Psychology of Learning

Volume 1

SAGE LIBRARY *of*
COGNITIVE AND EXPERIMENTAL PSYCHOLOGY

SAGE LIBRARY IN COGNITIVE AND EXPERIMENTAL PSYCHOLOGY

PSYCHOLOGY OF LEARNING

VOLUME I

EDITED BY
DAVID SHANKS

SAGE

Los Angeles • London • New Delhi • Singapore • Washington DC

First published 2009

SAGE Publications Ltd
1 Oliver's Yard
55 City Road
London ECIY ISP

SAGE Publication Inc.
2455 Teller Road
Thousand Oaks, California 91320

SAGE Publications India Pvt Ltd
B 1/I 1 Mohan Cooperative Industrial Estate
Mathura Road, Post Bag 7
New Delhi 110 044

SAGE Publications Asia-Pacific Pte Ltd
33 Pekin Street #02-01
Far East Square
Singapore 048763

British Library Cataloguing in Publication Data

A catalogue record for this book is available from the British Library

ISBN: 978-1-84787-268-5 (set of five volumes)

Library of Congress Control Number: 2008924785

Typeset by Star Compugraphics Private Limited, Delhi
Printed on paper from sustainable resources
Printed in Great Britain by the MPG Books Group, Bodmin and King's Lynn

Appendix of Sources

All articles and chapters have been reproduced exactly as they were first published. All cross-references can be found in the original source of publication.

Grateful acknowledgement is made to the following sources for permission to reproduce material for this book.

1. 'Pavlovian Conditioning and Its Proper Control Procedures', *Robert A. Rescorla*
   Psychological Review, 74(1) (1967): 71–80.
   Copyright © 1967 by American Psychological Association. Reprinted with permission.

2. 'A Theory of Pavlovian Conditioning: Variations in the Effectiveness of Reinforcement and Nonreinforcement', *Robert A. Rescorla and Allan R. Wagner*
   A.H. Black and W.F. Prokasy (eds), *Classical Conditioning II: Current Research and Theory* (New York: Appleton-Century-Crofts, 1972), pp. 64–99.
   Published by Appleton-Century-Crofts, New York. Reprinted with the permission of the author.

3. 'A Theory of Attention: Variations in the Associability of Stimuli with Reinforcement', *N.J. Mackintosh*
   Psychological Review, 82(4) (1975): 276–298.
   Copyright © 1975 by American Psychological Association. Reprinted with permission.

4. 'A Model for Pavlovian Learning: Variations in the Effectiveness of Conditioned but Not of Unconditioned Stimuli', *John M. Pearce and Geoffrey Hall*
   Psychological Review, 87(6) (1980): 532–552.
   Copyright © 1980 by American Psychological Association, Inc. Reprinted with permission.

5. 'A Model for Stimulus Generalization in Pavlovian Conditioning', *John M. Pearce*
   Psychological Review, 94(1) (1987): 61–73.
   Copyright © 1987 by American Psychological Association, Inc. Reprinted with permission.

6. 'Distributed Memory and the Representation of General and Specific Information', *James L. McClelland and David E. Rumelhart*
   Journal of Experimental Psychology: General, 114(2) (1985): 159–188.
   Copyright © 1985 by American Psychological Association, Inc. Reprinted with permission.

7. 'Why There Are Complementary Learning Systems in the Hippocampus and Neocortex: Insights from the Successes and Failures of Connectionist Models of Learning and Memory', *James L. McClelland, Bruce L. McNaughton and Randall C. O'Reilly*
   Psychological Review, 102(3) (1995): 419–457.
   Copyright © 1995 by American Psychological Association, Inc. Reprinted with permission.

8. 'A Neural Substrate of Prediction and Reward', *Wolfram Schultz, Peter Dayan and P. Read Montague*
Science, 275 (1997): 1593–1599.
Copyright (1997) AAAS. Reprinted with permission from American Association for the Advancement of Science.

9. 'Context, Time, and Memory Retrieval in the Interference Paradigms of Pavlovian Learning', *Mark E. Bouton*
Psychological Bulletin, 114(1) (1993): 80–99.
Copyright © 1993 by American Psychological Association, Inc. Reprinted with permission.

10. 'Relation of Cue to Consequence in Avoidance Learning', *John Garcia and Robert A. Koelling*
Psychonomic Science, 4 (1966): 123–124.
Published by Psychonomic Society Publications. Reprinted with permission.

11. 'Specific Hungers and Poison Avoidance as Adaptive Specializations of Learning', *Paul Rozin and James W. Kalat*
Psychological Review, 78(6) (1971): 459–486.
Copyright © 1971 by American Psychological Association, Inc. Reprinted with permission.

12. 'Two-Process Learning Theory: Relationships between Pavlovian Conditioning and Instrumental Learning', *Robert A. Rescorla and Richard L. Solomon*
Psychological Review, 74(3) (1967): 151–182.
Copyright © 1967 by American Psychological Association. Reprinted with permission.

13. 'Auto-Shaping of the Pigeon's Key-Peck', *Paul L. Brown and Herbert M. Jenkins*
Journal of the Experimental Analysis of Behavior, 11(1) (1968): 1–8.
Copyright (1968) by the Society for the Experimental Analysis of Behavior, Inc. Reprinted with permission.

14. 'On the Law of Effect', *R. J. Herrnstein*
Journal of the Experimental Analysis of Behavior, 13(2) (1970): 243–266.
Copyright (1970) by the Society for the Experimental Analysis of Behavior, Inc. Reprinted with permission.

15. 'Selection by Consequences', *B.F. Skinner*
Science, 213(4507) (1981): 501–504.
Copyright © 1981 AAAS. Reprinted with permission from American Association for the Advancement of Science.

16. 'There Is No Convincing Evidence for Operant or Classical Conditioning in Adult Humans', *William F. Brewer*
W.B. Weimer and D.S. Palermo (eds), *Cognition and the Symbolic Processes* (Hillsdale, NJ: Erlbaum, 1974), pp. 1–42.
Reprinted with permission of Taylor & Francis Group LLC-Books, permission conveyed through Copyright Clearance Center, Inc.

17. 'Implicit Learning and Tacit Knowledge', *Arthur S. Reber*
Journal of Experimental Psychology: General, 118(3) (1989): 219–235.

18. 'A Case of Syntactical Learning and Judgment: How Conscious and How Abstract?', *Don E. Dulany, Richard A. Carlson and Gerald I. Dewey*
Journal of Experimental Psychology: General, 113(4) (1984): 541–555.

19. 'Preserved Learning and Retention of Pattern-Analyzing Skill in Amnesia: Dissociation of Knowing How and Knowing That', *Neal J. Cohen and Larry R. Squire*
Science, 210 (1980): 207–210.

20. 'Deciding Advantageously before Knowing the Advantageous Strategy', *Antoine Bechara, Hanna Damasio, Daniel Tranel and Antonio R. Damasio*
Science, 275 (1997): 1293–1295.

21. 'Acquisition of Cognitive Skill', *John R. Anderson*
Psychological Review, 89(4) (1982): 369–406.

22. 'Context Theory of Classification Learning', *Douglas L. Medin and Marguerite M. Schaffer*
Psychological Review, 85(3) (1978): 207–238.

23. 'Rule-Plus-Exception Model of Classification Learning', *Robert M. Nosofsky, Thomas J. Palmeri and Stephen C. McKinley*
Psychological Review, 101(1) (1994): 53–79.

24. 'From Conditioning to Category Learning: An Adaptive Network Model', *Mark A. Gluck and Gordon H. Bower*
Journal of Experimental Psychology: General, 117(3) (1988): 227–247.

25. 'Cue Competition in Causality Judgments: The Role of Nonpresentation of Compound Stimulus Elements', *Linda J. Van Hamme and Edward A. Wasserman*
Learning and Motivation, 25 (1994): 127–151.

26. 'Automatic and Effortful Processes in Memory', *Lynn Hasher and Rose T. Zacks*
Journal of Experimental Psychology: General, 108(3) (1979): 356–388.

27. 'Toward an Instance Theory of Automatization', *Gordon D. Logan*
Psychological Review, 95(4) (1988): 492–527.

28. 'The Neural Basis of Cognitive Development: A Constructivist Manifesto', *Steven R. Quartz and Terrence J. Sejnowski*
Behavioral and Brain Sciences, 20 (1997): 537–556.

29. 'Psychobiology of Plasticity: Effects of Training and Experience on Brain and Behavior', *Mark R. Rosenzweig and Edward L. Bennett*
Behavioural Brain Research, 78 (1996): 57–65.

30. 'Motor Sequence Learning: A Study with Positron Emission Tomography', *I.H. Jenkins, D.J. Brooks, P.D. Nixon, R.S.J. Frackowiak and R.E. Passingham*
The Journal of Neuroscience, 14(6) (1994): 3775–3790.

31. 'Why Faces Are and Are Not Special: An Effect of Expertise', *Rhea Diamond and Susan Carey*
Journal of Experimental Psychology: General, 115(2) (1986): 107–117.

32. 'The Role of Deliberate Practice in the Acquisition of Expert Performance', *K. Anders Ericsson, Ralf Th. Krampe and Clemens Tesch-Römer*
Psychological Review, 100(3) (1993): 363–406.

33. 'Three Aspects of Cognitive Development', *Robert S. Siegler*
Cognitive Psychology, 8 (1976): 481–520.

34. 'On Learning the Past Tenses of English Verbs', *D.E. Rumelhart and J.L. McClelland*
D.E. Rumelhart, J.L. McClelland and the PDP Research Group (eds), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Vol. 2: Psychological and Biological Models* (Cambridge, MA: MIT Press, 1986), pp. 216–271.

35. 'Rules of Language', *Steven Pinker*
Science, 253 (1991): 530–535.

36. 'A Distributed, Developmental Model of Word Recognition and Naming', *Mark S. Seidenberg and James L. McClelland*
Psychological Review, 96(4) (1989): 523–568.

37. 'Finding Structure in Time', *Jeffrey L. Elman*
Cognitive Science, 14 (1990): 179–211.

38. 'Children's Sensitivity to Constraints on Word Meaning: Taxonomic versus Thematic Relations', *Ellen M. Markman and Jean E. Hutchinson*
Cognitive Psychology, 16 (1984): 1–27.

39. 'The Phonological Loop as a Language Learning Device', *Alan Baddeley, Susan Gathercole and Costanza Papagno*
Psychological Review, 105(1) (1998): 158–173.

40. 'Statistical Learning by 8-Month-Old Infants', *Jenny R. Saffran, Richard N. Aslin and Elissa L. Newport*
Science, 274 (1996): 1926–1928.

41. 'Origins of Knowledge', *Elizabeth S. Spelke, Karen Breinlinger, Janet Macomber and Kristen Jacobson*
Psychological Review, 99(4) (1992): 605–632.

42. 'Maturational Constraints on Language Learning', *Elissa L. Newport*
Cognitive Science, 14 (1990): 11–28.

43. 'Categories and Induction in Young Children', *Susan A. Gelman and Ellen M. Markman*
Cognition, 23 (1986): 183–209.

44. 'Can Decline in Adult Intellectual Functioning Be Reversed?', *K. Warner Schaie and Sherry L. Willis*
Developmental Psychology, 22(2) (1986): 223–232.

45. 'Learned Helplessness: Theory and Evidence', *Steven F. Maier and Martin E.P. Seligman*
Journal of Experimental Psychology: General, 105(1) (1976): 3–46.

46. 'Fears, Phobias, and Preparedness: Toward an Evolved Module of Fear and Fear Learning', *Arne Öhman and Susan Mineka*
Psychological Review, 108(3) (2001): 483–522.

47. 'A Modern Learning Theory Perspective on the Etiology of Panic Disorder', *Mark E. Bouton, Susan Mineka and David H. Barlow*
Psychological Review, 108(1) (2001): 4–32.
Copyright © 2001 by American Psychological Association, Inc. Reprinted with permission.

48. 'Learned Industriousness', *Robert Eisenberger*
Psychological Review, 99(2) (1992): 248–267.
Copyright © 1992 by American Psychological Association, Inc. Reprinted with permission.

49. 'Commitment, Choice and Self-Control', *Howard Rachlin and Leonard Green*
Journal of the Experimental Analysis of Behavior, 17(1) (1972): 15–22.
Copyright (1972) by the Society for the Experimental Analysis of Behavior, Inc. Reprinted with permission.

50. 'Evidence from Rats that Morphine Tolerance Is a Learned Response', *Shepard Siegel*
Journal of Comparative and Physiological Psychology, 89(5) (1975): 498–506.
Copyright © 1975 by American Psychological Association. Reprinted with permission.

51. 'The Neural Basis of Drug Craving: An Incentive-Sensitization Theory of Addiction', *Terry E. Robinson and Kent C. Berridge*
Brain Research Reviews, 18 (1993): 247–291.
Copyright © 1993 Elsevier Science Publishers B.V. Reprinted with permission from Elsevier via Copyright Clearance Center's Rightslink service.

52. 'Confirmation, Disconfirmation, and Information in Hypothesis Testing', *Joshua Klayman and Young-Won Ha*
Psychological Review, 94(2) (1987): 211–228.
Copyright © 1987 by American Psychological Association, Inc. Reprinted with permission.

53. 'Discrimination Learning with and without "Errors"', *H.S. Terrace*
Journal of the Experimental Analysis of Behavior, 6(1) (1963): 1–27.
Copyright (1963) by the Society for the Experimental Analysis of Behavior, Inc. Reprinted with permission.

54. 'New Conceptualizations of Practice: Common Principles in Three Paradigms Suggest New Concepts for Training', *Richard A. Schmidt and Robert A. Bjork*
Psychological Science, 3(4) (1992): 207–217.
Published by Blackwell Publishing. Reprinted with permission.

55. 'Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria', *Ido Erev and Alvin E. Roth*
The American Economic Review, 88(4) (1998): 848–881.
Published by American Economic Association. Reprinted with permission.

Editor's Introduction: Psychology of Learning

David R. Shanks

The origins of the psychological study of learning are firmly rooted in 19th Century comparative psychology. The impact of the theory of evolution on reflection about the mind and behaviour led to a common view that while evolution determined instincts – patterns of largely invariant behaviours triggered by specific stimuli – it was the capacity for learning and plasticity that explained the rich and flexible behaviours of intelligent organisms. The ability to learn, to be sure, was itself an evolved capacity, but adaptive behaviour could not be fully understood without a detailed description of the learning faculty. Thus, from the work of Lloyd Morgan and others in the 1890's began to emerge the psychological sub-discipline of learning theory.

Probably the greatest achievements of the field followed in the early part of the 20th Century when Pavlov and Thorndike began not only to develop experimental techniques for studying learning but also proposed some of the seminal theoretical ideas to explain what later came to be described as classical and instrumental conditioning. Through the 1930's to 1960's, the theoretical analysis of these forms of learning developed enormously in the work of Skinner, Spence, Hull, and many others. At the same time, a considerable body of research built up on more practical aspects of learning, such as the role of errors in learning. Thorndike in particular exemplifies the bridge between theoretical and applied research.

In the past 40 or so years, research on learning has diversified enormously. As well as a continuing thread of work on theories of conditioning, learning has now become a major focus in relation to such diverse topics as language acquisition, expertise, neurobiology, behavioural economics, concept acquisition, and many more. In this collection, I have assembled a set of articles that cover a broad range of classic contributions across these areas. Although the diversity in the modern field emphasizes the central importance of learning to psychology, it does of course create difficulties with respect to knowing where to draw the line between learning and other very related topics such as choice and memory. I have therefore only included topics whose relevance is predominantly with respect to learning. For instance, I have not included classic issues such as spacing effects (the finding that learning trials generally lead to better retention when spread across time rather than bunched together) as the principal relevance of this is in relation to memory rather than learning.[1]

The present chapter provides an introductory overview, attempting to highlight the historical context that gave each article its unique significance. Although any effort to divide the area up into neat domains is bound to be subjective and to cut across important linkages, I have grouped the articles into 4 areas: learning theory and connectionism; learning and cognition; development, language acquisition, and plasticity; and practical aspects of learning.

Learning Theory and Connectionism

The associationist interpretation of learning has dominated research for the last century. Up until the publication of the first article in this collection (Rescorla, 1967), it was generally assumed that learning was driven by contiguous pairings of events. Although the article at face value is predominantly concerned with a methodological issue – what is the appropriate control group for measuring conditioning? – Rescorla's real interest was in arguing that it is informational value or contingency rather than contiguity that is critical. In the context of Pavlovian conditioning, where a conditioned stimulus (CS, e.g., a tone) is paired with a motivationally-significant unconditioned stimulus (US, e.g., food), conditioned responses (CRs) develop when the CS and US are correlated, regardless of whether contiguous pairings occur. Thus even when pairings are present, learning will not occur if there are additional presentations of the CS alone and US alone which render the contingency or correlation between CS and US zero. Ever since this pioneering analysis, the goal of learning theory has shifted: it is now taken to be the problem of understanding how information or contingency is extracted from temporally-distributed series of events.

Rescorla went on from this to develop with Wagner his well-known theory of learning (Rescorla & Wagner, 1972), one of the great contributions to modern psychology. This theory describes how the strength of association between a CS and US develops over trials, and explains how it is that raw contiguous pairings can yield information about contingency. At the heart of the theory is what has subsequently come to be described as "error-correction", the notion that learning depends on the extent to which there is a discrepancy between the expected and observed outcome of a learning trial. This concept has been central to many subsequent theoretical embellishments and developments of the Rescorla-Wagner theory. Amongst the theory's great accomplishments has been the ability to explain many selective learning phenomena, whereby learning about one event may compete with or block learning about other events.

Although they played little explicit role in Rescorla and Wagner's theory, attentional processes have been known for many decades to be important for learning. Mackintosh (1975) presented an alternative model of learning which explained selective learning by variations in attention rather than competition for limited associative strength. Mackintosh assumed that attention is a limited resource, and hence an informationally-significant stimulus may restrict attention to, and learning about, other stimuli. Pearce and Hall (1980) made diametrically the opposite assumption: that attention wanes to stimuli which have stable predictive value. Thus even a CS which is perfectly correlated with a US will lose attention in the Pearce-Hall theory. It is no exaggeration to say that a large proportion of basic learning research since these theories were published has been devoted to discriminating between them.

An important orthogonal development relates to the representation of stimulus compounds and the nature of generalization between stimuli. The Rescorla-Wagner theory assumed that a stimulus compound, AB, composed of two distinct stimuli A and B, would have an associative strength equal to the linear sum of A and B's associative strengths. Although Rescorla and Wagner themselves recognised early on the weakness of this assumption, it continued to be a useful simplifying assumption until Pearce (1987) proposed a configural model of learning in which no intrinsic relationship is assumed between compounds and the stimuli that comprise them. Pearce's model functions very like an exemplar theory of the sort described in the next section.

Pearce later couched his model explicitly in connectionist terms (Pearce, 1994, 2002). Such models, in which aspects of cognition are explained via the parallel operation of numerous simple processing elements, began to be explored in the 1950's but came to much greater prominence in the 1970's and 1980's. An important early contribution is McClelland and Rumelhart (1985) who showed how a simple connectionist model could explain many aspects of the learning of both specific and general features of stimuli. They described, in particular, how multiple memory traces could be stored in a superimposed fashion in a single network, how strong responding to an unfamiliar prototype could emerge, and how abstract representations of familiar items like words can be dispensed with. McClelland, O'Reilly, and McNaughton (1995) later developed related ideas in the context of a neurally-inspired view of one-shot learning (controlled by the hippocampus) versus slower, gradual learning of statistical regularities (in the neocortex). This is an exceptionally wide-ranging article that presents formal ideas about consolidation, semantic memory, interference, amnesia, and other important topics.

Discrepancies between obtained and expected rewards lie at the heart of all of the models discussed above, but these accounts are noncommittal about the origin and basis of such discrepances. A detailed formal account of how prediction errors drive learning, via the activity of dopamine neurons, is provided by Schultz, Dayan, and Montague (1997). This article considers not only classical conditioning but also action selection in instrumental behaviour.

Much research on learning has been concerned with the processes that cause forgetting. Bouton's (1993) important article reviews work on this issue and presents an analysis of how changes in context (which are assumed to include the lapse of time) interact with learning and retention. The result is a theory employing a small number of conceptual principles and applicable to a range of learning phenomena such as extinction, counterconditioning, and latent inhibition.

Developments in the basic theory of learning have been heavily driven by ecological findings from animal experiments. One such example was Garcia and Koelling's (1966) observation of selective learning. Garcia and Koelling made the that if rats drank "bright-noisy" water (water accompanied by a light and clicker) and experienced a shock, then an aversion to the water developed, whereas no such aversion developed if the water was flavoured rather than bright and noisy. In contrast, an aversion was readily induced to flavoured but not bright-noisy water if the animals were made ill by a toxin or radiation. Thus, contrary to the widespread view that the laws of learning are general across events, some stimuli have greater intrinsic propensity to be associated with a given unconditioned stimulus than others. The idea of 'belongingness' – that there may be particular evolutionary affordances associated with particular CS-US pairings – was advanced and evaluated by Rozin and Kalat (1971) in their theoretical discussion of taste-aversion learning.

The relationship between classical (Pavlovian) and instrumental (Thorndikean) conditioning has been a central focus of discussion during the whole history of learning theory. Attempts to reduce classical conditioning to principles of instrumental reinforcement, or vice versa, had proven sterile by the time Rescorla and Solomon (1967) wrote their influential review of two-process theory. This argues that stimulus-response (S-R) associations strengthened by a reinforcer in instrumental learning are necessarily accompanied and mediated by concomitant CS-US associations. In particular, CRs energise or inhibit instrumental behaviour. A typical example would be the mediation of instrumental avoidance responding by a CR such as conditioned fear. Rescorla and Solomon's article outlines the several lines of evidence supporting this two-process approach.

Another ecological finding of considerable importance was Brown and Jenkins' (1968) discovery of autoshaping in pigeons, the finding that pigeons will come to peck a key that signals food reinforcement even though that behaviour has no instrumental value. Autoshaping has played an important role in the theoretical analysis of differences between classical and instrumental conditioning. Herrnstein's (1970) famous analysis of the law of effect placed the concept of reinforcement of instrumental actions on a firm footing which has influenced much subsequent work. Herrnstein proposed a simple relationship – the matching law – between response rate and reinforcement rate and generalized this law to choice as well as strength of responding. The relationship between matching and the maximization of reinforcements has since been a prominent topic in both animal and human learning and relates to topics such as addiction (see below). The more general role of consequences in establishing and maintaining behaviour is famously stated by Skinner (1981). This article served as a manifesto for the idea that learning should be considered alongside evolution from an explanatory perspective. Just as the consequences (fitness) of some adaptation serve to promote or hinder its propagation via natural selection, so Skinner argued that behaviour is selected by reinforcement.

Learning and Cognition

For much of the past century of scientific progress in the psychology of learning, a deep problem has been the relationship between learning and cognition, where the latter refers to the processing of conceptual information. As highlighted in the articles in the previous section, learning has often been viewed from a theoretical standpoint as based on simple component processes such as activation, attention, excitation, inhibition, weight change, reinforcement, and so on. How do these basic processes relate to meaning, reasoning, inference, and other 'higher-level' operations? One idea, originating with Thorndike and others, would be that elementary forms of learning are detached from thinking and cognition, and occur via the automatic strengthening of stimulus-response bonds. Put simply, this is a 2-system view of the mind. Alternatively, learning and cognition might emerge from a common set of underlying processes and be based ultimately on the same mental 'stuff'.

Many have argued forcefully for the latter perspective, that all mental processes, including complex symbolic cognitive ones, are ultimately realised in systems composed of simple processing units. The view that learning and cognition depend on a common set of processes was famously argued by Brewer (1974). However, in contrast to the idea that all of cognition can be subsumed within the elementary processes of connectionism, Brewer argued exactly the opposite. He proposed that even Pavlovian conditioning was actually controlled by symbolic, cognitive operations. In Brewer's analysis, there are no aspects of conditioning for which associationist principles are required.

For Brewer, there are no learning processes which proceed independently of awareness – all learning is cognitive and conscious. The field of implicit learning explicitly attempts to test this claim. As far back as the 1920's, Thorndike had argued that unaware learning could be demonstrated. A significant development in this area was the invention by Arthur Reber of an experimental task – artificial grammar learning – which allowed this form of learning to be studied experimentally. A broad range of findings from this task, together with an overarching theoretical context, is reviewed in Reber (1989). The artificial grammar learning task bears a close resemblance to language acquisition, in that it requires participants to try to learn a set of abstract rules governing the legal sequences of a set of symbols, much as the rules of a grammar determine legal sequences of words in a natural language. With this task, Reber was able to demonstrate that participants can learn to discriminate grammatical from ungrammatical sequences, but apparently without being aware of the grammatical rules on which their behaviour was based.

Reber's claims for implicit learning were challenged by Dulany, Carlson, and Dewey (1984). These authors suggested that participants in the artificial grammar learning task do in fact consciously learn and respond on the basis of rules, but that each participant's rule set is partial and not necessarily veridical. The debate concerning the reality of implicit learning has continued to this day, and much of the research has focused on neuropsychological findings.

For example, Cohen and Squire (1980) argued on the basis of data from individuals with amnesia for a distinction between declarative and procedural learning. The former is thought to be conscious and able to be 'declared', while the latter is unconscious and not available for verbal report. Although previous studies had demonstrated preserved learning in amnesia, these had largely involved perceptual-motor skills such as tracking. Cohen and Squire's important discovery was that even pattern-analyzing skills may be preserved, and this led them to a different conceptualization of the properties of the 2 learning systems ('knowing that' versus 'knowing how'). Similarly, Bechara, Damasio, Tranel, and Damasio (1997) presented evidence that individuals can learn to predict which of various choice alternatives is optimal, but that they do so on the basis of unconscious 'somatic marker' signals. These are bodily reactions to risky choices which can be detected physiologically even though individuals may not be aware of the reasons for their decisions.

The procedural/declarative distinction has proven very influential in cognitive psychology and was incorporated quite explicitly into a model of skill acquisition by Anderson (1982). This model employed a very specific architecture called ACT which is implemented in computer software for modelling human cognition and behaviour. Anderson built on the procedural/declarative distinction to explain aspects of skill acquisition and expertise. Anderson's work has had a major influence on the development of computer-based learning systems (e.g., expert systems) in many applied contexts.

One of the central components of cognition is conceptual knowledge, our knowledge about objects, events, things, and so on. The study of concepts is a large field in its own right, but one particular aspect makes substantial contact with learning, and this is the issue of how we acquire conceptual or category knowledge. We have already seen some indications of how connectionist models might learn concepts in McClelland and Rumelhart's (1985) article (ch. 6) which considered the relationship between prototypes and exemplars. Prototypes represent hypothetical mental abstractions corresponding to the average of a set of specific examples of a category that an individual has learned about. It is well-known that the prototype of a category tends to elicit very accurate categorization responses even when it is unfamiliar and this has been taken as evidence that prototypes are automatically extracted and mentally represented in the course of learning a new category. McClelland and Rumelhart argued that prototypes do not need to be extracted for this behaviour to emerge, and in this claim they were building on a seminal article by Medin and Schaffer (1978). These authors presented a range of arguments in favour of the view that categories are represented by memorization of their exemplars, with responding to new objects being based on their similarity to the stored exemplars. Hence category learning is just the accumulation of category exemplars.

Medin and Schaffer's theory formed the starting point for arguably the most substantial strand of cumulative theory development in the whole of cognitive psychology, which is particularly embodied in the work of Robert Nosofsky. In Nosofsky, Palmeri, and McKinley (1994) a model of classification learning is presented which moves away from the idea of exemplar representation and instead proposes that individuals form category rules which divide the stimulus space up into distinct response regions, with exceptions being separately coded. This model is shown to fit an enormous breadth of empirical findings.

The linkage between connectionist learning models and concept acquisition was re-asserted by Gluck and Bower (1988), but these authors also drew the connection between these two areas and the elementary process of conditioning. Thus the theoretical model they proposed is intended to have exceptionally broad scope, connecting high-level concept learning with basic learning processes. At the time their article was published, this connection was more tacit than explicit. When attention was drawn to the linkage, much controversy ensued which continues to this day.

We saw in the chapter by Rescorla and Wagner (ch. 2) that a simple error-driven learning rule can capture a broad range of features of conditioning, and that this rule has also provided much stimulation in the study of human learning (ch. 24). The article by Van Hamme and Wasserman (1994) was important in three respects. This research showed that an important form of human learning, namely causal inference, can be illuminated by associative theories. Although Van Hamme and Wasserman were not the first to make this link, their study provides an elegant example of how human learning analogues of animal conditioning experiments may be designed. Secondly, this work demonstrated a significant boundary condition on the Rescorla-Wagner theory, in that changes in associative strength appear to take place even for cues that are absent on the learning trial. Again, although there had been precursors to this result, Van Hamme and Wasserman's demonstration was the first to show these effects on a trial-by-trial basis. Lastly, these researchers were able to formulate an elegantly simple extension of the Rescorla-Wagner theory which allows their data to be explained. In this modified version of the theory, absent cues are assigned negative values of the learning rate parameter.

Much subsequent work has followed up and confirmed that 'retrospective' changes in associative strength can occur in both humans and animals. It is symptomatic of the interplay between research in animals and humans that a theory originally developed to explain animal conditioning was subsequently applied to human learning where a phenomenon was found that was inconsistent with the theory. When researchers then sought to obtain that same phenomenon in animal behaviour, they were eventually able to find evidence of it, providing further reassurance that the laws of learning – however complex they may be – have some generality across species.

Research in learning has always had a close connection to the study of attentional processes, as described earlier. One influential claim about the automaticity of learning was made by Hasher and Zacks (1979), namely that the encoding of such things as the spatial properties and frequencies of events can occur automatically. Frequency information underpins a range of other judgments and behaviours such as estimates of risk, so this claim and the research it initiated have been of considerable significance in cognitive psychology. Hasher and Zacks provided a detailed exposition of what they considered to be the defining features of automatic and effortful processes, and speculated about individual differences in these processes. A related paper by Logan (1988) proposed a particular theoretical framework for thinking about automatic processes, via the same accumulation of instances that we have seen so prominently in work on category learning. Logan showed that his model accounted for important regularities in learning such as the power law of practice. In contrast to the common view that a lack of knowledge characterizes novice performance, Logan claimed that it is lack of attentional resources that more often marks the early stages of learning.

Development, Language Acquisition, and Plasticity

At the heart of research on learning and development over the past 50 or so years has been the central issue of innate versus acquired capacities. A seminal contribution to this debate is Quartz and Sejnowski's (1997) theoretical article which offers a neurally-inspired view of development in which the brain is in constant interaction with the environment and the learning input to extract regularities and enrich its representations so as to handle more complex inputs. These authors argued against domain-specific modules, a topic of much controversy in psychology.

Issues of brain plasticity and its relationship to training and experience are described in an important article by Rosenzweig and Bennett (1996) summarizing several decades of research on this topic, and concluding that enriched experience can induce a cascade of brain effects and behavioural outcomes. The understanding of brain plasticity has been enormously enhanced by the development in recent years of neuroimaging techniques such as PET (positron emission tomography) and fMRI (functional magnetic resonance imaging). Jenkins, Brooks, Nixon, Frackowiak, and Passingham's (1994) article is one of the earliest reports of mechanisms of motor sequence learning in the human brain using PET. These authors, who were amongst the pioneers in the technical development of this imaging method, found both increases and decreases in neural activity for learning motor sequences, in areas including the cerebellum and prefrontal cortex.

Diamond and Carey (1986) consider our ability to rapidly and efficiently recognize faces and asked whether this is evidence of a set of special-purpose mechanisms or whether instead it is the result of high levels of expertise. They note that face recognition, in relation to recognition of other common objects like houses, is particularly harmed by inversion (turning the object upside down). To understand the role of general learning versus a specific face module, they compared face recognition in normal participants with the recognition of dogs by dog experts. The high recognition ability of dog experts is, of course, due to learning. At variance with the special-purpose account, Diamond and Carey found that dog experts (but not novices) were severely