Schneider W and Bjorklund DF (1998) *Memory*. In: Kuhn D and Siegler RS (eds) *Handbook of Child Psychology*, vol. 2, *Cognition, Perception, and Language*, 5th edn, pp. 467–521. New York, NY: John Wiley.

Schneider W and Pressley M (1997) *Memory Development between 2 and 20*, 2nd edn. Mahwah, NJ: Lawrence Erlbaum.

Memory, Long-term

Introductory article

Alan Richardson-Klavehn, Goldsmiths College, University of London, London, UK
Robert A Bjork, University of California, Los Angeles, USA

CONTENTS

Long-term memory is central to cognitive functioning. Taking a wide variety of forms, from skills to general knowledge to memory for personal experiences, it is characterized by dynamic interactions between encoding and retrieval processes and by constructive processes, and thus differs fundamentally from current human-made information storage systems.

DEFINITION AND CLASSIFICATION OF LONG-TERM MEMORY

The ability to retain information over long periods is fundamental to intelligent thought and behavior. Memory is the 'glue', in effect, that holds our intellectual processes together, from perception, attention, and language, to reasoning, decision-making, and problem-solving. Memory also plays a critical role in social and emotional functioning, because our sense of who we and other people are is distilled from factual and autobiographical information in our memories. The study of memory, therefore, occupies a central position in the cognitive sciences.

Broadly, memory can be broken into three stages of information processing: (1) encoding, the transformation of information into a form retainable in memory; (2) storage, the holding of information in memory across a time interval; and (3) retrieval, the accessing of information from storage after a time interval and the use of that information to guide thought and behavior. This distinction between stages is important but – as clarified later – encoding and retrieval processes are intimately interconnected and cannot be understood in isolation from each other. (*See* **Information Processing**)

Distinguishing between Short-term and Long-term Memory

In everyday discourse, long-term memory is usually distinguished from short-term memory in terms of the time that has elapsed since information was encoded. Moreover, it is not unusual to find memory that persists over days or weeks being described as short-term memory. In psychology, however, the terms long-term and short-term memory have come to have specialized meanings that stem from a distinction made by William James in 1890. James observed that our consciousness is not just of the immediate present: the information that we currently attend to lingers in consciousness for some period of time. He called this lingering consciousness primary memory, and distinguished it from secondary memory, which occurs when information has left consciousness but returns to it again later. Thus, secondary memory involves retrieval in a way that primary memory does not.

James's distinction is not simply one of retention interval. It would, theoretically, be possible to retain information in primary memory indefinitely as long as one's attention remained focused on that information (i.e., as long as the information was rehearsed). Conversely, information that leaves consciousness and then returns to it is retrieved from secondary memory, even if retrieval occurs only seconds later. Thus the distinction between short-term (primary) and long-term (secondary) memory as used by psychologists is a qualitative one, not a simple quantitative one based on retention interval. In recent years, Alan Baddeley has introduced the term working memory to refer to short-term memory, which emphasizes its role in manipulating – as well as maintaining in consciousness – a variety of kinds of information. (*See* **James, William**; **Working Memory**)

Varied Forms of Long-term Memory

As with the short- versus long-term memory distinction, the main distinctions between forms of long-term memory involve reference to consciousness. The distinction between declarative and procedural memory originated in computer science, where stored data structures were distinguished from stored programs specifying how the data were manipulated. Psychologists borrowed these terms to capture a distinction made by the philosopher Gilbert Ryle in 1949, between *knowing that* and *knowing how*. Declarative memory involves knowing consciously *that* particular events happened in one's past, or *that* particular facts are true (e.g., Paris is the capital of France). Procedural memory, on the other hand, involves knowing *how* to manipulate mental or physical objects. Such knowledge is not necessarily consciously accessible and very difficult to communicate verbally. Explaining to someone how to ride a bicycle, for example, offers them scant assistance in learning that skill. Practicing such a skill is essential to its learning. The declarative/procedural distinction is closely associated with John R. Anderson and Larry Squire. (*See* **ACT**; **Skill Acquisition: Models**; **Automaticity**; **Implicit Learning**; **Skill Learning**; **Knowledge Representation, Psychology of**)

Within declarative memory, episodic memory is distinguished from semantic memory. Episodic memory involves awareness of particular events in one's personal autobiography, whereas semantic memory involves knowledge of language, categories and concepts, and facts. This distinction is closely associated with Endel Tulving. Within episodic memory, in turn, recollection is distinguished from familiarity. Recollection involves re-experiencing the particular contextual details of a past event, such as the tone of voice in which a statement was uttered in the kitchen at nine o'clock yesterday morning. Familiarity involves the knowledge that a current situation bears some relationship to a past event, without awareness of the particular contextual details of that event. For example, we sometimes experience the strong sense that we have met someone before, without being able to recollect where and when we met them, or anything else about them. The recollection/familiarity distinction is closely associated with George Mandler and Larry Jacoby. A related distinction, between remembering and knowing, has been made by Tulving and by John Gardiner. (*See* **Semantic Memory: Computational Models**; **Episodic Memory, Computational Models of**; **Autobiographical Memory**; **Knowledge Representation, Psychology of**)

A final important distinction is between explicit and implicit memory. Explicit memory refers to conscious awareness of events in one's personal past that accompanies deliberate attempts to think back to those events. Implicit memory refers to influences of past events on one's current behavior that occur involuntarily or unintentionally, often without any current awareness of the relevant prior events. This distinction, closely associated with Daniel Schacter and Peter Graf, can be traced back to similar distinctions by Hermann Ebbinghaus, who published the first experimental studies of memory in 1885, and by a number of other influential thinkers going back to René Descartes in 1649. (*See* **Descartes, René**; **Ebbinghaus, Hermann**)

It must be noted that none of the foregoing distinctions is universally accepted. First, none is entirely clear-cut. For example, Paul Kolers and Henry Roediger have questioned the procedural/declarative distinction, arguing that all forms of memory involve the modification of procedures for manipulating information. And Schacter, Alan Richardson-Klavehn and others have pointed out that the explicit/implicit memory distinction is blurred by cases when conscious awareness of events in one's personal past comes about without any deliberate attempt to retrieve those events, a phenomenon termed involuntary explicit memory or involuntary conscious memory.

Second, a controversial question is whether these distinctions imply different information-processing mechanisms, with different bases in the structure and function of the brain. Support for the latter view comes from research by Brenda Milner, Elizabeth Warrington, Lawrence Weiskrantz and others on the amnesia (memory loss) that results from damage to limbic system structures in the brain (the hippocampus, portions of the thalamus, and connected structures). This memory loss is selective, resulting in dissociations between different measures of memory. For example, short-term memory is largely spared, whereas the acquisition of new long-term memories is severely impaired, and the declarative and explicit forms of long-term memory are impaired much more than the procedural and implicit forms. Dissociations that are similar in some respects are observed in dementias such as Alzheimer and Huntington diseases, as well as in the memory loss that accompanies normal aging. Such dissociations have led some, including Schacter, Squire, and Tulving, to argue that the brain has distinct memory systems, which may have had different evolutionary histories. Others, such as Kolers, Roediger, Mary Sue Weldon, and Bruce Whittlesea, argue that a unitary memory system – in which similar information-processing mechanisms handle a wide variety of kinds of information – can account for such dissociations.

As argued by Morris Moscovitch and others, resolving these issues will involve clarifying the extent to which the varied forms of memory involve different versus common information-processing components and on understanding the relationship between these components and brain structure and function. Recent advances in imaging the activity of the living brain (neuroimaging) are making an important contribution in these respects. Whatever their interpretation, however, the selective memory impairments that have fueled the current controversies offer a striking illustration of the complexity and variety of memory, and of the centrality of memory to intellectual and social functioning. Such memory impairments often have a catastrophic effect on an individual's ability to hold down a job, remain informed about ongoing affairs in the world, and maintain normal social relationships. (*See* **Human Cognition**; **Aging and Cognition**; **Memory: Implicit versus Explicit**; **Memory, Development of**; **Neural Basis of Memory: Systems Level**; **Amnesia**; **Alzheimer Disease**; **Huntington Disease**; **Neuroimaging**)

THE DYNAMIC CHARACTER OF LONG-TERM MEMORY

The remainder of this article focuses on the cognitive processes involved when new information is added to long-term memory and later retrieved. At first thought, libraries and computers might seem useful metaphors for understanding these processes. In computers, for example, files are created (encoding), held on disk (storage), and subsequently made active again (retrieval). Such metaphors, however, can be highly misleading. They suggest that encoding and retrieval are strictly sequential, and that encoding new information does not involve retrieving information that is already stored. With human memory, by contrast, encoding new information *depends* on retrieval of information already in memory. The computer metaphor also suggests that the act of retrieving an item makes that item no more and no less accessible in the future, and does not affect the accessibility of other items. With human memory, by contrast, retrieval renders the retrieved information more accessible in future, and can have either positive or negative effects on the retrievability of other information, depending on circumstances. Furthermore, such metaphors suggest that memories are stored in specific spatial locations, whereas human memories appear not to be stored in specific locations in the brain, but in distributed networks of brain cells (neurons), each of which participates in the storage of many memories.

The key to understanding the unique properties of human long-term memory is to appreciate that it has a dynamic character not shared by current human-made information storage systems. That is, the state of memory is constantly changing as a result of encoding and retrieval processes that are intimately interdependent. This unique character is a product of the properties of the brain as an information-processing device and may reflect the evolution of memory from the perceptual mechanisms of the brain.

The Interdependence of Encoding and Retrieval

Levels of processing, encoding specificity, and resonance

Research from the 1970s onwards has greatly enhanced our understanding of encoding and retrieval processes in long-term memory. One important principle to emerge is that the primary determinant of long-term retention is the level of cognitive processing when new material is encoded – irrespective of intention to learn, or amount of repetition or rehearsal, both of which have little impact on long-term retention. Shallow levels of processing involve attending to the physical characteristics (typically, appearance or sound) of material, whereas deep levels of processing involve attending to the meaning of material, with deep processing usually resulting in superior retention. Both shallow and deep processing involve retrieving pre-existing knowledge about appearance, sound, or meaning; the resulting memory trace is a by-product of the processing involved in retrieving that knowledge. This levels-of-processing principle originated with Fergus Craik and Robert Lockhart.

A second principle to emerge is that when encoded material is later retrieved, the stimuli present in the retrieval environment (retrieval cues) also play an important role in whether that material is retrievable. For successful retrieval, it is critical that the information provided by the retrieval cues matches the information in the memory trace, which will in turn reflect the type of processing engaged at encoding. This encoding specificity principle originated with Endel Tulving and Donald Thomson.

The interaction between level of processing at encoding and the cues present at retrieval is illustrated by the results of an experiment reported by Ronald Fisher and Fergus Craik. They asked people to study pairs of words, with the words in the pairs related either by meaning (e.g., *sleet–hail*) or by sound (e.g., *pail–hail*). The meaning relationship led to a deep level of processing, whereas the sound relationship led to a shallow level of processing. Later, people's ability to recall the second word from each pair was tested via a cued recall test, in which cues were provided to assist with retrieval. The cues either involved the first word in the pair presented at encoding (e.g., *associated with sleet* and *rhymes with pail*, respectively), a similar word (e.g., *associated with snow* and *rhymes with bail*, respectively), or a different word (e.g., *rhymes with bail* and *associated with snow*, respectively).

When the original first word in the pair was presented as a retrieval cue, there was a substantial recall advantage for deep (meaning) over shallow (sound) processing at encoding (54 percent versus 24 percent), but when the cue was a similar word the advantage for deep processing was reduced (36 percent versus 18 percent), and when the cue was a different word the advantage was almost eliminated (22 percent versus 16 percent). This result demonstrates that both the level of processing at encoding and the cues present at retrieval are critical. If it were simply the case that deep processing produces stronger or longer lasting memory traces than shallow processing, then the advantage for deep processing would be observed regardless of retrieval conditions. And if, on the other hand, the only important factor is the match between retrieval cues and memory traces, then increasing the degree of match should benefit performance regardless of the level of processing at encoding. Instead this result suggests that deep processing produces memory traces containing information that is distinctive in comparison with the information contained in other memory traces. As a consequence, when a retrieval cue matches a trace that resulted from deep processing, only that trace is likely to be activated. By contrast, when a retrieval cue matches a trace that resulted from shallow processing, many other traces become active, because they also contain information matching the retrieval cue. In consequence, there is interference that impairs retrieval.

Combining the principles of levels of processing and encoding specificity leads to a more general principle that can be called the principle of selective resonance. The resonance idea is drawn from physics. A 440 Hz tone emitted near the undamped strings of a piano, for example, will lead to sympathetic resonance in the strings tuned to 440 Hz and, to a lesser extent, in the strings tuned to frequencies that have harmonic relationships to 440 Hz (e.g., 880 Hz, 220 Hz). Retrieval can be thought of as resembling resonance. Memory traces are 'tuned' to specific frequencies, based on the information encoded into them. At retrieval, they 'resonate' to the extent that they share information with retrieval cues. Retrieval succeeds when the resonance is unique to relevant traces, and not shared with irrelevant traces.

The selective resonance idea can be traced to the little-known memory theorist Richard Semon, who coined the term ecphory in 1921 to describe the process whereby memory traces (or engrams) resonate in response to retrieval cues, and to Hedwig Von Restorff, who demonstrated the importance of distinctiveness for memory in 1933. More recently, Roger Ratcliff has demonstrated that the resonance concept forms a realistic basis for formal mathematical models of retrieval. In addition, the principle of selective resonance is a natural property of recent models of memory that postulate networks of interconnected units analogous to networks of neurons in the brain (connectionist models). These models, developed by James A. Anderson, James McClelland and others, solve the conundrum of where memories are stored by postulating that they are stored in a distributed form, as changes in the connections between many units, with each unit participating in the storage of many memories. Finally, processes akin to selective resonance appear to be a fundamental property of neural information processing, starting with perception, where groups of neurons are 'tuned' to respond selectively to particular features of stimuli impinging on the senses. The resonance principle thus suggests that our memory capabilities may have evolved as an extension of our perceptual capabilities, with some of the underlying neural information-processing principles being carried through. (*See* **Human Cognition**; **Connectionism**; **Distributed Representations**; **Encoding**

and Retrieval, Neural Basis of; **Learning and Memory, Models of**; **Pattern Vision, Neural Basis of**; **Perceptual Systems: The Visual Model**; **Memory Models**; **Perception: Overview**)

Perceptual memory tests and transfer appropriate processing

Deep processing at encoding is typically superior to shallow processing in supporting long-term retention but – as discussed earlier – such superiority can disappear when the cues at test mismatch those present at encoding. Might there then actually be circumstances in which shallow processing is superior to deep processing? The answer is yes. Tests such as cued recall, recognition memory, in which people are asked to discriminate previously presented material from material not previously presented, and free recall, in which retrieval is not aided by external cues, show advantages for deep processing because succeeding on such tests requires semantic, meaning-related, processing. These tests are classified as conceptual tests. In special circumstances, however, when a memory test demands retrieval of information concerning the perceptual characteristics of previously encountered information (perceptual tests), the advantage of deep over shallow processing can be reversed. For example, Barry Stein showed that shallow processing at encoding was superior when people later had to recognize whether or not words had been shown to them in a particular configuration of upper and lower case letters. Such reversals can occur because shallow processing results in memory traces that contain more distinctive information about the perceptual – as opposed to semantic – characteristics of studied items than does deep processing.

In the light of such findings, John Bransford, Jeffrey Franks and others have argued that the levels-of-processing principle – which can be taken to imply universally superior retention for deep levels of processing – should be reformulated as the transfer appropriate processing principle. This principle retains Craik and Lockhart's fundamental insight that the content of the memory trace is a by-product of cognitive processing at encoding, but it asserts that the value of a particular level of processing at encoding is relative to the kind of processing that is later required at retrieval. (*See* **Memory: Implicit versus Explicit**)

Retrieval as an encoding event

The fundamental interplay of encoding and retrieval processes is further illustrated by experiments on the impact of retrieval on later memory performance. The generation effect, first systematically explored by Norman Slamecka and Peter Graf, is a particularly good example. Inducing people to actively generate items from semantically related cues (e.g., *horse-c_ _t*, the generated item being *cart*) produces better later memory for those items than does simply reading them (e.g., *horse-cart*). Generating, like deep processing, creates more semantically distinctive memory traces than does reading. Another illustration is the impact of retrieving newly acquired material on the later retrieval of that material. Thomas Landauer, Robert Bjork and others have demonstrated that testing people on newly acquired material – versus providing an additional exposure to the material – can result in superior later recall. Moreover, provided that recall succeeds, the more difficult or involved the recall is, the greater its positive effects on later recall. Thus retrieval modifies the state of memory, acting as an additional encoding event, such that retrieved material is rendered more accessible later. These positive effects of retrieval are known as test effects. (*See* **Memory**; **Learning Aids and Strategies**; **Education, Learning in**)

Environmental, mental state, and temporal context effects

The principle that reinstating the kind of cognitive processing engaged at encoding is critical for later retrieval extends to the influence on cognitive processing of environmental context (e.g., location and other ambient environmental stimuli) and of mental state context (e.g., mood states and drug states). Alan Baddeley, Eric Eich, Steven Smith and others have shown that retrieval is often less successful if these forms of context are changed between encoding and retrieval. However, such context effects are not always observed. They appear to be more likely when retrieval occurs in the absence of explicit cues, as in free recall tests; when people are unable to mentally reinstate the context present at encoding; and when the context either becomes explicitly associated with the to-be-remembered material at encoding or exerts an explicit or an implicit influence on the semantic interpretation of the material.

Changes in context can have positive as well as negative effects, as revealed in research by Smith, Arthur Glenberg, Robert Bjork and others. Material encountered on multiple occasions in different contexts is more retrievable later than is material always encountered in the same context. In addition, material encountered on multiple occasions is much more retrievable later when those occasions are spaced apart in time rather than massed together, a temporal context effect termed the spacing effect. These benefits appear to arise, in part, from a common mechanism: encoding variability. That is, changes in context across encounters vary the kind of cognitive processing involved when the material is encoded. In the case of spaced repetition, such variation occurs as a result of a drift in environmental and mental state context over time – an idea first formalized by William Estes in his 1955 stimulus fluctuation theory. Variable encoding benefits memory because it increases the likelihood that some aspect of the cognitive processing engaged at retrieval will match information in the memory trace.

Another factor in the enhanced retention that results from encoding variability is that retrieval is also an encoding event. On the second and subsequent encounters with the material, it is necessary – if the repetition of the material is to count as such psychologically – that the material is recognized as having been encountered earlier. Such recognition is more difficult, and thus the retrieval involved more powerful as an encoding event, when context changes across the successive encounters with the material. This retrieval-based explanation of the benefit of encoding variability again illustrates the intimate relationship between encoding and retrieval processes. (*See* **Memory**; **Learning Aids and Strategies**; **Learning, Psychology of**; **Education, Learning in**; **Mathematical Psychology**)

Forgetting, Interference, and Inhibition

Any theory of memory must explain why information is often forgotten over time. The most obvious hypothesis relies on a further metaphor for memory, and the one that is perhaps most intuitive: memory traces are like characters engraved on a stone or wax tablet, or like footprints in the sand, and these imprints weather away over time. Stated more scientifically, the hypothesis is that memory traces have strengths that decay with time. As with the strength interpretation of level-of-processing effects, however, this hypothesis cannot explain forgetting. One illustration of the inadequacy of this strength–decay idea is the recognition failure of recallable words, a phenomenon discovered by Endel Tulving and Donald Thomson. They showed that words (e.g., *queen*) studied in the context of weakly related words (e.g., *woman–queen*) were often forgotten on a recognition test when presented in the context of strongly related words (e.g., *king–queen*). However, these forgotten items were often successfully retrieved on a subsequent cued recall test when the weakly related word presented at encoding (e.g., *woman*) was provided as the cue.

The importance of this finding is that the original recognition failure cannot be attributed to decay of trace strength, because it would then have been impossible for the items to be retrieved on the later – and nominally more difficult – cued recall test. Instead, the recognition failure occurred because the target word was presented in a changed associative context in the recognition test – a powerful demonstration of the encoding specificity principle. John McGeoch was the first to argue, in 1932, that forgetting from long-term memory, rather than being a consequence of the decay of memory traces, reflects an inability to retrieve those traces. He argued that such retrieval failures are caused by (1) changes in associative, environmental, and mental-state context over time that he termed altered stimulating conditions, and (2) interference between competing traces in memory. McGeoch's two-factor theory of forgetting has stood the test of time, and fits well with the resonance conception of retrieval described earlier.

Such interference effects were a major focus of research on human learning carried out in the behaviorist tradition (ca. 1900–1970). Conclusions from this research are (1) that interference is the greater the more the similarity in content between the interfering materials; (2) that new learning interferes with old learning (retroactive interference), but, that as time passes, old learning recovers to interfere with the new learning (proactive interference); and (3) that interference can take the form of competition between the sets of materials, evidenced as a confusion at retrieval about which set of materials is which, or as unlearning of the materials, evidenced as the inability to bring a particular set of materials to mind.

Recent research suggests, in addition, that the process of retrieving information from memory can itself cause forgetting. Successful retrieval, as discussed earlier, makes the retrieved material more accessible later, but at an apparent cost: other material associated to the cues guiding retrieval can become less accessible later. Michael Anderson, Robert Bjork, Elizabeth Bjork, and Barbara Spellman have demonstrated such retrieval-induced forgetting by showing that the repeated retrieval of particular members of a category of previously studied words can inhibit subsequent access to the other nonretrieved members of that category. Such inhibited access is apparently a consequence of the need to suppress – that is, not overtly respond with – those items during the earlier retrievals of the target items.

Retrieval-induced forgetting and related effects point to another dynamic property of memory. In order to avoid catastrophic interference as a result of the large amount of information stored, only a limited portion of the information can be accessible for retrieval at any given time. From an adaptive point of view, the portion of information that is most accessible at any given time should be the portion retrieved in the recent past, because that information is most likely to be relevant in the near future. When forgetting is viewed in this way, it can be seen to be essential to the efficient functioning of memory, and thus far from a negative phenomenon. This adaptive theory of interference, inhibition, and forgetting has been formulated as a new theory of disuse by Robert and Elizabeth Bjork, integrating and extending ideas put forward by Edward Thorndike in 1914 and William Estes in 1955. (*See* **Learning, Psychology of**; **Rational Models of Cognition**)

THE CONSTRUCTIVE CHARACTER OF LONG-TERM MEMORY

The library and computer storage metaphors for human memory, which are misleading for the reasons suggested earlier, are misleading in another respect. They suggest that the storage capacity of human memory is gradually used up as more material is stored. By contrast, there is no known limit to the human capacity to acquire new information. Indeed, acquiring new information in the form of organized knowledge *creates* further capacity. This ever-expanding capacity reflects constructive processes that are unique to human memory. While these processes have the positive effect of enabling the retention of astounding quantities of information, they can also have serious negative effects, by leading to memory distortions and illusions. Such negative effects show that human memory – in contrast to a videotape recorder (another misleading metaphor) – is not a literal record of previously encountered information. Even vivid memories of personally experienced events can be attributions of currently experienced mental events to the past.

Organization, Chunking, and Expertise

Encoding, as discussed earlier, involves bringing pre-existing knowledge in memory to bear on the interpretation of new information. To understand our unlimited capacity to acquire new information, Fergus Craik has suggested that memory, rather than being thought of as a library, computer, or videotape, should be thought of as a scaffolding. The scaffolding is the organized information in memory, which forms a framework for the interpretation of new information, and which permits new information to be attached. It follows that the more scaffolding there is, the greater the capacity to attach (encode) new information. Such organized information plays an important role in retrieval as well: it permits reconstruction of the likely properties of the material.

There are many experimental demonstrations of the powerful positive effects of organizing new material in terms of existing semantic knowledge. For example, Gordon Bower and his colleagues gave people four opportunities to study and freely recall 112 words drawn from various semantic categories. When the words were presented separated into the semantic categories (e.g., *minerals–metals–alloys: bronze, steel, brass; minerals–stones–precious: sapphire, emerald, diamond*), recall of this very large number of words was almost perfect by the second study and recall attempt, and perfect by the third. By contrast, when the identical 112 words were presented in a randomly intermixed fashion for the same amount of study time, recall reached only around 60 percent by the fourth study and recall attempt.

When confronted with new material to learn, as a student for example, a common difficulty is that the material appears to lack organization, and is therefore meaningless. Effective learning requires abstracting the structure of to-be-learned material. One important component of such abstraction is 'chunking', a term coined by George Miller in 1956. A famous example of the importance of chunking, originating with Karl Lashley in 1951, is the following French sentence: *Pas de lieux Rhône que nous*. Even to a French speaker, this sentence is not memorable, because it is nonsensical. However, sounding the sentence out a few times (with the correct French pronunciation) soon leads to a reorganization of its constituent units that renders it instantly memorable – but as an English sentence. Most apparent examples of learning by rote repetition – for example, learning of scripts by actors, and even of sequences of steps by dancers – actually involve some form of chunking.

Chunking also plays a central role in creating differences in memory ability between experts and novices in a particular field. Adriaan De Groot, for example, showed that expert chess players remembered chess positions much more accurately than novice chess players, but that this advantage was not attributable to the experts having better overall memory capabilities: when the chess pieces were randomly arranged, the memory difference between experts and novices disappeared. Instead, the experts' superior memory depended on their ability to chunk groups of pieces according to their knowledge of positions that might be expected to occur in a game.

Exceptional and apparently astounding memory abilities can also be acquired by developing sophisticated chunking strategies. For example, William Chase and Anders Ericsson trained an individual to recall over 80 sequentially presented random digits, even though he could initially recall only an average seven. His ability reflected strategies for grouping the numbers into meaningful chunks, not a general increase in 'memory power' with training: When the task was switched to remembering random letter sequences, he could once again recall only around seven. Naturalistic studies of memory experts who perform in public also usually confirm that they have learned sophisticated chunking strategies to perform their apparently photographic feats. Thus, memory is not like a muscle that can be 'strengthened' purely by repetitive exercise. (*See* **Lashley, Karl S.**; **Memory**; **Learning Aids and Strategies**; **Expertise**; **Memory Mnemonics**; **Education, Learning in**; **Learning and Memory, the Ecology of**)

Memory Distortions, Illusions, and Attributions

The inevitable role of existing world knowledge in the encoding and retrieval of new information has negative as well as positive consequences, as first reported by Frederick Bartlett in 1932. The then-prevailing experimental practice, originating with Ebbinghaus, was to attempt to examine memory for new information in isolation from pre-experimental knowledge by asking people to learn simple and often meaningless materials. Bartlett, by contrast, asked people to learn stories, which gave maximum scope for the influence of pre-experimental knowledge to be revealed. In recall of the stories, details were omitted, leaving memory for the gist, or main structural elements. Bartlett also found that new elements were introduced and existing elements distorted in accordance with knowledge – including cultural and social preconceptions – about the kind of events likely to have occurred in the story.

The kinds of general world knowledge that influenced and distorted the recollections of the participants in Bartlett's experiments are termed schemas. More recent experimental research with textual materials, as well as naturalistic research on memory for real-world events, conducted by Gordon Bower, John Bransford, Jeffery Franks, Ulric Neisser and others, has clarified the processes by which schemas produce these distortions and additions. At encoding, the specific information provided is elaborated in terms of schemas (e.g., assumptions and inferences are made), and these elaborations become part of the memory trace for the material. At retrieval, information provided by retrieval cues, by the schemas, and by specific information retrieved, which includes the elaborations made at encoding, combines to produce a reconstruction of the previously encountered information – one that can contain significant distortions. Once again, encoding and retrieval processes are intimately interwoven. (*See* **Bartlett, Frederic Charles**; **Schemas in Psychology**; **Memory Distortions and Forgetting**; **Learning and Memory, the Ecology of**)

Such distortions are also evident when people 'remember' recent well-circumscribed events that in fact never occurred. A striking recent example was reported by Henry Roediger and Kathleen McDermott, who updated an experimental procedure introduced by James Deese in 1959. People studied a number of lists of words, with the words within each list consisting of semantic associates (e.g., *mad*, *fear*, *hate*, *rage*, *temper*, *fury*) of a particular prototype word (e.g., *anger*), which was not itself studied. On a later recognition test, previously studied semantic associates were mixed with the nonstudied prototype words, and with other nonstudied words that were unrelated to the previously studied words (e.g., *bread*). People were well able to reject the unrelated words as not having been studied. But they were just as likely to endorse the nonpresented prototype words as having been studied as they were the actually presented words. Critically, they did not just guess that the nonstudied prototypes could plausibly have been studied. They not only 'recognized' them with high confidence, but also claimed to re-experience vividly the details of their prior occurrence (e.g., what they were thinking at the time). This false memory phenomenon, therefore, is not just a memory distortion, but a memory illusion.

This and other memory illusions illustrate that our understanding – or misunderstanding – of our own memory processes, which is termed metamemory, plays a critical role in long-term memory. There is considerable evidence, for example, that we monitor the fluency with which information is currently processed. When that fluency exceeds the fluency we would expect – based on our knowledge of how fluently that kind of information is normally processed – we face a problem of attribution. Where does that unexpected fluency come from? We could attribute it to current external conditions that make the information especially easy to process, the information being well established in memory, a recent encounter with the information, or some other factor.

These attributions about the source of current processing fluency are adaptive in the sense that they are usually valid. But they can sometimes be mistaken, causing striking memory illusions of various kinds. For example, Larry Jacoby and his colleagues found that increasing the identifiability of sentences spoken against a noise background, by presenting those sentences earlier in a supposedly unrelated task, causes a misattribution of that increased identifiability to a lowered level of the background noise. And Lynne Reder and her colleagues found that prefamiliarizing key words in general knowledge questions, such as the words *golf* and *par* in the question *What is the term in golf for scoring one under par?*, increases the likelihood that people judge that they know the answer to the questions – but without improving their actual ability to answer them. Similar illusions of knowledge, based on general familiarity with a subject domain, can be evident in students' judgments of their comprehension and future memory of textual material, as shown by Arthur Glenberg, William Epstein and their colleagues.

Finally, with regard to the illusory recognition phenomenon described earlier, Bruce Whittlesea has recently theorized that there is unexpected fluency in the semantic processing of the nonpresented prototypes when they appear on the recognition test, caused by the prior presentation of their associates. This unexpected fluency results in the automatic construction of vivid 'recollections' of earlier encounters with the prototypes. Whether or not this particular theory stands the test of time, such attributional theories of memory are important more generally because they raise the possibility that all memories – whether veridical or illusory – are constructions based on current cognitive processing. (*See* **Memory**; **Learning Aids and Strategies**; **Metacognition**; **False Memory**; **Education, Learning in**)

CONCLUSION

The overall picture that emerges from just over a century of scientific research is that human long-term memory is exceptionally complex and sophisticated: it is varied, dynamic, and constructive, and quite unlike current human-made memory devices in virtually every important respect. The resulting capacities of human long-term memory are stunning, which can make its limitations – in terms of forgetting, distortions, and illusions – seem equally stunning. These limitations, however, are part and parcel of a neural information-processing system that is remarkably well adapted to cope with the demands of living in a constantly changing and ever more complex world.

Further Reading

Anderson JR (2000) *Learning and Memory: An Integrated Approach*. New York, NY: John Wiley.

Baddeley AD (1999) *Essentials of Human Memory*. Hove, UK: Psychology Press.

Bjork EL and Bjork RA (eds) (1996) *Handbook of Perception and Cognition*, vol. 10: *Memory*. New York, NY: Academic Press.

Bjork RA and Bjork EL (1992) A new theory of disuse and an old theory of stimulus fluctuation. In: Healy AF, Kosslyn SM and Shiffrin RM (eds) *From Learning Processes to Cognitive Processes: Essays in Honor of William K. Estes*, vol. 2, pp. 35–67. Hillsdale, NJ: Lawrence Erlbaum.

Bjork RA and Richardson-Klavehn A (1989) On the puzzling relationship between environmental context and human memory. In: Izawa C (ed) *Current Issues in Cognitive Processes: The Tulane Flowerree Symposium on Cognition*, pp. 313–344. Hillsdale, NJ: Lawrence Erlbaum.

Buckner RL and Wheeler ME (2001) The cognitive neuroscience of remembering. *Nature Reviews: Neuroscience* **2**: 624–634.

De Groot AD, Gobet F and Jongman RW (1996) *Perception and Memory in Chess: Studies in the Heuristics of the Professional Eye*. Assen, Netherlands: Van Gorcum.

Draaisma D (2000) *Metaphors of Memory: A History of Ideas About the Mind*. New York, NY: Cambridge University Press.

Eichenbaum H and Cohen NJ (2001) *From Conditioning to Conscious Recollection: Memory Systems of the Brain*. New York, NY: Oxford University Press.

Foster JK and Jelicic M (eds) (1999) *Memory: Systems, Process, or Function?* Oxford, UK: Oxford University Press.

Higbee KL (2001) *Your Memory: How It Works and How to Improve It*. New York, NY: Marlowe.

Naveh-Benjamin M, Moscovitch M and Roediger HL III (eds) (2001) *Perspectives on Human Memory and Cognitive Aging: Essays in Honor of Fergus Craik*. New York, NY: Psychology Press.

Parkin AJ (1999) *Memory and Amnesia*. Hove, UK: Psychology Press.

Parkin AJ (2000) *Memory: A Guide for Professionals*. New York, NY: John Wiley.

Richardson-Klavehn A and Bjork RA (1988) Measures of memory. *Annual Review of Psychology* **39**: 475–543.

Schacter DL (1996) *Searching for Memory: The Brain, the Mind, and the Past*. New York, NY: Basic Books.

Schacter DL (2001) *The Seven Sins of Memory: How the Mind Forgets and Remembers*. Boston, MA: Houghton Mifflin.
Schacter DL (2001) *Forgotten Ideas, Neglected Pioneers: Richard Semon and the Story of Memory*. Philadelphia, PA: Psychology Press.
Squire LR (ed.) (1992) *Encyclopedia of Learning and Memory*. New York, NY: Macmillan.
Squire LR and Kandel ER (1999) *Memory: From Mind to Molecules*. New York, NY: Scientific American Books.
Tulving E and Craik FIM (eds) (2000) *The Oxford Handbook of Memory*. New York, NY: Oxford University Press.
Whittlesea BWA (1997) Production, evaluation, and preservation of experiences: constructive processing in remembering and performance tasks. In: Medin DL (ed.) *The Psychology of Learning and Motivation*, vol. 37, pp. 211–264. New York, NY: Academic Press.
Wilding J and Valentine E (1997) *Superior Memory*. Hove, UK: Psychology Press.

Memory, Neural Basis of: Cellular and Molecular Mechanisms

Introductory article

Mark R Rosenzweig, University of California, Berkeley, California, USA

CONTENTS

Through the ages, there has been much speculation about how memory works. The current consensus is that memories are stored in circuits of neurons by plastic neurochemical and neuroanatomical changes at junctions between neurons (synapses), but it remains to be determined which of these changes are necessary and sufficient for memory formation.

INTRODUCTION

Recorded speculations about mechanisms of memory go back over two thousand years, but only in the last fifty years has substantial empirical progress been made in solving this mystery. In Greek mythology, memory was the province of the goddess Mnemosyne. She was the mother of the nine Muses, goddesses who presided over learning and the arts and sciences. This relationship demonstrated the necessity of memory for creativity. We invoke the name of Mnemosyne whenever we call methods to aid memory 'mnemonics' or 'mnemonic devices'.

Speculations about the bodily mechanisms of memory have been related to the technology of each period of history. Thus, classical Greek and Latin authors used as their models or metaphors of memory processes the then-current technology of wax slates and of signet rings impressing wax seals. Socrates assumed that there is a block of wax in our souls, the gift of Mnemosyne. He suggested that the wax varied in quality in different individuals, with finer wax allowing sharper, more detailed impressions. More recent models of memory mechanisms have ranged from telephone exchanges, to computers, to storage of genetic information.

In the Renaissance, when water was used to activate mechanical devices, nerves were thought to be tubes that conducted a fine fluid called 'animal (or animate) spirits'. By the mid-nineteenth century, nerve cells were visualized with the aid of microscopes and dyes, and in the latter part of the century, psychologists had begun to speculate that training could cause the proliferation of contacts between neurons. Such speculations appeared even before the formal announcement of the 'neuron doctrine' – that is, that neurons are separate cells that can affect other neurons but do not interpenetrate them. Some likened the nervous system to a telegraph system, but after telephones entered into commercial use in the 1880s, others likened the nervous system to a telephone system where connections can be made or broken. Synaptic junctions between neurons were named only near the end of the nineteenth century. Not until the middle of the twentieth century was it accepted that transmission at most synaptic junctions is accomplished