S. Groes (Ed.), Memory in the Twenty-First Century: New Critical Perspectives from the Arts, Humanities, and Sciences, (pp251-262). Hampshire, England. Palmgrave MacMillan 2016

28

Remembering

Larry R. Squire and John T. Wixted

Memory is a large topic, growing out of the fundamental fact that the experiences we have can modify the nervous system such that our mental life and our behaviour can be different than they were in the past. The study of memory ranges widely – from cellular and molecular questions about the nature of synaptic change to questions about what memory is: whether it is one thing or many, which brain systems support memory, and how those systems operate. We will consider in particular the structure and organization of memory with a focus on brain systems.

The idea that functions of the nervous system can be localized was well accepted by the end of the nineteenth century. Yet these ideas concerned mainly sensory-motor functions and language and did not speak to the topic of memory itself. In the early twentieth century, an influential programme of research in the rat concluded that memory is not localized but is distributed through the neocortex (the outer layer of the cerebral hemispheres of the brain of mammals involved in higher functions such as sensory perception, attention, memory, and action), such that each region contributes equivalently to the whole (Lashley 1929). Memory was thought to be distributed and well integrated with intellectual and perceptual functions, and no particular brain region was thought to be dedicated to memory function.

All of this changed in the 1950s when profound effects on memory were reported following a bilateral medial temporal lobe resection (the removal of the inner structures of the temporal lobe) carried out in the patient known as H. M. (Scoville and Milner 1957). This experimental surgery successfully relieved H. M.'s severe epilepsy, as was intended, but it also resulted in severe and debilitating forgetfulness, which occurred against a background of apparently intact intellectual and perceptual functions. For example, the patient could copy a complex drawing as well as controls, suggesting that his ability to perceive visual information was intact; and he could continuously rehearse (and then repeat back) a string of five or six digits as well as controls, suggesting that his 'working memory' was also intact. But when his attention was diverted, he soon forgot the drawing and the digits. Early descriptions of H. M. can be said to have inaugurated the modern era of memory research and strongly influenced the direction of subsequent work.

Most significantly, this work identified for the first time a particular area of the brain as important for memory.

H. M.'s bilateral lesion included the hippocampus, amygdala, and the adjacent parahippocampal gyrus. The immediate question was of which structures within this large surgical removal were responsible for his circumscribed memory impairment; that is, which structures and connections within the human temporal lobe have dedicated memory functions? These matters became understood gradually during the 1980s following the successful development of an animal model of human amnesia in the nonhuman primate (Mishkin 1978). The important structures proved to be the hippocampus and the adjacent entorhinal, perirhinal and parahippocampal cortices, which make up much of the parahippocampal gyrus (Figure 28.1). (Anatomically related structures in the thalamus and hypothalamus in the diencephalic midline, an area not part of H.M.'s lesion, are also important for memory, but these will not be discussed.) Damage limited to the hippocampus itself causes moderately severe memory impairment, but the impairment is greatly exacerbated when the damage extends to and includes the parahippocampal gyrus (as was the case with H.M.) (Zola-Morgan, Squire and Amaral 1986; Rempel-Clower et al. 1996). In all cases, the disorder is characterized most prominently by an impaired ability to form new memories (anterograde amnesia), but also by difficulty in accessing some memories acquired before the onset of the impairment (retrograde amnesia). Memories acquired shortly before the occurrence of a brain lesion (such as during the previous year) tend to be more impaired than memories acquired in the distant past. Thus, the structures that compose the medial temporal lobe memory system are essential for the initial formation of enduring long-term memories as well as for their maintenance and retrieval for a time after learning. The fact that very remote memory tends to be preserved after medial temporal lobe damage indicates that these structures are not the ultimate repository of long-term memory.

Once the important structures of the medial temporal lobe had been identified, the question naturally arose of whether the different structures have specialized roles. An early view held that the hippocampus plays an especially important role in spatial memory (O'Keefe and Nadel 1978). This idea was based on the common finding that rodents with selective hippocampal lesions are severely impaired in spatial learning tasks, such as learning to navigate a maze. However, subsequent work involving humans and monkeys with selective hippocampal lesions demonstrated pronounced spatial and nonspatial memory impairment. For example, patients with hippocampal lesions were impaired in their ability to recognize words that had appeared in an earlier list – a task with no obvious spatial component (Reed and Squire 1997). Findings like these suggest that the hippocampus plays a broader role in memory encoding and consolidation (the gradual process by which a temporary, labile memory is transformed into a more stable, longlasting form).

Another popular idea about specialization of function within the medial temporal lobe was based on a long-standing psychological distinction between familiarity and recollection (Atkinson and Juola 1974; Mandler 1980). Familiarity

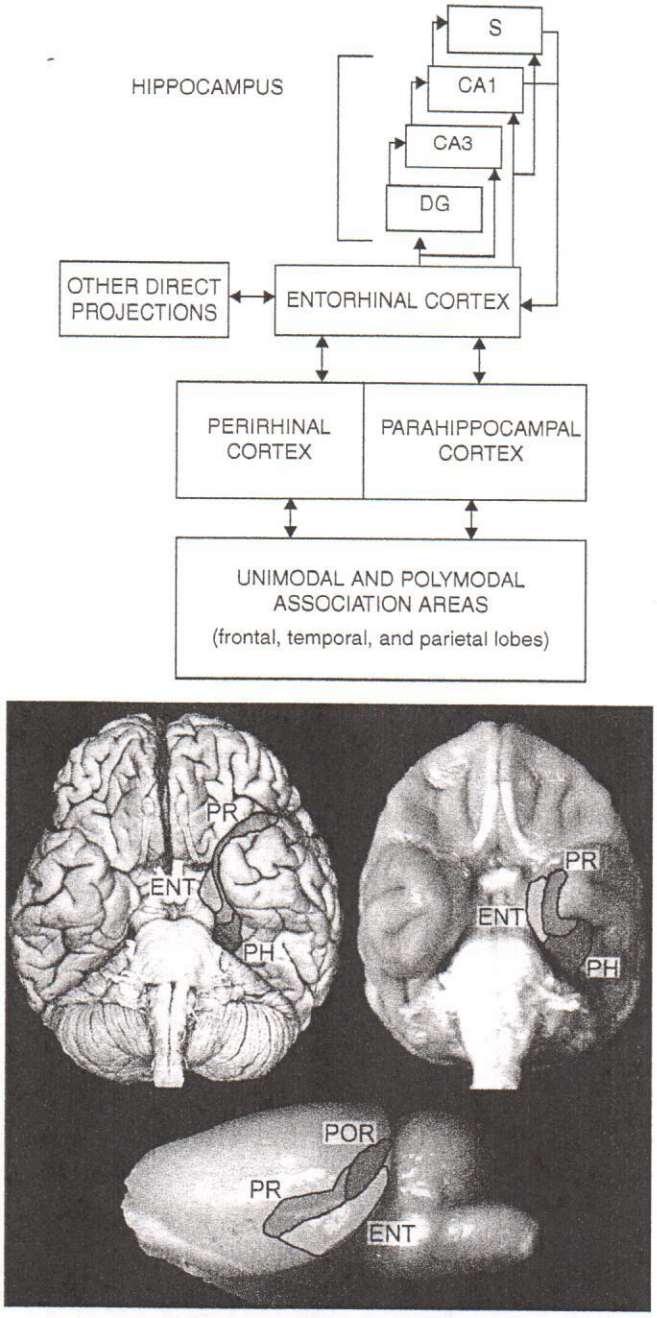

Figure 28.1 The medial temporal lobe memory system. Top: Schematic view of the memory system, which is composed of the hippocampus and the perirhinal, entorhinal and parahippocampal cortices. In addition to the connections shown here, there are also weak projections from the perirhinal and parahippocampal cortices to the CA1-subiculum border. Bottom: Ventral view of a human brain (upper left) and a monkey brain (upper right) and a lateral view of a rat brain (lower centre). The major cortical components of the medial temporal lobe are highlighted and outlined. The hippocampus is not visible from the surface and, in the human, lies beneath the structures of the medial temporal lobe. Its anterior extent lies below the posterior entorhinal and perirhinal cortices, and the main body of the hippocampus lies beneath the parahippocampal cortex. In the rat, the parahippocampal cortex is termed the post rhinal cortex. Abbreviations: dg, dentate gyrus; ent, entorhinal cortex; ph, parahippocampal cortex; por, post rhinal cortex; pr, perirhinal cortex; S, subicular complex

Source: Adapted from Figure 2 in Larry R. Squire and John T. Wixted, 'The Cognitive Neuroscience of Memory since H.M.', *Annual Review of Neuroscience* 34 (2011), 259–288.

involves knowing only that an item has been previously encountered (for example, when you recognize a face but cannot recall who the person is), and recollection involves recalling specific details about the prior encounter (such as recalling where and when you met the familiar person). Initially, a number of findings were interpreted to mean that hippocampal lesions selectively impair the recollection process but leave memory based on familiarity intact (Brown and Aggleton 2001). In addition, neuroimaging studies were often interpreted to mean that recollection-based decisions generate elevated activity in the hippocampus, whereas familiarity-based decisions generate elevated activity in other medial temporal lobe structures, particularly the perirhinal cortex (Eichenbaum et al. 2007). However, subsequent studies found that bilateral hippocampal lesions in humans have comparable effects on recollection and familiarity, and neuroimaging studies found that both familiarity-based and recollection-based recognition generate elevated hippocampal activity when both kinds of memory are strong (Song et al. 2011; Smith, Wixted and Squire 2011; Song, Jeneson and Squire 2011). Thus, the specialization of function within the medial temporal lobe does not seem to be informed by this distinction.

Because the functions of the different medial temporal lobe structures do not apparently divide up along the lines of spatial versus nonspatial memory or recollection versus familiarity, we must look elsewhere to identify functional differences between the structures. An important consideration is the fact that the inputs to each structure are quite different (Squire 1986; Suzuki and Amaral 1994). For example, the perirhinal cortex receives the majority of its cortical input from areas supporting visual object perception. Thus, the perirhinal cortex may be particularly important for forming memories of visual objects. Similarly, the parahippocampal cortex receives significant input from areas supporting spatial processing (for example, the ability to perceive that objects A and B are closer together than objects C and D). This area may therefore be particularly important for forming memories about the spatial locations of objects. A growing body of evidence is consistent with these ideas (Buffalo et al. 2006; Staresina et al. 2011; Liang et al. 2013; Staresina et al. 2013). That is, the functional specialization of different medial temporal lobe structures is sensibly related to the domain of information they process – information that is carried to these structures from upstream regions supporting different kinds of perceptual processing (Wixted and Squire 2011).

Within the medial temporal lobe, the hippocampus is the ultimate recipient of convergent projections from the entorhinal, perirhinal and parahippocampal cortices. Thus, the hippocampus itself is in a position to play a role in the encoding and consolidation of all aspects of an experience (its visual, spatial, auditory and olfactory qualities, as well as other contextual information). These anatomical facts can therefore explain why damage to the hippocampus results in broad memory impairment that covers all modalities and extends across multiple domains. Current studies are using new genetic methods in mice and other techniques to analyse the separate contributions of specific connections and cell types within the hippocampus (Yassa and Stark 2011; Xu Lieu et al. 2012).

The memory impairment associated with medial temporal lobe lesions is narrower than once thought, because not all forms of learning and memory are affected. The first clue came in 1962 when H. M. was found capable of acquiring a motor skill (mirror drawing) over a period of three days, though he could not recall these periods of practice. While this finding showed that memory is not unitary, discussions at the time tended to set aside motor skills as a special case representing a less cognitive form of memory. The suggestion was that the rest of memory is of one piece and is dependent on medial temporal lobe structures.

Yet during the subsequent years, it was discovered that motor-skill learning is but one example of a large domain of abilities that are independent of the medial temporal lobe. An early discovery was that perceptual and cognitive skills – not just motor skills – are intact in patients like H. M. Thus, memory-impaired patients acquired at a normal rate the skill of reading mirror-reversed words, despite poor memory for the words themselves (Cohen and Squire 1980). This finding led to the proposal of a brain-based distinction between declarative and procedural knowledge. Declarative knowledge referred to knowledge available as conscious recollections about facts and events. Procedural knowledge referred to skill-based information: knowledge expressed through performance rather than recollection.

Soon after this discovery was made, the phenomenon of priming was also found to be spared in amnesia (Tulving and Schacter 1990; Warrington and McCarthy 1987; Schacter and Buckner 1998). Priming refers to an improved ability to detect or identify stimuli based on a recent encounter with the same or related stimuli. For example, memory-impaired patients could (like healthy volunteers) name recently presented object drawings 100 milliseconds faster than new drawings, despite having poor memory for the drawings themselves (Backer Cave and Squire 1992). Perhaps the most compelling evidence for the independence of priming and ordinary memory ability was that severely amnesic patients can exhibit fully intact priming for words while performing only at chance levels on conventional recognition memory tests for the same words (Hamann and Squire 1997).

Another important insight was the idea that the neostriatum (a subcortical region of the brain that includes the caudate nucleus and putamen), and not the medial temporal lobe, is important for the sort of gradual, feedback-guided learning that results in habit memory (Mishkin et al. 1984). For example, memory-impaired patients learned tasks at a normal rate when the outcome of each learning trial was determined probabilistically, and performance therefore needed to be based on a gut feeling rather than on conscious memory of past events (Knowlton et al. 1996). Work with experimental animals was also the source of new insights, including the discovery in the early 1980s that the cerebellum is essential for delay eyeblink conditioning,[16] a kind of learning entirely preserved after hippocampal lesions (Clark and Squire 2000; Christian and Thompson 2003). Still other types of learning, which involve attaching a positive or negative valence to a stimulus (as in fear conditioning), depend on the amygdala (Ledoux 1996). Given the variety of tasks explored in these

studies and the number of brain structures implicated, an account of memory based on a two-part dichotomy (declarative versus procedural) began to seem too simplistic. Accordingly, the perspective eventually shifted to a framework that accommodated more than two memory systems. At that time, the umbrella term 'non-declarative memory' was introduced with the intention of distinguishing between declarative memory (which refers to one memory system) and other types of memory (in which several additional systems are involved) (Squire and Zola-Morgan 1988). Figure 28.2 illustrates this idea.[17] Declarative memory is what the term 'memory' signifies when we use it in everyday language. The stored representations are flexible and thought to be accessible to conscious awareness. Declarative memory is representational; it provides a way to model the external world and is either true or false. In contrast, nondeclarative memory is neither true nor false: it is dispositional and occurs as modifications within specialized performance systems. Thus, the various memory systems can be distinguished in terms of the different kinds of information they process and the principles by which they operate. These systems work in parallel to support behaviour. For example, an aversive event in childhood (such as being knocked down by a large dog) can lead to an enduring declarative memory of the event itself (dependent on the hippocampus and related structures) as well as a long-lasting, nondeclarative

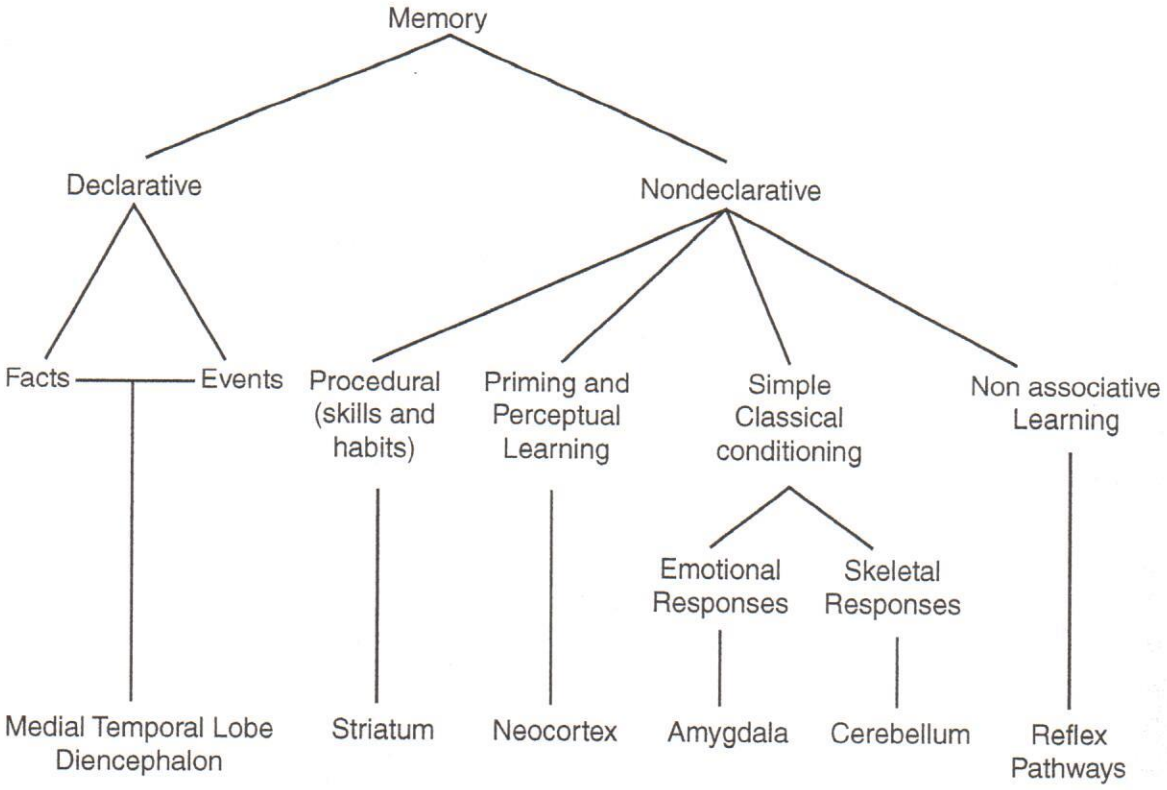

Figure 28.2 Organization of mammalian long-term memory systems. The figure lists the brain structures thought to be especially important for each form of declarative and nondeclarative memory. In addition to its central role in emotional learning, the amygdala is able to modulate the strength of both declarative and nondeclarative memory

fear of dogs (a phobia, dependent on the amygdala) that is experienced as part of the personality rather than as a memory.

The hippocampus and related structures in the medial temporal lobe have a time-limited role in the formation and storage of memory. Two lines of work underlie this idea. First, damage to these structures typically spares remote memory and impairs more recent memory in a temporally graded fashion. In humans, hippocampal lesions affect memory for up to a few years after learning. In experimental animals (usually rats or mice), similar damage impairs memory for up to 30 days after learning (Squire and Bayley 2007). Thus, long-term, stable memory develops more slowly in humans than in experimental animals. Discussion in the field continues about the possible special status of spatial memory and autobiographical memory in humans and the idea that these forms of memory might depend on medial temporal lobe structures as long as memory persists (Moscovitch et al. 2006). Yet there are reports of patients with medial temporal lobe lesions in whom remote spatial and autobiographical memory has been spared (Squire and Bayley 2007). The second line of work involves studies of experimental animals that track neural activity or structural changes in the hippocampus and neocortex after learning. For example, expression patterns of activity related genes like c-Fos describe gradually decreasing activity in the hippocampus after learning and parallel increases in activity in a number of cortical regions (Frankland and Bontempi 2005). These findings and others describe the increasing importance of distributed cortical regions for the representation of memory as time passes after learning (Restivo et al. 2009). Similar findings have been obtained in neuroimaging studies; for example, when volunteers attempt to recall news events that occurred anywhere from one to 30 years earlier (Smith and Squire 2009). The idea is not that memory is literally transferred from the hippocampus to the neocortex. Memory is always in the neocortex, but gradual changes occur to increase the complexity, distribution and connectivity of memory representations among multiple cortical regions. At the same time the role of the hippocampus gradually diminishes.

One way to view this process is to suppose that a time-and-place-specific new memory (a so-called episodic memory) is represented initially by an ensemble of distributed changes in the neocortex and by changes in the hippocampus (and anatomically related structures) as well. The neocortical ensemble is viable so long as the episode is maintained within active memory. However, when one's attention is directed elsewhere, a problem arises. How can the unique distribution of sites that represent this new memory be revivified by unaided recall or after the presentation of a partial reminder? The notion is that remembering becomes possible because medial temporal lobe structures, by way of their widespread, divergent connections to the neocortex, effectively bind together the distributed neocortical sites that together constitute the new memory. This connectivity supports the capacity for remembering during the consolidation process until the connectivity among the relevant cortical sites becomes strong enough to represent a stable memory without the support of the medial temporal lobe.

A long-standing idea, which has received renewed attention in recent years, is that retrieval of memory provides an opportunity for updating or modulating what was originally learned and even the possibility of severely disrupting it (Nader et al. 2000; Loftus 2005; Lee 2009; Dudai 2012; St. Jacques et al. 2013). The process by which a long-term memory transiently returns to a labile state (and then re-stabilizes) has been termed reconsolidation. Although it is clear that memory can be modified or distorted by memory retrieval, questions remain about the conditions under which memory can actually be abolished. Some studies in experimental animals report that a reactivated memory can be impaired but that the disruption is transient (Lattal and Abel 2006). Other studies in animals report that only recent memories (ones that are one or seven days old, but not 14 or 28 days old) can be impaired after reactivation (Milekic and Alberini 2002).

Consolidation presumably requires some relatively long-lasting form of communication between the medial temporal lobe and the neocortex. One proposal for how this could be accomplished is through the phenomenon of neural replay. Recordings of neural activity in rodents showed that firing sequences of hippocampal neurons during waking behaviour are then spontaneously replayed during subsequent slow-wave sleep (Wilson and McNaughton 1994). Later it was found that hippocampal replay was coordinated with firing patterns in the visual cortex, which is consistent with the idea that a dialogue occurs between hippocampus and neo-cortex (Ji and Wilson 2007). This coordination could be part of the process by which recent memories eventually become consolidated remote memories. Interestingly, disrupting replay activity in rodents during a rest period (filled by quiet wakefulness and slow-wave sleep) following spatial learning impairs later memory for the task (Ego-Stengel and Wilson 2010).

These studies with rodents led to conceptually similar studies with humans. For example, volunteers memorized the locations of card pairs on a computer screen while being exposed to a particular odour (the smell of a rose). Later, odour re-exposure, specifically during slow-wave sleep, increased hippocampal activity (measured by neuroimaging) and lessened forgetting of the card pair locations following sleep (Rasch et al. 2007). In another study, the hippocampus and para-hippocampal gyrus were active while participants learned routes in a virtual reality environment and were active again during subsequent slow-wave sleep (Peigneux et al. 2004). The degree of activation during slow-wave sleep correlated with memory performance the next day. Studies like these have been interpreted to mean that consolidation results from the reactivation of newly encoded hippocampal representations, specifically during slow-wave sleep (Inostroza and Born 2013).

An important question is whether neural replay and the consolidation process are specific to slow-wave sleep or whether these events might occur whenever the brain is not actively encoding new memories, such as during quiet wakefulness (Mednick et al. 2011). In rodents, neural replay can occur during wakefulness (Karlsson and Frank 2009). Moreover, in a neuroimaging study with humans, coordinated hippocampal-cortical activity occurred during a rest period that

followed learning, and this activity predicted later memory performance (Tambini et al. 2010). Accordingly, an intriguing possibility is that the neural replay activity proposed to underlie memory consolidation may occur whenever the brain is in a quiet state (not just during slow-wave sleep).

Where are memories ultimately stored in the brain? A variety of evidence has converged on the view that the different aspects of remembered information are stored in the same regions of the brain that initially perform the processing and analysis of that information. According to this view, remembering a previous experience consists of the coordinated reactivation of the distributed neocortical regions that were activated during initial perceptual processing (Renzi 1982; Mishkin 1982; Squire 1987; and Damasio 1989). While the memory is still new, this reactivation of distributed cortical activity depends on the hippocampus and other medial temporal lobe structures, but once memory is fully consolidated, reactivation can occur within the neocortex itself. Each neocortical region operates within a specific domain and stores only the features of an experience – such as visual, auditory or spatial information – that belong to that domain. Thus, as proposed by psychologist Karl Lashley long ago, memories are distributed throughout the neocortex (Lashley 1929). However, contrary to his view, memory is not uniformly distributed. Some areas are more important for storing the visual aspects of an experience, and other areas are more important for storing other aspects.

An implication of this view is that neocortical lesions that selectively impair perceptual processing in a particular domain (such as the perceptual processing of colour) should also cause correspondingly specific anterograde and retrograde memory impairment within the same domain. This circumstance is illustrated by 'The Case of the Colorblind Painter', a case described by the neurologist Oliver Sacks (1995). An accomplished painter was involved in an automobile accident at the age of 65, which rendered him colour-blind. The disability was striking: he could discriminate between wavelengths of light, even though the different wavelengths gave rise to the perception of various shades of grey rather than the perception of different colours. Because his condition was acquired (it was not congenital), it was possible to interrogate not only his ability to form new colour memories, but also the status of previously established memories that had once included the subjective experience of colour. The case description leaves little doubt that the patient's experience – both going forward and looking back – was now completely (and selectively) devoid of colour. Although he retained abstract semantic knowledge of colour, he could neither perceive nor later remember the colour of objects presented to him (anterograde impairment). In addition, he could not subjectively experience colour in his earlier (and once chromatic) memories. For example, he knew that his lawn was green, but he reported that he could no longer visualize it in green when he tried to remember what it once looked like.

Note the difference between the effect of this cortical lesion on memory and the effect of bilateral medial temporal lobe lesions. With respect to remote memories that have already been fully consolidated, medial temporal lobe lesions have

little effect. In contrast, focal cortical lesions can selectively abolish one feature (like colour) of a long-consolidated memory. With respect to new experiences, bilateral medial temporal lesions lead to severe anterograde amnesia (no subsequent memory for a recent experience). In contrast, focal cortical lesions of the kind suffered by the painter prevent the encoding and retrieval of only one aspect of the experience (colour in his case). Because the processing of colour in the painter's neocortex was impaired, his experience of colour was eliminated in both perception and memory.

Selective deficits in long-term knowledge of the kind suffered by the painter are not limited to perceptual experience. Semantic knowledge (knowledge about objects, facts and word meanings) is also stored in neocortical regions that can be selectively damaged (Warrington and McCarthy 1987). Thus, damage limited to lateral regions of patients' temporal lobe (close to, but not including, medial temporal lobe structures) can disrupt previously stored information – such as what an animal looks or sounds like. Such patients have difficulty naming pictures of animals and providing information about them. Other patients with damage to the parietal cortex can have difficulty identifying small manipulable objects (like spoons and brushes) and knowing how to use them. Neuroimaging studies support the findings from lesion studies and show that the properties of objects, together with how they are perceived and used, influence which brain areas store long-term knowledge about their identity (Martin 2007).

The information in the preceding sections helps illuminate some of the memory deficits associated with normal ageing and dementia. One of the most common experiences associated with normal ageing is the decline in memory function. Often, the memory difficulty is characterized as poor 'short-term' memory. In its common usage, a short-term memory problem means having trouble remembering recent experiences (such as when someone tells a story for the second time without remembering having told it before) while at the same time having no trouble remembering events from decades ago. Older adults who exhibit these symptoms are having difficulty encoding and consolidating new memories, while memories that were acquired and consolidated long ago are easy to retrieve. These changes in memory ability are related to changes within medial temporal lobe structures. In experimental animals, the dentate gyrus within the hippocampus is most sensitive to the effects of ageing (Small et al. 2004). Studies in humans have reported between one and two percent annual hippocampal atrophy in non-demented adults older than 55 years (Jack et al. 1998). Aerobic exercise can reverse age-related volume loss by one to two years (Erickson et al. 2011).

Alzheimer's disease, the most common form of dementia, is a progressive neurodegenerative condition. It is a distinct condition, not an acceleration of the normal ageing process. The first targets of the disease are the entorhinal cortex and the CA1 field of the hippocampus, which explains why memory is especially affected in its early stages (Hyman et al. 1984; West et al. 1994). The rate of hippocampal volume loss is at least 2.5 times greater in Alzheimer's disease than in normal ageing (Jack et al. 1998). The disease progresses to involve intellectual functions quite broadly. The neocortex becomes involved (though sensory and

motor areas are relatively spared) and patients develop difficulty with language, problem solving, calculation and judgment.

Semantic dementia, another progressive disorder, begins elsewhere in the brain and is associated with a different pattern of symptoms (Hodges and Graham 2001). This condition prominently involves atrophy of the anterior and lateral temporal lobes (Levy et al. 2004; Patterson et al. 2007). Unlike patients with Alzheimer's disease, these patients have severe loss of previously stored and long consolidated semantic knowledge (that is, loss of conceptual knowledge about objects, facts and word meanings). Yet their ability to form new memories can be relatively spared. Thus, patients could recognize which drawings of animals they had seen recently but failed at tests of conceptual knowledge about the same items (Graham et al. 1997). Not just the name of the item is lost – the concept itself is degraded.

The understanding of memory has changed in ways that might have seemed revolutionary to Karl Lashley when he searched for sites of memory storage in the brains of rats (Lashley 1929). All that has been learned about the structure and organization of memory and about brain systems is the result of basic, fundamental research, mostly in rodents, monkeys and humans. Although we did not review it here, much has also been learned from studies of the cellular and molecular basis of memory, an enterprise that has depended heavily on mice as well as invertebrate animals like Aplysia and Drosophila. As this work continues, one can expect not only new insights into how memory operates but also improved understanding of human health and disease, including improved ways to diagnose, treat and prevent the diseases that affect memory.[18]

Notes

1. Viktor Mayer-Schönberger, *Delete: The Virtue of Forgetting in the Digital Age* (New Jersey: Princeton University Press, 2009), Loc. 62–70.
2. Fredric Jameson, *Postmodernism, or, the Cultural Logic of Late Capitalism* (New York and London: Verso, 1992), ix.
3. Gilles Deleuze and Felix Guattari, *A Thousand Plateaus: Capitalism and Schizophrenia*, trans. Brian Massumi (London and New York: Continuum, 2004), 15.
4. Marita Sturken, *Tangled Memories: The Vietnam War, the AIDS epidemic, and the Politics of Remembering* (Berkeley and Los Angeles, Oakland, CA: University of California Press, 1997), 16–17.
5. Wolfgang Ernst, *Digital Memory and the Archive*, ed. Jussi Parikka (Minneapolis and London: University of Minnesota Press, 2013), 195.
6. Mark Currie, *Postmodern Narrative Theory* (Houndmills: Palgrave, 1998), 97.
7. J. Bernlef, *Out of Mind*, trans. Adrienne Dixon (London: Faber, 1989), 127.
8. Evgeny Morozov, *To Save Everything, Click Here: Technology, Solutionism, and the Urge to Fix Problems That Don't Exist* (London: Penguin, 2013). For more examples of similar projects and technologies see Stacey Pitsillides' essay in this volume.
9. Douwe Draaisma, *Metaphors of Memory: A History of Ideas About the Mind*, trans. Paul Vincent (Cambridge: Cambridge University Press, 2000), 155.
10. Michael J. Sandel, 'The Case Against Imperfection', in *Human Enhancement*, ed. Nick Bostrom and Julian Savulescu (Oxford: Oxford University Press, 2009), 74–75.
11. Draaisma, *Metaphors of Memory*, 155.

12. Gerald Edelman, *Bright Air, Brilliant Fire* (Basic: New York, 1992), 206.
13. Catherine Malabou, *Ontology of the Accident: An Essay on Destructive Plasticity*, trans. Carolyn Shread (Cambridge: Polity, 2012), 18.
14. Malabou, *Ontology*, 18–19.
15. Catherine Malabou, *What Should We Do with Our Brain?*, trans. Sebastian Rand (New York: Fordham University Press, 2008), 9.
16. Delay eyeblink conditioning is a form of Pavlovian conditioning in which a conditioned stimulus (such as a tone) is presented and remains on until the unconditioned stimulus (such as a puff of air to the eye) is presented. The two stimuli overlap and co-terminate. D.A. McCormick et al., 'Initial Localization of the Memory Trace for a Basic Form of Learning', *Proceedings of the National Academy of Sciences* 79 (1982), 2731–2735.
17. For earlier versions of this diagram, see Squire, 'Mechanisms of Memory'.
18. This work appeared originally in the 2015 Winter Issue of *Daedalus*, the journal of the American Academy of Arts and Sciences.