N.E. (1991). Connectionist representation techniques. *Artificial Intelligence Review*, 5, 143-167.

Shepard, R.N. (1958). Stimulus and response generalization: Deduction of the generalization gradient from a trace model. *Psychological Review*, 65, 242-256.

Shepard, R.N. (1962). The analysis of proximities: Multidimensional scaling with an unknown distance function I, *Psychometrika*, 27, 125-140.

Shepard, R.N. (1987). Toward a universal law of generalization for psychological science. *Science*, 237, 1317-1323.

Shields, M.D. (1984). A predecisional approach to the measurement of the demand for information in a performance report. *Accounting, Organizations and Society*, 9, 355-363.

Simnett, R., & Trotman, K.T. (1989). Auditor versus model: Information choice and information processing. *The Accounting Review*, 64, 514-528.

Simon, H.A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69, 99-118.

Simon, H.A. (1979). Information-processing theory of human problem solving. In Estes, W. (Ed.), *Handbook of Learning and Cognitive Processes* (5). Hillsdale, NJ: Erlbaum.

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.

Slovic, P. (1990). Choice. In Osherson, D.N., & Smith, E.E. (Eds.), *Thinking: An invitation to cognitive science* , 3 (pp. 89-115). Cambridge, MA: MIT Press.

Slovic, P., Fischhoff, B., & Lichtenstein, S. (1977). Behavioral decision theory. *Annual Review of Psychology*, 28, 1-39.

Smith, E.E. (1990). Categorization. In Osherson, D.L., & Smith, E.E. (Eds.), *Thinking: An invitation to cognitive science*, 3 (pp. 33-53). Cambridge, MA: MIT Press.

Smith, E.E., & Medin, D. (1981). *Categories and concepts*. Cambridge, MA: Harvard University Press.

Smith, E.E., Langston, C., & Nisbett, R.E. (1992). The case for rules in reasoning. *Cognitive Science*, 16, 1-40.

Smith, M. (1993). *Neural networks for statistical modelling*. New York: Van Nostrand Reinhold.

Smolensky, P. (1986). Information processing in dynamical systems: Foundations of harmony theory. In Rumelhart, D.E., & McClelland, J.L. (Eds.), *Parallel distributed processing; Explorations in the microstructure of cognition*. Cambridge, MA: MIT Press.

Smolensky, P. (1988). On the proper treatment of connectionism. *The Behavioral and Brain Sciences*, 11, 1-74.

SPSS Inc. (1990). *SPSS Reference Guide*. Chicago, IL.: SPSS Inc.

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: Association for Computing Machinery.

Steiner, M., & Wittkemper, H.G. (1995). Neural networks as an alternative stock market model. In Refenes, A.P. (Ed.), *Neural networks in the capital markets* (pp. 137-147). New York: Wiley.

Stone, M. (1974). Cross validatory choice and assessment of statistical predictions. *Journal of the Royal Statistical Society*, 36(b), 111-147.

Surkan, A.J., & Ying, X. (1991). Bond rating formulas derived through simplifying a trained neural network. *Proceedings of the International Joint Conference on Neural Networks* (pp. 1566-1570). New York: IEEE.

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.

Tank, D.W., & Hopfield, J.J. (1987). Collective computation in neuronlike circuits. *Scientific American*, December 1987, 62-70.

Taraban, R., McDonald, J., & MacWhinney, B. (1989). Category learning in a connectionist model: Learning to decline the German definite article. In Corrigan, R., Eckman, F., &

Noonan, N. (Eds.), *Linguistic categorization* (pp. 163-193). Philadelphia, PA: Benjamins.

Taraban, R., & Palacios, J.M. (1993). Exemplar models and weighted cue models in category learning. *The Psychology of Learning and Motivation*, 29, 91-127.

The Economist. (1987). What the brain builders have in mind. *The Economist*, may 1987, 86-88.

Touretzky, D.S., & Pomerleau, D.A. (1994). Reconstructing physical symbol systems. *Cognitive Science*, 18, 345-353.

Towell, G.G., & Shavlik, J.W. (1993). Extracting refined rules from knowledge-based neural networks. *Machine Learning*, 13, 71-101.

Trigueiros, D., & Berry, R. (1990). *The application of neural network based methods to the extraction of knowledge from accounting reports.* Poster session presented at the INNC-90, Paris.

Trippi, R.R., & Turban, E. (Eds.). (1993). *Neural networks in finance and investing.* Chicago, IL: Probus Publishing.

Trotman, K.T., & Sng, J. (1989). The effect of hypothesis framing, prior expectations and cue diagnosticity on auditors' information choice. *Accounting, Organizations and Society*, 14, 565-576.

Tsibouris, G., & Zeidenberg, M. (1995). Testing the efficient markets hypothesis with gradient descent algorithms. In Refenes, A.P. (Ed.), *Neural networks in the capital markets* (pp. 127-136). New York: Wiley.

Tversky, A. (1977). Features of similarity. *Psychological Review*, 84, 327-352.

Tversky, A., & Kahneman, D. (1974). Judgement under uncertainty: Heuristics and biases. *Science*, 185, 1124-1131.

Udo, G. (1993). Neural network performance on the bankruptcy classification problem. *Computers and Industrial Engineering*, 25, 377-380.

Usher, M., & Zakay, D. (1993). A neural network model for attribute-based decision processes. *Cognitive Science*, 17, 349-396.

Utans, J., & Moody, J. (1991). Selecting neural network architectures via the prediction risk: Application to corporate bond rating prediction. *Proceedings on the First International Conference on Artificial Intelligence Applications on Wall Street.* Los Alamitos, CA: IEEE Computer Society Press.

van Gelder, T. (1993). Is cognition categorization ? *The Psychology of Learning and Motivation*, 29, 469-494.

Wagner, A.R. (1978). Expectancies and the priming of STM. In Hulse, S.H., Fower, H., & Honig, W.H. (Eds.), *Cognitive processes in animal behaviour* (pp. 177-210). Hillsdale, NJ: Erlbaum.

Waller, W.S., & Felix, W.L. (1989). Auditors' causal judgements: Effects of forward vs. backward inference on information processing. *Accounting, Organizations and Society*, 14, 179-200.

Wasserman, P.D. (1989). *Neural computing: Theory and practice.* New York: van Nostrand Reinhold.

Wasserman, P.D. (1993). *Advanced methods in neural computing*. New York: van Nostrand Reinhold.

Weber, E.U., Böckenholt, U., Hilton, D.J., & Wallace, B. (1993). Determinants of diagnostic hypothesis generation: Effects of information, base rates and experience. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 19, 1151-1164.

Weigend, A.S., Huberman, B.A., & Rumelhart, D.E. (1990). Predicting the future: A connectionist approach. *International Journal of Neural Systems*, 1, 193-209.

Weigend, A.S., Rumelhart, D.E., & Huberman, B.A. (1991). Generalization by weight-elimination with application to forecasting. *Advances in Neural Information Processing Systems*, 3, 875-882.

Werbos, P.J. (1974). *Beyond regression: New tools for prediction and analysis in the behavioral sciences,* Unpublished doctoral dissertation, Harvard University, MA.

White, G.I., Sondhi, A.C. & Fried, D. (1994) *The analysis and use of financial statements*, New York: Wiley.

White, H. (1988). Economic prediction using neural networks: The case of IBM daily stock returns, Proceedings of the IEEE International Conference on Neural Networks (pp. II451-II458). Reprinted in Trippi, R.R., & Turban, E. (Eds.), *Neural networks in finance and investing* (pp. 315-328). Chicago, IL: Probus Publishing.

White, H. (1989). Learning in artificial neural networks: A statistical perspective. *Neural Computation*, 1, 425-464.

White, H. (1990). Connectionist nonparametric regression: Multilayer feedforward networks can learn arbitrary mappings. *Neural Networks*, 3, 535-549.

White, H. (1992). *Artificial neural networks : approximation and learning theory.* Cambridge, MA : Blackwell.

Whittlesea, B.W.A., & Brooks, L.R. (1988). Critical influence of particular experiences in the perception of letters, words and phrases. *Memory and Cognition*, 16, 387-399.

Widrow, B., & Hoff, M.E. (1960). Adaptive switching circuits, *IRE WESCON Conv. Record* , 4, 96-104.

Widrow, B., Winter, R.G., & Baxter, R.A. (1987). Learning phenomena in layered neural networks. *IEEE First International Conference on Neural Networks, II* (pp. 411-429), San Diego, CA: SOS Printing.

Wilcox, J.W. (1973). A prediction of business failure using accounting data. *Journal of Accounting Research*, 11 (supplement), 163-171.

Wilcox, J.W. (1976). The gamblers ruin approach to business risk. *Sloan Management Review*, 18, 33-46.

Wilson, R., & Sharda, R. (1994). Bankruptcy prediction using neural networks. *Decision Support Systems*, 11, 545-557.

Windsor, C.G., & Harker, A.H. (1990). Multi-variate financial index prediction - A neural network study. *Proceedings of the INNC -90- International Neural Network Conference* (pp. 357-360). Dordrecht: Kluwer.

Winograd, T., & Flores, F. (1986). *Understanding computers and cognition*. Norwood, NJ: Ablex.

Wittgenstein, L. (1953). *Philosophical investigations*. New York: Macmillian.

Wong, F.S., Wang, P.Z., Goh, T.H., & Quek, B.K. (1992). Fuzzy neural systems for stock selection. *Financial Analyst Journal*, January-February 1992, 47-80.

Wray, B., Palmer, A., & Bejou, D. (1994). Using neural network analysis to evaluate buyer-seller relationships. *European Journal of Marketing*, 28, 32-48.

Wright, W.F. (1977). Financial information processing models: An empirical study. *The Accounting Review*, 52, 676-689.

Yamamoto, Y., & Zenios, S.A. (1993). Predicting prepayment rates for mortgages using the cascade-correlation learning algorithm. *The Journal of Fixed Income*, March 1993, 86-96.

Yoon, Y., & Swales, G. (1991). Predicting stock price performance: A neural network approach. *Proceedings of the IEEE 24th Annual International Conference of Systems Sciences* (pp. 156-162). San Diego, CA.: IEEE.

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.

Yoon, Y., Guimaraes, T., & Swales, G. (1994). Integrating artificial neural networks with rule-based expert systems. *Decision Support Systems*, 11, 497-507.

Zadeh, L.A. (1965). Fuzzy sets. *Information and Control*, 8, 338-353.

Zavgren, C.V. (1985). Assessing the vulnerability to failure of American industrial firms: A logistic analysis. *Journal of Banking and Finance*, 12, 19-45.

Zavgren, C.V., & Friedman, G.E. (1988). Are bankruptcy prediction models worthwhile? An application in securities analysis. *Management International Review*, 28, 34-44.

Zeidenberg, M. (1990). *Neural networks in artificial intelligence*. New York: Ellis Horwood.

Zipser, D. (1990). Modelling cortical computation with backpropagation. In Gluck, M.A., & Rumelhart, D.E. (Eds.), *Neuroscience and connectionist theory* (pp. 355-383). Hillsdale, NJ: Erlbaum.

Ziebart, D.A. (1987). The effects of annual accounting data on stock returns and trading activity: A causal model study. *Journal of the American Statistical Association*, 82 (applications), 733-738.

Ziebart, D.A., & Reiter, S.A. (1992). Bond rating, bond yields and financial information. *Contemporary Accounting Research,* 9, 252-282.

Zimmer, I. (1980). A lens study of the prediction of corporate failure by bank loan officers. *Journal of Accounting Research*, 18, 629-639.

Appendix A. Summary of applications of neural networks to problems in business administration.

Reference	Application area[1]	Model	Benchmark	RA[2]	Cog[3]
Altman, Marco & Varetto, 1994	Bankruptcy prediction[4]	Backpropagation	Discriminant analysis	N	N
Baestens & van der Bergh, 1995	Stock index prediction	Backpropagation	Regression analysis	y[5]	N
Barker, 1990	Financial analysis	N/A	None	N	N
Binks & Allison, 1991	Financial data recognition	Backpropagation and Self-organizing map	None	N	N
Brown, 1992	Consumer schema modelling	Constraint satisfaction model	None	N	Y
Chakaborty, Mehrotra, Mohan & Ranka, 1992	Forecasting flour prices	Backpropagation	ARMA(1,1) and AR(2)	N	N
Chang, Sheu & Thomas, 1993	Stock price prediction	Backpropagation	None	N	N
Coats & Fant, 1993	Recognizing financial distress patterns	Cascade correlation	Discriminant analysis	N	N
Collins, Gosh & Scofield, 1988	Emulation of mortgage under-writing judgements	Restricted Coulomb Energy NN	None	Y	n[6]
Erxleben, Baetge, Feidicker, Koch, Krause & Mertens, 1992	Bankruptcy prediction	Backpropagation	Discriminant analysis	N	N
Deng, 1993	Commercial loan evaluation	Backpropagation	None	N	N
Diamond, Shadbolt, Barac & Refenes, 1993	Tactical asset allocation	Backpropagation	Weighted portfolio	N	N
Dutta & Shekhar, 1988	Bond rating	Backpropagation	Regression analysis	N	N
Grudnitski & Osburn, 1993	Forecasting futures prices	Backpropagation	None	N	N
Grudnitski & Do, 1995	Forecasting futures prices	Backpropagation	None	Y	N
Hruschka, 1993	Estimating market response functions	Backpropagation	Regression analysis	y[7]	N
Hsu, Hsu & Tenorio, 1993	Predicting currency exchange rates	Supervised Clustering Network	None	N	N
Hutchinson, Lo & Poggio, 1994	Pricing and hedging derivative securities (Black-Scholes test)	Radial basis function networks and backpropagation	Regression analysis and projection pursuit	Y[8]	N
Jensen, 1992	Credit scoring	Backpropagation	None	N	N
Jung & Burns, 1993	Managerial problem diagnosis	Backpropagation	None	Y	n[9]
Kamijo & Tanigawa, 1990	Stock price pattern recognition	Recurrent NN	None	Y	n[10]
Kimoto, Asakawa, Yoda & Takeoka, 1990	Stock price prediction	Backpropagation	Regression analysis	Y	N
Kim, Weistroffer & Redmond, 1993	Bond rating	Backpropagation	Discriminant, regression and logistic analysis, ID3	N	N

1 In this summary, the *author's terms* indicating the application area are generally used.

2 Representational analysis is marked Y if the study contains an analysis of how the representations of the neural network performs the vector mappings. Lowercase letters are used to indicate doubt about the classification.

3 Cognitive is marked Y if the study refers to the neural network as a cognitive model or compares it to a cognitive model of the task performance. Lowercase letters are used to indicate doubt about the classification.

4 The authors use the term "corporate distress diagnosis".

5 Baestens and van der Bergh (1995) compute a "decisiveness" measure of input variables to evaluate the sensitivity of the response to variations in these inputs.

6 However, Collins et al. (1988) state that: "The system was trained on several thousand previous underwriter judgements and learned to mimic their underwriting skills".

7 Hruschka (1993) performs sensitivity analysis to investigate the input/output mapping of the neural network.

8 Comprehensive analysis of prediction errors.

9 Jung and Burns, (1993) discuss the knowledge representation of connectionist systems in a systems perspective, and make no comments on hypothetical resemblance to the cognitive system of humans.

10 However, the task of Kamijo and Tanigawa (1990) is a typical pattern recognition task performed by experts in technical analysis.

Kryzanowski, Galler & Wright, 1993	Stock return classification	Boltzmann Machine	None	N	N
Liang, Chandler, Han & Roan, 1992	LIFO/FIFO classification	Backpropagation	Probit analysis and ID3	N	N
Malliaris & Salchenberger, 1993	Estimation of option prices	Backpropagation	Black-Scholes	N	N
Martin-del-Brio & Serrano-Cinca, 1993	Bankruptcy prediction[1]	Selforganizing map	None	Y	N
Mehta, 1995	Exchange rate prediction	Backpropagation	None	N	N
Moody & Utans, 1995	Bond rating	Backpropagation	Regression analysis	n[2]	N
Nottola, Condamin & Naim, 1992	Company evaluation	Backpropagation	ID3	y[3]	N
Odom & Sharda, 1990	Bankruptcy prediction	Backpropagation	Discriminant analysis	N	N
Piramuthu, Shaw & Gentry, 1994	Loan evaluation	Backpropagation and 2. order modification	Probit analysis and ID3	N	N
Poddig, 1995	Bankruptcy prediction	Backpropagation and LVQ	Discriminant analysis	N	N
Raghupathi, Schkade & Raju, 1991	Bankruptcy prediction	Backpropagation		N	n[4]
Rahimian, Singh, Thammachote & Virmani, 1993	Bankruptcy prediction	Backpropagation, Athena and simple perceptron	None	N	N
Refenes, Azema-Barac, Chen & Karoussos, 1993	Currency exchange rate prediction	Backpropagation	None	N	N
Refenes, 1993	Exchange rate prediction	CLS+	AR(4) and exponential smoothing	N	N
Refenes, Zapranis & Francis, 1995	Stock price prediction	Backpropagation with variations	Regression analysis	y[5]	N
Rehkugler & Poddig, 1991	Stock price prediction	Backpropagation and Boltzmann machine	"Naive prognose"	N	N
Romaniuk & Hall, 1992	Evaluation of creditworthiness	Feed forward network with cell recruitment learning	None	n[6]	N
Salchenberger, Cinar & Lash, 1992	Bankruptcy prediction[7]	Backpropagation	Logit model	N	N
Schöneburg, 1990	Stock price prediction	Adaline, Madaline, Simple perceptron and Backpropagation	None	N	N
Sen, Oliver & Sen, 1995	Corporate merger prediction	Backpropagation	Logistic regression analysis	Y	N
Sharda & Patil, 1992	Time series prediction	Backpropagation	AUTOBOX[8]	N	N
Singleton & Surkan, 1995	Bond rating changes	Backpropagation	Discriminant analysis	N	y[9]
Srivastava, 1992	Business loan evaluation	Backpropagation	None	N	n[10]
Steiner & Wittkemper, 1995	Stock price prediction	Backpropagation	Regression analysis	N	N
Surkan & Ying, 1991	Bond rating	Backpropagation	None	Y	N

[1] Martin-del-Brio and Serrano-Cinca (1993) study *bank* classification and bankruptcy prediction.
[2] Moody and Utans (1995) use sensitivity analysis to determine the importance of input units, and thus, study the input/output mappings of corresponding variables.
[3] Nottola et. al (1992) use ID3 to extract rules from the input to hidden unit response mapping.
[4] However, Raghupati et al. (1991) state: "Various financial ratios may be giving some intermediate features such as immediate financial health of the company, long-term financial health, recent revenue generating trends, and others. Based on these higher-level features, the network may be arriving at a categorizing decision".
[5] Refenes et al. (1995) perform sensitivity analysis to investigate the input/output mappings.
[6] Romaniuk and Hall (1992) give examples of rules extracted from the neural network by "traversing" the network. The exact method of this "traversing" is not explained.
[7] Salchenberger et al. (1992) study failure of *thrift institutions.*
[8] The specific ARIMA model selected by AUTOBOX is not reported in Sharda and Patil (1992).
[9] Singleton and Surkan (1995) state that "Neural network success........ suggests that neural networks may have captured some of the judgement exercised by these analysts".
[10] However, Srivastava (1993) states about the model: "It simulates human judgement and integrates it with mathematical analytical tools".

Surkan & Singleton, 1990	Bond rating	Backpropagation	Discriminant analysis	N	y[1]
Tam & Kiang, 1992	Bank failure prediction	Backpropagation	Discriminant analysis, logistic regression, KNN and ID3	N	N
Tam, 1991	Bank bankruptcy prediction	Backpropagation	See Tam & Kiang, 1992	N	N
Trigueiros & Berry, 1993	Modelling industry homogeneity	Backpropagation	Discriminant analysis	Y	n[2]
Tsibouris & Zeidenberg, 1995	Stock price prediction	Backpropagation and Temporal Difference NN	None	N	N
Udo, 1993	Bankruptcy classification	Backpropagation	Regression analysis	N	N
Utans & Moody, 1991	Bond rating	Backpropagation	Regression analysis	N	N
Weigend, Huberman & Rumelhart, 1990	Economic time series prediction	Backpropagation	AR(2)	N	N
White, 1988	Stock price prediction	Backpropagation	None	N	N
Wilson & Sharda 1994	Bankruptcy prediction	Backpropagation	Discriminant analysis	N	N
Windsor & Harker, 1990	Financial index prediction	Backpropagation	Regression analysis and AR(4)	N	N
Wong, Wang, Goh & Quek, 1992	Stock ranking/stock selection	Backpropagation	None[3]	N	N
Wray, Palmer & Bejou, 1994	Modelling buyer-seller relationships	Backpropagation	Regression analysis	N	N
Yamamotò & Zenios, 1993	Predicting prepayment rates for mortgages	Cascade correlation	Naive and econometric model[4]	N	N
Yoon & Swales, 1991	Stock price classification	Backpropagation	Discriminant analysis	N	N
Yoon, Swales & Margavio, 1993	Stock price classification	Backpropagation	Discriminant analysis	y[5]	N
Yoon, Guimaraes & Swales, 1994	Stock price classification	Backpropagation	Discriminant analysis[6]	N	N

[1] Surkan and Singleton (1990) state: "There is a hope that some of the intermediate representations may be identified with concepts used by humans to analyze this bond classification problem".

[2] In their original paper presented at INNC 1990, they state: "The emerging organization reproduces the way an expert in ratio analysis chooses variables...Experts put together several points of view around a few significant variables. And extended ratios seem to be trying the same sort of procedure" (Trigueiros & Berry, 1990, p.12).

[3] Wong et al. (1992) focus on integrating the artificial neural network with an expert system.

[4] Benchmark models are not explicitly formulated but reference to an "econometric model" is given.

[5] Yoon et al. (1993) investigate the effect of the different inputs on the classification, not the representation as such.

[6] Focus is on integrating the artificial neural network with a rule-based expert system.

Appendix B. Introductory text, stimulus material, and response form.

FINANSIELL DIAGNOSE

Denne teksten blir lest høyt i klassen:

I dette eksperimentet skal du stille diagnose av tre selskapers lønnsomhet, finansiering, likviditet og soliditet, samt dets generelle situasjon og utvikling. Dette er for at vi skal vite mer om hvordan dere stiller finansielle diagnoser før kurset starter, og er helt anonym. Det er viktig at du følger prinsippene som blir forklart i denne innledningen.

Diagnosen gjennomføres ved at du benytter svarskjemaet som følger etter presentasjonen av hvert selskap. Hvert selskap med tilhørende svarskjema utgjør en dobbeltside i materialet (se neste sider).

På den venstre siden finner du reultatregnskap og balanse for selskapets to siste driftsår, endel beregnede nøkkeltall, samt finansieringsanalyse for siste driftsår. En forklaring til enkelte av regnskapsstørrelsene og formler som er benyttet for beregning av nøkkeltallene finner du på siste side i materialet.

På den høyre siden finner du noen generelle opplysninger, plass for å utføre beregninger (dersom du vil regne ut andre størrelser enn de som allerede er beregnet), og svarskjema for selve diagnosen. Svarskjemaet har fem deler, en del for hvert av de områdene du skal stille diagnose: Lønnsomhet, finansiering, likviditet, soliditet og en generell totalvurdering.

Du skal stille diagnose for selskapets situasjon, og for dets utvikling. I tillegg er det viktig at du karakteriserer situasjonen for de ulike områdene med ett eller få ord. Her skal du ikke bruke adjektiver, slik som for eksempel tragisk lønnsomhet. I stedet skal du bruke substantiver som karakteriserer selskapet, slik som for eksempel krisebedrift, konkursbedrift, suksessbedrift eller liknende.

Til slutt skal du for hvert område angi inntil 4 opplysninger i datagrunnlaget som i størst grad understøtter din karakteristikk av selskapets situasjon. Du skal maksimalt sette tall i 4 ruter. Forhold mellom tall angir du ved å sette tall i to ruter, og trekker en strek mellom rutene.

Før diagnosen starter ber vi deg fylle ut noen generelle data, slik at vi kan vite litt mer om den som har utført diagnosen.

Hva er din høyeste utdanning før du startet på høyere revisorstudium ?

Siviløkonom NHH, BI eller SiB	Juridisk embetseksamen
Siviløkonom fra utlandet	Revisorstudiet DH eller BI
Annen utdanning, vennligst spesifiser:______________	Sosialøkonomisk embetseksamen

Hvor mange års yrkeserfaring har du etter at du gjennomførte denne utdanningen ?

______________ År

Hvor mange års yrkeserfaring har du totalt ?

______________ År

Da kan du starte diagnosen. Vær så grundig som mulig, og husk at du totalt har 40 minutter til rådighet for å stille diagnose for de tre selskapene.

(Introductory text)

A/S A

Resultatregnskap

	1989	1988
Salgsinntekter:	13756944	12429718
Varekostnad:	10648347	9485488
Andre var. kostnader:	0	0
Dekningsbidrag:	3108597	2944230
Andre driftsinntekter:	274366	335932
Lønninger:	2093721	1748167
Andre driftskostnader:	1146485	944675
Avskrivninger:	333109	279801
Tap på fordringer:	0	0
Driftsresultat:	-190352	307519
Finansinntekter:	1248	3394
Finanskostnader:	190023	189522
Resultat før e.o. poster:	-379127	121391
Ekstraord. inntekter:	245533	103779
Ekstraord. kostnader:	123781	0
Resultat før års.oppgj. disp.:	-257375	225170
Skattem. meravskrivning:	-55584	159997
Skatter:	2101	21388
Andre årsoppgjørsdisp.:	0	14000
Årsresultat:	-203892	29785
Foreslått utbytte:	0	

Balanse

	1989	1988
Kasse, bank o.l.:	115682	274886
Kundefordringer:	43960	49881
Andre korts. fordringer:	3690	7901
Varelager:	926400	801000
Omløpsmidler:	1089732	1133668
Langsiktige fordringer:	53889	48808
Maskiner, inventar o.l.:	260960	356235
Bygninger:	354400	379170
Annen fast eiendom:	502550	791815
Anleggsmidler	1171799	1576028
Sum eiendeler:	2261531	2709696
Kassekreditt:	191561	19772
Leverandørgjeld:	651768	725941
Offentlig gjeld:	266574	318237
Annen kortsiktig gjeld:	199142	259200
Sum kortsiktig gjeld:	1309045	1323150
Langsiktig gjeld:	865192	1095360
Anleggsreserve:	0	0
Varelagerreserve:	0	0
Konsolideringsfond:	35200	35200
Andre bet. sk.fr. avs.:	0	0
Sum bet. skattefr. avs.:	35200	35200
Aksjekapital:	125000	125000
Reservefond:	81500	81500
Frie fond:	-154406	49486
Sum egenkapital:	52094	255986
Sum gjeld og egenkap.:	2261531	2709696

Nøkkeltall

	1989	1988
Omsetningsvekst:	10.7%	
Vekst faste kostnader:	20.2%	
Dekningsgrad:	22.6%	23.7%
Omsetning pr. ansatt:	982639	
Dekningsbidrag pr. ansatt:	222043	
Lønnsomhetsmargin:	-1.4%	2.5%
Overskuddsgrad:	-2.8%	1.0%
Kapitalens omløpshast.:	6.1	4.6
Rentabilitet totalkapital:	-8.4%	11.5%
Rentabilitet egenkapital:	-727.8%	47.4%
Gj. sn. gjeldsrente:	8.7%	7.8%
Rentedekningsgrad:	-99.5%	164.1%
Langsiktig lagerfinansiering:	-23.7%	-23.7%
Gj. sn. lagertid (dager):	32	31
Gj. sn. kundekredittid:	1	1
Gj. sn. leverandørkredittid:	18	23
Likviditetsgrad 1: (Current ratio)	83.2%	85.7%
Likviditetsgrad 2: (Acid test)	12.5%	25.1%
Egenkapitalandel:	3.1%	10.1%

Finansieringsanalyse

	1989
Resultat før årsoppgj. disp.:	-257375
- Skatter:	2101
+ Ordinære avskrivninger:	333109
- Utbytte	0
= Tilført fra årets drift:	73633
+ Ny aksjekapital	0
= Egenfinansiering:	73633
Endring langsiktig gjeld:	-230168
Endring anleggsmidler:	-126704
Endring arbeidskapital:	-29831
Endring kundefordringer:	-5921
Endring andre korts. fordr.:	-4211
Endring varelager:	125400
Endring leverandørgjeld:	-74173
Endring offentlig gjeld:	-51663
Endring annen korts. gjeld:	-60058
Netto likviditetsendring:	-330993
Endring kassekreditt:	171789
Endring kasse, bank:	-159204

(Stimulus material)

Ytterligere opplysninger

A/S A er en bedrift lokalisert i Oppland ☐. Selskapet driver primært med produksjon av næringsmidler ☐.
I 1989 hadde selskapet totalt 14 ☐ fast ansatte. Selskapet er ikke tilknyttet noe konsern, og det er forøvrig ikke noe spesielt å bemerke vedrørende driften.

Plass til evt. utregninger:

______

______

______

______

______

Lønnsomhet

Hvordan vil du karakterisere selskapets lønnsomhet i 1989?	Svært dårlig ☐	Dårlig ☐	Hverken dårlig eller god ☐	God ☐	Svært god ☐
Hvordan vil du karakterisere den lønnsomhetsmessige utvikling selskapet er inne i?	Svært negativ ☐	Negativ ☐	Hverken positiv eller negativ ☐	Positiv ☐	Svært positiv ☐

Karakteriser lønnsomhetssituasjonen med ett eller få ord. ______

Sett tallet **1** i inntil 4 ruter i datagrunnlaget som du mener i størst grad støtter din karakteristikk av lønnsomhetssituasjonen.

Finansiering

Hvordan vil du karakterisere selskapets finansielle situasjon i 1989?	Svært dårlig ☐	Dårlig ☐	Hverken dårlig eller god ☐	God ☐	Svært god ☐
Hvordan vil du karakterisere den finansielle utvikling selskapet er inne i?	Svært negativ ☐	Negativ ☐	Hverken positiv eller negativ ☐	Positiv ☐	Svært positiv ☐

Karakteriser den finansielle situasjon med ett eller få ord. ______

Sett tallet **2** i inntil 4 ruter i datagrunnlaget som du mener i størst grad støtter din karakteristikk av finansieringen.

Likviditet

Hvordan vil du karakterisere selskapets likviditet i 1989?	Svært dårlig ☐	Dårlig ☐	Hverken dårlig eller god ☐	God ☐	Svært god ☐
Hvordan vil du karakterisere den likviditetsmessige utvikling selskapet er inne i?	Svært negativ ☐	Negativ ☐	Hverken positiv eller negativ ☐	Positiv ☐	Svært positiv ☐

Karakteriser likviditetssituasjonen med ett eller få ord. ______

Sett tallet **3** i inntil 4 ruter i datagrunnlaget som du mener i størst grad støtter din karakteristikk av likviditetssituasjonen

Soliditet

Hvordan vil du karakterisere selskapets soliditet i 1989?	Svært dårlig ☐	Dårlig ☐	Hverken dårlig eller god ☐	God ☐	Svært god ☐
Hvordan vil du karakterisere den soliditetsmessige utvikling selskapet er inne i?	Svært negativ ☐	Negativ ☐	Hverken positiv eller negativ ☐	Positiv ☐	Svært positiv ☐

Karakteriser soliditetssituasjonen med ett eller få ord. ______

Sett tallet **4** i inntil 4 ruter i datagrunnlaget som du mener i størst grad støtter din karakteristikk av soliditetssituasjonen.

Generell sammenfatning

Hvordan vil du karakterisere selskapets økonomiske situasjon i 1989?	Svært dårlig ☐	Dårlig ☐	Hverken dårlig eller god ☐	God ☐	Svært god ☐
Hvordan vil du karakterisere den økonomiske utvikling selskapet er inne i?	Svært negativ ☐	Negativ ☐	Hverken positiv eller negativ ☐	Positiv ☐	Svært positiv ☐

Karakteriser selskapets totale økonomiske situasjon med ett eller få ord. ______

Sett tallet **5** i inntil 4 ruter i datagrunnlaget som du mener i størst grad støtter din karakteristikk av totalsituasjonen.

(Response form)

Appendix C. Companies with complete financial statements and their code.

Company	Location	Code
Olaf Bryhn A/S	2436 Våler i Solør	A/S A
Frekhaug Støperi A/S	5110 Frekhaug	A/S B
Friva A/S	1820 Spydeberg	A/S C
Fotlandsvåg Fabrikk A/S	5255 Fotlandsvåg	A/S D
Joh. Fredheim A/S	8200 Fauske	A/S E
Foss Snekkeri A/S	7396 Jerpstad	A/S F
Forus Industri A/S	4033 Forus	A/S G
Formular Service A/S	1060 OSLO 6	A/S H
Einersen Trykkeri A/S	1060 OSLO 6	A/S I
Magnus Engmark A/S	0975 OSLO 9	A/S J
E Trykk	0578 OSLO 5	A/S K
A/S Demokraten	1600 Fredrikstad	A/S L
Sigurd Ecklund A/S	1010 OSLO 1	A/S M
Eidsvold Blad A/S	2081 Eidsvoll	A/S N
Claussen & Heyerdahl A/S	0482 OSLO 4	A/S O
A/S Bygg og Innbu	6762 Almenningen	A/S P
Central Plast A/S	8056 Saltstraumen	A/S Q
Haakon Burø A/S	1010 Oslo 1	A/S R
Brødr. Brøste A/S	6330 Verma	A/S S
Bryne Offset A/S	4341 Bryne	A/S T
Ingebjørg Almankås A/S	3800 Bø i Telemark	A/S U
Audna Bruk A/S	4520 Sør-Audnedal	A/S V
Aanonsen Sats A/S	1010 OSLO 1	A/S W
Falleth A/S	1600 Fredrikstad	A/S X
Fokus A/S	4909 Songe	A/S Y
Folkestad KVV-Service A/S	3800 Bø	A/S Z
Eik Sølv-Plett A/S	3101 Tønsberg	A/S AA
Eker Cementvarefabrikk	3300 Hokksund	A/S AB
Elvarmovner A/S	4001 Stavanger	A/S AC
O.C. Akselsen Fabrikker A/S	4400 Flekkefjord	A/S AD
Thor Berntsen og søn. A/S	1900 Fetsund	A/S AE
Binders A/S	4000 Stavanger	A/S AF
Brandbu Pølsemakeri A/S	2760 Brandbu	A/S AG
Brattværfisk A/S	6580 Vestsmøla	A/S AH
Brevik Blikkvarefabr. A/S	3950 Brevik	A/S AI
Brd. Gilstad Sag & Høvl. A/S	7600 Levanger	A/S AJ
Gjerde Bruk A/S	5700 Voss	A/S AK
Goman-Bakeriet A/S	4301 Sandnes	A/S AL
Grimstad Adressetid. A/S	4890 Grimstad	A/S AM
H.O. Grindheim A/S	5000 Bergen	A/S AN
Grovane Sagbruk A/S	4700 Vennesla	A/S AO
Grønland Grafiske A/S	1081 OSLO 10	A/S AP
Hanssen & Whist A/S	2310 Stange	A/S AQ
Olav Haug Møbelfabr. A/S	2400 Elverum	A/S AR
Hedpall A/S	2340 Løten	A/S AS
Hillesvåg Ullvare A/S	5164 Hjelmås	A/S AT
Hollung Stålindustri A/S	1600 Fredrikstad	A/S AU
Holmen betong A/S	9322 Karlstad	A/S AV
Holten & Asgård møbelfabr.	6652 Surna	A/S AW
John Holvik A/S	6800 Førde	A/S AX
Instrumentbyrået A/S	2200 Kongsvinger	A/S AY
I. C. Iversen eftf. A/S	1040 OSLO 4	A/S AZ
Trygve Jespersen A/S	8500 Narvik	A/S BA
C.A. Johanson Snekker A/S	1060 OSLO 6	A/S BB

Johnson Controls A/S	0667 OSLO 6	A/S BC
Kilen Trevaresalg A/S	3100 Tønsberg	A/S BD
Kirkeby A/S	0661 OSLO 6	A/S BE
Langklopp & Halgunset A/S	7391 Berkåk	A/S BF
Larsen & Mortensen A/S	4000 Stavanger	A/S BG
Lerøy Metallindustri A/S	5250 Lonevåg	A/S BH
A/S Lettbetong	1827 Hobøl	A/S BI
Lie Jærplast A/S	4301 Sandnes	A/S BJ
C.A. Ljungmann & sønn A/S	1060 OSLO 6	A/S BK
Mandal Teppeveveri A/S	4500 Mandal	A/S BL
Moss Jern- og stans. A/S	1580 Rygge	A/S BM
Møre Skofabrikk A/S	6138 Steinsvik	A/S BN
N.K. Nielsen Jernstøp. A/S	1060 OSLO 6	A/S BO
NOFI SVolvær A/S	8301 Svolvær	A/S BP
Nordtveit Skipsbyggeri A/S	5677 Nordtveitgrend	A/S BQ
Nordheimsund Mek. A/S	5600 Nordheimsund	A/S BR
Olsens Vognfabrikk A/S	1870 Ørje	A/S BS
Pla-NY A/S	6083 Haugsbygda	A/S BT
Polar-Boats A/S	4818 Færvik	A/S BU
Protectors A/S	3000 Drammen	A/S BV
Norpower A/S	6500 Kristiansund N.	A/S BW
Randsfjord Glassverk A/S	2700 Jevnaker	A/S BX
Rekon A/S	3200 Sandefjord	A/S BY
Ringerike møbel & tref. A/S	3503 Tyristrand	A/S BZ
Rubb Motor A/S	5420 Rubbestadneset	A/S CA
ScanMatic A/S	4920 Staubø	A/S CB
Servoteknikk A/S	0502 OSLO 5	A/S CC
Signalco A/S	1040 OSLO 4	A/S CD
Slagterborg. Felless. A/S	1010 OSLO 1	A/S CE
Solli Plast A/S	4994 Åkland	A/S CF
Stavprodukter A/S	7500 Stjørdal	A/S CG

Appendix D. Ratio formulas

$SGROWTH = \frac{Sales_2 - Sales_1}{Sales_1}$	Sales growth
$CGROWTH = \frac{Costs_2 - Costs_2}{Costs_1}$	Costs growth
$CONTPR_i = \frac{Contribution\_margin_i}{Sales_i}$	Contribution margin (%)
$PROMARG_i = \frac{Operating\_profit_i + Financial\_revenues_i}{Sales_i}$	Profit margin
$OPMARG_i = \frac{Profit\_before\_extraordinary\_items_i}{Sales_i}$	Operating margin
$ASSTURN_i = \frac{Sales_i}{Total\_assets_i}$	Assets turnover (times)
$ROI_i = \frac{Operating\_profit_i + Financial\_revenues_i}{Total\_assets_i}$	Return on assets
$ROE_i = \frac{Profit\_before\_extraordinary\_items_i}{Equity_i}$	Return on equity
$AIR_i = \frac{Financial\_costs_i}{Interest\_related\_debt_i}$	Average interest rate
$ICOV_i = \frac{Operating\_profit_i + Financial\_revenues_i}{Financial\_costs_i}$	Interest coverage
$LTINV_i = \frac{Working\_capital_i}{Cost\_of\_goods_i}$	Long term invent. financing
$ITURN_i = \frac{Inventory_i * 365}{Cost\_of\_goods_i}$	Invent. turnover time (days)
$ART_i = \frac{Accounts\_receivable_i * 365}{Sales_i * 1.2}$	Collection period (days)
$APT_i = \frac{Accounts\_payable_i * 365}{Goods\_purchased_i * 1.2}$	Accounts payable per. (days)
$CURR_i = \frac{Current\_assets_i}{Current\_debts_i}$	Current ratio
$ACID_i = \frac{Quick\_assets_i}{Current\_debts_i}$	Acid test
$BER_i = \frac{Adjusted\_equity_i}{Total\_assets_i}$	Equity ratio

Appendix E. Summary statistics of independent variables

Variable	Mean	Std. Dev.	Minimum	Maximum	K/S	N
SGROWTH	0.18	0.18	-0.30	0.75	0.10	75
CGROWTH	0.17	0.22	-0.20	1.15	0.15**	75
CONTPR1	0.51	0.19	0.17	0.91	0.053	75
CONTPR2	0.50	0.18	0.11	0.91	0.053	75
PROMARG1	0.05	0.05	-0.06	0.26	0.14**	75
PROMARG2	0.04	0.04	-0.04	0.20	0.17**	75
OPMARG1	0.02	0.05	-0.08	0.21	0.13**	75
OPMARG2	0.02	0.04	-0.04	0.16	0.12**	75
ASSTURN1	2.52	1.09	0.79	7.08	0.15**	75
ASSTURN2	2.56	1.11	1.04	6.23	0.17**	75
ROI1	0.10	0.09	-0.09	0.44	0.10*	75
ROI2	0.09	0.07	-0.08	0.29	0.08	75
ROE1	0.69	1.31	-1.65	7.88	0.18**	75
ROE2	0.39	1.86	-7.28	9.38	0.21**	75
AIR1	0.06	0.04	0.00	0.16	0.06	75
AIR2	0.06	0.04	0.00	0.18	0.06	75
ICOV1	5.02	14.56	-8.02	86.67	0.35**	75
ICOV2	6.97	21.15	-7.17	127.84	0.36**	75
LTINV1	0.70	1.50	-1.72	6.86	0.19**	75
LTINV2	0.65	1.91	-5.95	6.97	0.20**	75
ITURN1	102.77	81.08	0.00	442.54	0.09	75
ITURN2	89.98	79.47	0.00	474.96	0.15**	75
ART1	36.35	18.60	1.18	83.39	0.05	75
ART2	38.04	18.51	0.03	78.75	0.05	75
APT1	69.45	49.90	0.51	262.50	0.18**	75
APT2	69.62	42.20	2.80	221.08	0.12*	75
CURR1	1.27	0.43	0.67	2.79	0.17**	75
CURR2	1.27	0.51	0.24	3.90	0.16**	75
ACID1	0.78	0.34	0.21	2.02	0.09	75
ACID2	0.82	0.38	0.10	2.00	0.09	75
BER1	0.16	0.13	-0.06	0.56	0.10	75
BER2	0.16	0.12	-0.08	0.48	0.11*	75

Summary statistics of independent variables: Mean, standard deviation, minimum value, maximum value, Kolmogorov-Smirnov statistic for test of normality[1] (* and ** indicates significant at α=0.05 and 0.01 respectively), and number of observations.

[1] A Lilliefors version of the Kolmogorov-Smirnov test is used as in the EXAMINE procedure of SPSS.

Appendix F. Correlation matrix of independent variables

	SGRO-WTH	CGRO-WTH	CONT-PR1	CONT-PR2	PRO-MARG1	PRO-MARG2	OPMARG1	OPMARG2
SGROWTH	1.0000	.7045**	.1268	.0250	.1688	.2354*	.1473	.2967**
CGROWTH	.7045**	1.0000	.0773	.0749	.3431**	.1415	.3530**	.2087
CONTPR1	.1268	.0773	1.0000	.9715**	.3207**	.2620*	.2150	.1913
CONTPR2	.0250	.0749	.9715**	1.0000	.2864*	.2568*	.1876	.1912
PROMARG1	.1688	.3431**	.3207**	.2864*	1.0000	.6661**	.9154**	.5650**
PROMARG2	.2354*	.1415	.2620*	.2568*	.6661**	1.0000	.5139**	.8790**
OPMARG1	.1473	.3530**	.2150	.1876	.9154**	.5139**	1.0000	.5849**
OPMARG2	.2967**	.2087	.1913	.1912	.5650**	.8790**	.5849**	1.0000
ASSTURN1	-.2021	-.1313	-.2577*	-.2277*	-.3229**	-.3945**	-.1328	-.1628
ASSTURN2	-.0524	-.1026	-.3630**	-.3782**	-.3745**	-.4589**	-.1984	-.2304*
ROI1	-.0227	.1817	.3529**	.3347**	.8234**	.4100**	.8360**	.4044**
ROI2	.1771	.0225	.2217	.2200	.4593**	.8671**	.3838**	.8433**
ROE1	.1682	.2531*	.3012**	.2614*	.5787**	.4157**	.4997**	.3504**
ROE2	.1612	.1183	.2276*	.2231	.2659*	.4476**	.2314*	.5131**
AIR1	-.0626	-.1276	.3649**	.3487**	.2159	.3038**	-.1385	-.0594
AIR2	-.0903	-.2185	.0850	.0504	.1359	.2528*	-.1496	-.1586
ICOV1	-.1089	-.0145	.0351	.0554	.1685	.0889	.2931*	.2368*
ICOV2	-.0817	-.0257	-.0024	.0190	.0713	.0912	.1956	.2549*
LTINV1	.0911	.2159	.1176	.1226	.3464**	.1274	.3803**	.1704
LTINV2	.0732	.0928	.0514	.0443	.1192	.0204	.1686	.0930
ITURN1	.0535	-.2058	.4048**	.3660**	-.1815	.1197	-.3569**	-.0401
ITURN2	-.2178	-.2475*	.3620**	.3945**	-.1155	.1088	-.2480*	-.0183
ART1	.0919	.1251	.0021	-.0215	.3281**	.3385**	.3045**	.3316**
ART2	.2056	.1875	.1186	.0853	.1394	.2694*	.0817	.2527*
APT1	.3143**	.4236**	.1970	.1751	.3999**	.2486*	.3608**	.2364*
APT2	.1151	.2482*	.3501**	.3765**	.1855	.0909	.1776	.1084
CURR1	.0600	-.0271	-.0536	-.0519	.0389	.1434	.0224	.1297
CURR2	.0356	-.0783	-.0448	-.0530	-.0037	.1192	-.0272	.1145
ACID1	.0946	.1800	.0353	.0405	.3024**	.1833	.3866**	.2918*
ACID2	.1734	.1344	.0810	.0540	.1606	.1142	.2150	.2104
BER1	.0034	-.0386	.0646	.0594	.0375	.0097	.2359*	.2075
BER2	.0526	-.0647	.0607	.0448	.0576	.1709	.2377*	.3941**

cont...

	ASS-TURN1	ASS-TURN2	ROI1	ROI2	ROE1	ROE2	AIR1	AIR2
SGROWTH	-.2021	-.0524	-.0227	.1771	.1682	.1612	-.0626	-.0903
CGROWTH	-.1313	-.1026	.1817	.0225	.2531*	.1183	-.1276	-.2185
CONTPR1	-.2577*	-.3630**	.3529**	.2217	.3012**	.2276*	.3649**	.0850
CONTPR2	-.2277*	-.3782**	.3347**	.2200	.2614*	.2231	.3487**	.0504
PROMARG1	-.3229**	-.3745**	.8234**	.4593**	.5787**	.2659*	.2159	.1359
PROMARG2	-.3945**	-.4589**	.4100**	.8671**	.4157**	.4476**	.3038**	.2528*
OPMARG1	-.1328	-.1984	.8360**	.3838**	.4997**	.2314*	-.1385	-.1496
OPMARG2	-.1628	-.2304*	.4044**	.8433**	.3504**	.5131**	-.0594	-.1586
ASSTURN1	1.0000	.8772**	-.0049	-.1501	-.0928	-.0839	-.2492*	-.3246**
ASSTURN2	.8772**	1.0000	-.1439	-.2114	-.1347	-.1845	-.2635*	-.1810
ROI1	-.0049	-.1439	1.0000	.3759**	.5673**	.2086	.1584	.0066
ROI2	-.1501	-.2114	.3759**	1.0000	.4041**	.5544**	.1879	.1960
ROE1	-.0928	-.1347	.5673**	.4041**	1.0000	.4893**	.2244	.1614
ROE2	-.0839	-.1845	.2086	.5544**	.4893**	1.0000	.0654	-.0835
AIR1	-.2492*	-.2635*	.1584	.1879	.2244	.0654	1.0000	.7055**
AIR2	-.3246**	-.1810	.0066	.1960	.1614	-.0835	.7055**	1.0000
ICOV1	.2342*	.0958	.3077**	.1670	.0978	.0649	-.3407**	-.3465**
ICOV2	.2252	.0967	.1848	.1875	.0435	.0894	-.3649**	-.3788**
LTINV1	-.2112	-.2449*	.3745**	.0422	.1712	.0232	-.1060	-.1600
LTINV2	-.1089	-.0818	.1352	.0227	.0376	-.0034	-.1597	-.1394
ITURN1	-.3302**	-.3210**	-.2498*	.0836	-.1237	.0341	.4354**	.2942*
ITURN2	-.2553*	-.3488**	-.1421	.0589	-.1911	-.0175	.3989**	.1793
ART1	-.3271**	-.3454**	.1369	.2471*	.1962	.2103	-.0365	-.0801
ART2	-.2659*	-.3512**	-.0179	.1862	.2055	.2607*	.0202	-.1330
APT1	-.3590**	-.3120**	.1171	.1205	.3962**	.2499*	-.0965	-.1036
APT2	-.2764*	-.3875**	.0634	-.0599	.1287	.1695	-.0938	-.2452*
CURR1	-.3440**	-.3390**	-.0135	.0363	-.2061	-.0192	-.0288	-.0075
CURR2	-.2711*	-.2146	-.0618	.1110	-.1956	-.0026	-.0211	.0380
ACID1	-.1707	-.2377*	.2925*	.1055	.0920	.0896	-.2463*	-.2756*
ACID2	-.1554	-.1522	.1196	.1239	.0635	.1379	-.2282*	-.2110
BER1	-.0984	-.1120	.0730	-.0343	-.1921	-.0546	-.4443**	-.3594**
BER2	-.1084	-.0657	.0509	.2042	-.1149	.1054	-.4350**	-.3226**

cont...

	ICOV1	ICOV2	LTINV1	LTINV2	ITURN1	ITURN2	ART1	ART2
SGROWTH	-.1089	-.0817	.0911	.0732	.0535	-.2178	.0919	.2056
CGROWTH	-.0145	-.0257	.2159	.0928	-.2058	-.2475*	.1251	.1875
CONTPR1	.0351	-.0024	.1176	.0514	.4048**	.3620**	.0021	.1186
CONTPR2	.0554	.0190	.1226	.0443	.3660**	.3945**	-.0215	.0853
PROMARG1	.1685	.0713	.3464**	.1192	-.1815	-.1155	.3281**	.1394
PROMARG2	.0889	.0912	.1274	.0204	.1197	.1088	.3385**	.2694*
OPMARG1	.2931*	.1956	.3803**	.1686	-.3569**	-.2480*	.3045**	.0817
OPMARG2	.2368*	.2549*	.1704	.0930	-.0401	-.0183	.3316**	.2527*
ASSTURN1	.2342*	.2252	-.2112	-.1089	-.3302**	-.2553*	-.3271**	-.2659*
ASSTURN2	.0958	.0967	-.2449*	-.0818	-.3210**	-.3488**	-.3454**	-.3512**
ROI1	.3077**	.1848	.3745**	.1352	-.2498*	-.1421	.1369	-.0179
ROI2	.1670	.1875	.0422	.0227	.0836	.0589	.2471*	.1862
ROE1	.0978	.0435	.1712	.0376	-.1237	-.1911	.1962	.2055
ROE2	.0649	.0894	.0232	-.0034	.0341	-.0175	.2103	.2607*
AIR1	-.3407**	-.3649**	-.1060	-.1597	.4354**	.3989**	-.0365	.0202
AIR2	-.3465**	-.3788**	-.1600	-.1394	.2942*	.1793	-.0801	-.1330
ICOV1	1.0000	.9752**	.1343	.0302	-.1016	-.0406	-.1328	-.1449
ICOV2	.9752**	1.0000	.1030	.0605	-.0619	-.0273	-.1281	-.0986
LTINV1	.1343	.1030	1.0000	.7223**	-.2350*	-.2211	.1581	.0935
LTINV2	.0302	.0605	.7223**	1.0000	-.0500	-.1164	.1608	.1928
ITURN1	-.1016	-.0619	-.2350*	-.0500	1.0000	.8348**	.0245	.1106
ITURN2	-.0406	-.0273	-.2211	-.1164	.8348**	1.0000	.0589	.0436
ART1	-.1328	-.1281	.1581	.1608	.0245	.0589	1.0000	.7426**
ART2	-.1449	-.0986	.0935	.1928	.1106	.0436	.7426**	1.0000
APT1	-.1326	-.1365	.0417	.0563	-.0231	-.0902	.4440**	.3365**
APT2	-.1377	-.1487	.0961	.0194	.0101	.0541	.2845*	.3060**
CURR1	.1449	.1788	.4559**	.4423**	.1243	.0956	.0039	-.0359
CURR2	.1108	.1416	.3073**	.5301**	.2307*	.1767	.1698	.0547
ACID1	.1752	.1780	.7111**	.5852**	-.3404**	-.2823*	.2008	.1010
ACID2	.0903	.1263	.5430**	.7393**	-.1192	-.2697*	.2392*	.2429*
BER1	.2855*	.2875*	.2350*	.1764	-.0136	.0272	.0052	-.1810
BER2	.2954*	.3185**	.2089	.2299*	.0066	.0091	.0836	-.1145

cont...

	APT1	APT2	CURR1	CURR2	ACID1	ACID2	BER1	BER2
SGROWTH	.3143**	.1151	.0600	.0356	.0946	.1734	.0034	.0526
CGROWTH	.4236**	.2482*	-.0271	-.0783	.1800	.1344	-.0386	-.0647
CONTPR1	.1970	.3501**	-.0536	-.0448	.0353	.0810	.0646	.0607
CONTPR2	.1751	.3765**	-.0519	-.0530	.0405	.0540	.0594	.0448
PROMARG1	.3999**	.1855	.0389	-.0037	.3024**	.1606	.0375	.0576
PROMARG2	.2486*	.0909	.1434	.1192	.1833	.1142	.0097	.1709
OPMARG1	.3608**	.1776	.0224	-.0272	.3866**	.2150	.2359*	.2377*
OPMARG2	.2364*	.1084	.1297	.1145	.2918*	.2104	.2075	.3941**
ASSTURN1	-.3590**	-.2764*	-.3440**	-.2711*	-.1707	-.1554	-.0984	-.1084
ASSTURN2	-.3120**	-.3875**	-.3390**	-.2146	-.2377*	-.1522	-.1120	-.0657
ROI1	.1171	.0634	-.0135	-.0618	.2925*	.1196	.0730	.0509
ROI2	.1205	-.0599	.0363	.1110	.1055	.1239	-.0343	.2042
ROE1	.3962**	.1287	-.2061	-.1956	.0920	.0635	-.1921	-.1149
ROE2	.2499*	.1695	-.0192	-.0026	.0896	.1379	-.0546	.1054
AIR1	-.0965	-.0938	-.0288	-.0211	-.2463*	-.2282*	-.4443**	-.4350**
AIR2	-.1036	-.2452*	-.0075	.0380	-.2756*	-.2110	-.3594**	-.3226**
ICOV1	-.1326	-.1377	.1449	.1108	.1752	.0903	.2855*	.2954*
ICOV2	-.1365	-.1487	.1788	.1416	.1780	.1263	.2875*	.3185**
LTINV1	.0417	.0961	.4559**	.3073**	.7111**	.5430**	.2350*	.2089
LTINV2	.0563	.0194	.4423**	.5301**	.5852**	.7393**	.1764	.2299*
ITURN1	-.0231	.0101	.1243	.2307*	-.3404**	-.1192	-.0136	.0066
ITURN2	-.0902	.0541	.0956	.1767	-.2823*	-.2697*	.0272	.0091
ART1	.4440**	.2845*	.0039	.1698	.2008	.2392*	.0052	.0836
ART2	.3365**	.3060**	-.0359	.0547	.1010	.2429*	-.1810	-.1145
APT1	1.0000	.6357**	-.2494*	-.1442	-.0043	.0855	-.1598	-.1067
APT2	.6357**	1.0000	-.1827	-.2450*	.0483	-.0137	-.0987	-.1702
CURR1	-.2494*	-.1827	1.0000	.7953**	.7010**	.5792**	.5792**	.5223**
CURR2	-.1442	-.2450*	.7953**	1.0000	.5157**	.7052**	.3586**	.4123**
ACID1	-.0043	.0483	.7010**	.5157**	1.0000	.7829**	.5121**	.4820**
ACID2	.0855	-.0137	.5792**	.7052**	.7829**	1.0000	.3290**	.3927**
BER1	-.1598	-.0987	.5792**	.3586**	.5121**	.3290**	1.0000	.9279**
BER2	-.1067	-.1702	.5223**	.4123**	.4820**	.3927**	.9279**	1.0000

(* and ** indicates significant at α=0.05 and 0.01, respectively)

Appendix G. Factor loadings of first benchmark (A)

	FACTOR 1	FACTOR 2	FACTOR 3	FACTOR 4	FACTOR 5	FACTOR 6	FACTOR 7	FACTOR 8	FACTOR 9
SGROWTH	.08136	-.06628	.22839	.03701	.03591	.01493	.01053	-.07604	.90519
CGROWTH	.06698	.23641	-.00705	-.01378	.17767	-.03344	.00612	.03152	.86014
CONTPR1	.03301	.15868	.13987	.13585	.05335	-.00572	.92796	-.01947	.05754
CONTPR2	.02027	.15768	.11434	.12439	.06496	.00788	.93509	.00578	.00302
PROMARG1	.08464	.83786	.33483	.25757	.10944	-.04829	.09305	.04270	.12899
PROMARG2	.02096	.34031	.79575	.38219	.06020	-.00090	.04181	.01235	.08361
OPMARG1	.09853	.85067	.26807	.02779	.15664	.19209	.02287	.14114	.12190
OPMARG2	.06140	.28767	.81902	.07365	.14123	.26260	.02193	.14339	.14586
ASSTURN1	-.14725	-.10296	-.04437	-.81974	-.25192	-.06747	-.08040	.21330	-.15808
ASSTURN2	-.14474	-.16954	-.07751	-.77270	-.34804	-.06292	-.22019	.06718	-.01915
ROI1	.11439	.84726	.22080	-.02157	-.05400	-.06053	.21993	.15509	-.06666
ROI2	.01554	.19349	.90856	.11020	-.03248	-.02151	.05424	.07670	-.00127
ROE1	.01230	.53737	.40988	-.10666	.12987	-.30702	.19917	-.00170	.13770
ROE2	.03097	.04506	.68717	-.15800	.22176	-.01107	.21017	-.00637	.07091
AIR1	-.07888	.05064	.18152	.41592	-.29005	-.58199	.32237	-.30936	-.13819
AIR2	-.09952	.06559	.13438	.46792	-.40164	-.48439	-.01911	-.38055	-.14419
ICOV1	.05373	.19051	.07155	-.08000	-.11426	.14486	.01710	.93904	-.04908
ICOV2	.08082	.06364	.11289	-.08551	-.08786	.15751	-.00927	.94490	-.02808
LTINV1	.73940	.40354	-.10594	.01561	.07716	.03821	.07374	.04145	.07169
LTINV2	.86491	.04622	-.03293	-.07955	.11723	-.03777	.06439	.01323	.01691
ITURN1	-.07766	-.52821	.14055	.54080	-.05151	-.12161	.47017	.06603	-.06480
ITURN2	-.16431	-.38228	.06061	.54741	-.00209	-.02963	.46087	.12697	-.27914
ART1	.17532	.08773	.31095	.22022	.73305	-.01523	-.18387	-.08543	-.10661
ART2	.20400	-.13127	.29693	.10274	.74279	-.22292	-.02759	-.02280	.03207
APT1	-.10776	.25432	.07479	.11086	.66257	-.03821	.10429	-.09518	.37787
APT2	-.11555	.16084	-.15513	.05683	.69774	.07018	.37009	-.09807	.17151
CURR1	.69758	-.09268	.02371	.41760	-.23871	.35464	-.09089	.08327	.00394
CURR2	.73888	-.24770	.13454	.34470	-.13551	.14787	-.09520	.10998	-.05689
ACID1	.78963	.30695	.04275	.00286	.05374	.35726	-.04923	.02138	.05194
ACID2	.88875	.01808	.13303	-.06699	.12201	.15342	.00291	.00759	.09650
BER1	.26339	.04560	-.01131	.10663	-.13375	.89455	.05360	.12515	-.04477
BER2	.28534	-.01749	.24563	.05566	-.11773	.85671	.01788	.12271	-.04508

Appendix H. Factor loadings of second benchmark (B)

	FACTOR 1	FACTOR 2	FACTOR 3	FACTOR 4	FACTOR 5
CONTPRA	.33619	-.07195	.66571	.16113	.08514
PROMARGA	.90562	.15611	.15529	.14977	-.09498
OPMARGA	.87184	.17299	-.05022	.18752	.26170
ASSTURNA	-.16387	-.34328	-.57904	-.41760	.23256
ROIA	.95104	.08023	.05145	-.07009	.03634
ROEA	.69481	-.08286	.01491	.24313	-.02529
AIRA	.19456	-.08167	.40358	-.30192	-.79263
ICOVA	.24081	.02303	.00719	-.26931	.73082
LTINVA	.13324	.77946	-.13902	.14356	.02634
ITURNA	-.17501	-.07518	.87789	-.05870	-.07848
ARTA	.17405	.20148	.03295	.70182	-.13061
APTA	.17444	-.14378	.10311	.86813	.03593
CURRA	-.07757	.87392	.23012	-.18546	.08170
ACIDA	.16434	.87952	-.15030	.13717	.18346
BERA	.01911	.48256	.16948	-.10942	.65580

Appendix I. Summary statistics of dependent variables.

Variable	Mean	Std. Dev.	Minimum	Maximum	K/S	N
Levels:						
Profitability	2.85	.88	1.25	4.75	0.08	75
Financing	2.82	.89	1.40	4.67	0.13**	75
Liquidity	2.96	.79	1.33	4.67	0.06	75
Leverage	2.72	.87	1.00	4.67	0.09	75
General	2.94	.78	1.25	4.33	0.07	75
Trends:						
Profitability	2.86	.89	1.40	4.50	0.11*	75
Financing	2.80	.73	1.40	4.33	0.08	75
Liquidity	2.87	.77	1.40	4.33	0.08	75
Leverage	2.76	.80	1.00	4.33	0.09	75
General	2.75	.85	1.00	4.33	0.13**	75

Summary statistics of independent variables: Mean, standard deviation, minimum value, maximum value, Kolmogorov-Smirnov statistic for tests of normality[1] (* and ** indicates significant at α=0.05 and 0.01 respectively), and number of observations.

[1] A Lilliefors version of the Kolmogorov-Smirnov test is used as in the EXAMINE procedure of SPSS.

Appendix J. Manipulation check for variables level and trend

*** A N A L Y S I S   O F   V A R I A N C E ***

GENERAL LEVEL
by CASECODE

Source of Variation	Sum of Squares	DF	Mean Square	F	Sig of F
Main Effects	185.907	74	2.512	4.951	.000
CASECODE	185.907	74	2.512	4.951	.000
Explained	185.907	74	2.512	4.951	.000
Residual	118.733	234	.507		
Total	304.641	308	.989		

*** A N A L Y S I S   O F   V A R I A N C E ***

GENERAL TREND
by CASECODE

Source of Variation	Sum of Squares	DF	Mean Square	F	Sig of F
Main Effects	215.501	74	2.912	7.224	.000
CASECODE	215.501	74	2.912	7.224	.000
Explained	215.501	74	2.912	7.224	.000
Residual	93.117	231	.403		
Total	308.618	305	1.012		

Appendix K. Typical relationships between learning and cross validated errors

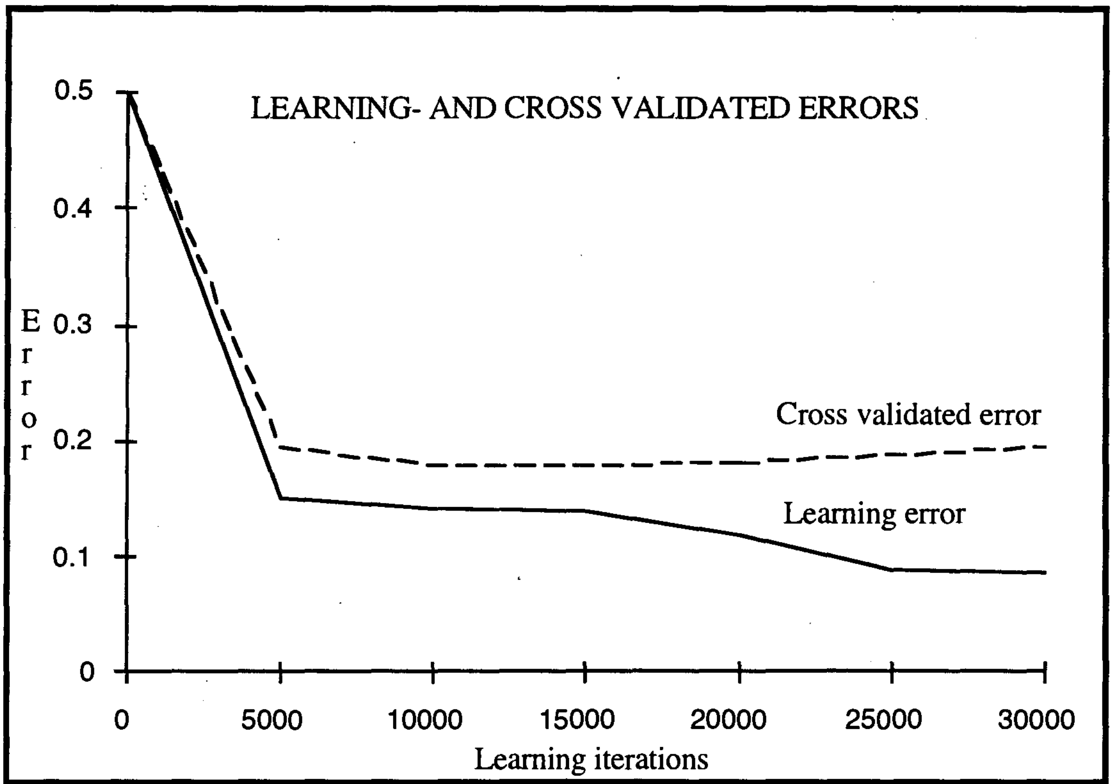

Learning error and cross validated error for level diagnosis of the combined model of chapter 8. The best model with two hidden units is used. Only errors for each 5000 learning iterations are reported

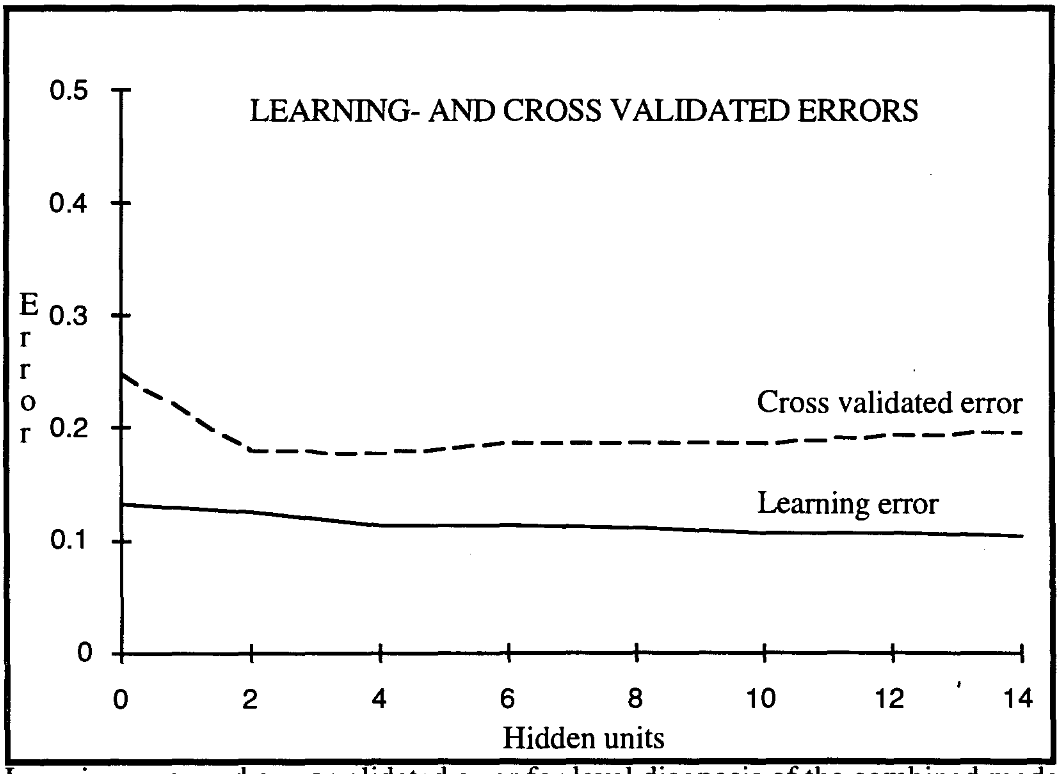

Learning error and cross validated error for level diagnosis of the combined model of chapter 8. The errors are reported after 10000 learning iterations. Only errors for every second hidden unit introduced are reported.