

Contents

Preface	x
Acknowledgements	xiv
CHAPTER 1 - CONNECTIONISM: AN HISTORICAL PERSPECTIVE	1
Early connectionism and the associationistic tradition	2
Formal neurons and their logical characteristics	9
Minsky and Papert and the EXOR problem	14
Early neural nets and pattern recognition: perceptrons	19
Perceptrons: Minsky and Papert's analysis	24
Diameter-limited perceptrons	29
Order-limited perceptrons	31
Perceptrons as computational devices	35
Coda	36
CHAPTER 2 - MEMORY AND LEARNING IN NEURAL NETWORKS	38
Parallel distributed processing and representation in neural networks	39
Distributed representation as coarse coding	41
Introduction to the formal properties of PDP systems	43
Preliminary comments about units and their interconnections	45
Weights and thresholds	46
Linear versus non-linear units	46

Pattern associators	49
Content-addressable memory	50
Learning in pattern associators	51
The Hebb rule	51
The delta rule	55
Back-propagation	56
Preliminary remarks	57
Gradient descent	58
Learning as gradient descent	59
Application of the back-propagation algorithm	66
Extensions to the standard back-propagation algorithm	68
Drawbacks and limitations of back-propagation	69
Probabilistic update of units	71
Hopfield nets	71
Boltzmann machines	76
Functional characteristics of the units in a Boltzmann machine	77
Concluding comments	81

CHAPTER 3 - ASPECTS OF VISION 83

Low-level vision	83
Visual information processing and levels of explanation	84
Stereopsis	85
Operation of the Marr–Poggio model	86
Appraisal of the Marr–Poggio model	87
New connectionism and the visual cortex	88
Description of Linsker's model	88
Appraisal of Linsker's model	92
Shape-from-shading	93
Architecture of the Lehky and Sejnowski model	93
Functional characteristics of the model	94
Appraisal of the Lehky and Sejnowski model	94
The biological implausibility of back-propagation	95
Edge extraction and figure–ground segmentation	97
The Kienker <i>et al.</i> model	97
Architecture of the model	98
Operation of the model	100
Grossberg's BCS/FCS Model	102
The model in detail	105
Appraisal of the model	107
Higher-level vision: object and shape recognition	109
Marr's account of object recognition	109
Hinton's (1981) account of shape recognition	112
Architecture of the model	112

Operation of the model	114
Hinton's model and the simulation of psychological phenomena	115
Hinton and Lang's simulations	116
General appraisal of Hinton's framework	119
Alternative connectionist accounts of object recognition	120
TRAFFIC: Object recognition as a process of 'transforming feature instances'	121
Appraisal of TRAFFIC	124
Object recognition on the basis of privileged views	126
The Poggio and Edelman model	126
The Edelman and Weinshall model	128
Concluding comments	130
CHAPTER 4 - ASPECTS OF LANGUAGE	132
Phoneme recognition	133
The model of Waibel and co-workers	133
Architecture of the model	134
Functional characteristics of the model	136
Initial appraisal of the model	138
General appraisal of the work on phoneme recognition	140
Context effects and the top-down/bottom-up distinction	142
TRACE	144
The architecture of TRACE	144
The operation of TRACE	147
General appraisal of TRACE	152
Massaro's critique of TRACE	154
New connectionist accounts of visual word recognition	157
Some functional characteristics of the IA model	158
General appraisal of the IA model	160
Interactive activation and word naming	162
Further issues concerning visual word recognition	163
Massaro's critique of the IA Model	166
Extensions to the early IA model of visual word recognition: processing more than one word at a time	167
CID	169
Architecture of CID	169
Functional characteristics of CID	170
BLIRNET	172
Critical appraisal of BLIRNET	175
The pronunciation of written words	177
Architecture of the Seidenberg and McClelland model	178
Training	180
Mature performance	181

Initial analysis of the Seidenberg and McClelland model	182
Lesioning the Seidenberg and McClelland model	186
General appraisal of the Seidenberg and McClelland model	189
Concluding comments	193

CHAPTER 5 - HIGHER-ORDER ASPECTS OF COGNITION 195

Accessing meaning from print: new connectionist accounts of deep dyslexia	196
Semantics 1: The Hinton and Sejnowski model	197
Architecture of the model	199
Operation of the model	200
Lesioning the model	201
Semantics 2: The Hinton and Shallice model	203
Architecture of the model	203
Semantic representation in the model	204
Lesioning the model	205
Further points about semantic classification	206
Category-specificity	207
Semantic classification without conscious identification	207
Concluding remarks about the psycholinguistics of lesioned nets	210
Double dissociations	210
Wood's account of double dissociations	212
Appraisal of Wood's model	214
Semantics 3: Distributed representations as lists of properties	217
McClelland and Rumelhart's model of prototype abstraction	217
Appraisal of the model	218
Smolensky's views on distributed representation	219
The implementation of classical symbols?	223
Hinton's account of conceptual structure	224
The role/filler distinction	226
Neural nets and reasoning	227
Context-independent representations	230
Limitations of the framework	230
Connectionist nets and rule-governed behaviour	232
McClelland and Rumelhart's model of past tense learning	233
Concluding comments	236

CHAPTER 6 - NEW CONNECTIONISM: A CRITICAL APPRAISAL 239

Neurological considerations	240
Crick and Asanuma's assessment of neurophysiological evidence	241
Additional hardware constraints	242
Problems with hidden units and the back-propagation algorithm	244

Massaro's critique	244
Norris's simulations	249
General morals	252
Stage-like models	252
Structured learning	253
Even more problems with the back-propagation account of human learning: Ratcliff's simulations	253
Error correction and human learning	255
Minsky and Papert's critique	257
Connectionism and thought: higher levels of cognition revisited	260
New connectionism and human reasoning	262
Concluding comments	265
References	267
Index	285