Learning and Memory

85

Richard Morris

Contents

R. Morris (✉)
Centre for Cognitive and Neural Systems, Neuroscience, The University of Edinburgh, Edinburgh, UK
e-mail: r.g.m.Morris@ed.ac.uk

D.W. Pfaff, N.D. Volkow (eds.), *Neuroscience in the 21st Century*,
DOI 10.1007/978-1-4939-3474-4_81

Abstract

Memory is fundamental to human life. Qualitatively distinct types of memory enable us to change behavior in response to experience, to acquire and use a repository of knowledge, to recollect events from the past, and to plan for the future. In many respects, memory defines human individuality, as the memories of one person are necessarily different from those of another. Where they overlap, as in the shared memories of a community or a nation, they form a cultural memory that is often ritualized into various art forms. The use of memory is changing, with a great deal of human knowledge now externalized and then sought on demand through the use of search engines on the internet. Nonetheless, the loss of memory remains greatly feared, perhaps because it is recognized that loss of private episodic memories would undermine the sense of self. Moreover, the inability to recollect events and episodes can develop from a minor irritation to a condition that undermines normal daily existence.

Keywords

Cellular consolidation • Delayed matching-to-place (DMP) • Encoding • Learning and memory. *See* Memory • Loss of memory • Memory • *N*-methyl-D-aspartate (NMDA) receptor • Retrieval, memory • Reversal learning • Semantic memory • Spatial memory • Spatial reference memory • Synaptic plasticity • Systems consolidation • Watermaze • Wisconsin General Testing Apparatus (WGTA)

Brief History

Memory is fundamental to human life. Qualitatively distinct types of memory enable us to change behavior in response to experience, to acquire and use a repository of knowledge, to recollect events from the past, and to plan for the future. In many respects, memory defines human individuality, as the memories of one person are necessarily different from those of another. Where they overlap, as in the shared memories of a community or a nation, they form a cultural memory that is often ritualized into various art forms. The use of memory is changing, with a great deal of human knowledge now externalized and then sought on demand through the use of search engines on the internet. Nonetheless, the loss of memory remains greatly feared, perhaps because it is recognized that loss of private episodic memories would undermine the sense of self. Moreover, the inability to recollect events and episodes can develop from a minor irritation to a condition that undermines normal daily existence.

Given the central role of learning and memory in cognition, a "Grand Challenge" for neuroscientists across the world is to understand the neural mechanisms of the capacity to encode, store, consolidate, and retrieve information. That understanding will, in time, influence new thinking about effective teaching in schools and new treatments for those afflicted by memory disorders. Over recent years, there has been an explosion of research that is gradually revealing the underlying psychological processes and neural mechanisms involved. Such research ranges from behavioral

and noninvasive imaging studies through to detailed investigations of the molecular mechanisms mediating specific intracellular signal transduction cascades. It builds upon the foundations in anatomy, physiology, and experimental psychology of the last century. Much has been learned, but there remain a number of unsolved issues.

This chapter addresses issues that are vital for an introductory understanding of key concepts in the neurobiology of learning and memory – with a focus on studies in animals. The reason for this focus is because animal studies enable invasive studies of the brain that are vital for a causal understanding of mechanism. The most important animal studies have, however, been inspired by an understanding of human memory and its disorders – and this is therefore the starting point. Starting from a discussion of amnesia, section "Multiple Types and Processes of Memory" considers both the concept of distinct *types* of human memory subserving different cognitive functions and the orthogonal issue that each of these types involves the distinct *processes* of encoding, storage, consolidation, and retrieval. It touches on the need for new techniques to address outstanding issues and new questions. Section "Making Declarative Memories" focuses on *declarative memory* and the distinction between semantic, episodic, and spatial memory. It includes details of the watermaze as one effective and relatively inexpensive method of studying spatial memory. Section "Physiological, Pharmacological, and Molecular Engineering Approaches to the Study of Memory" moves beyond monitoring behavior to consider *synaptic mechanisms of plasticity* and, specifically, outlines a simple *computational way of storing information* based on no more than altering the strength of a synapse within a distributed associative network. Section "The Structural and Functional Expression of Memory Traces" presents some developing ideas about likely *molecular mechanisms* of memory consolidation. The chapter is an introduction rather than a comprehensive survey, but will hopefully provide a start for the interested student – and enough to engage them in thinking rationally and creatively about the topic from their own individual vantage point.

Multiple Types and Processes of Memory

The Loss of Memory in Neurological Disorders Has Revealed Important Dissociations Between Different "Types" of Memory

It is a feature of human nature that the importance of something is often really recognized through its loss – as in the loss of a loved one. With respect to memory, the fact that it works is generally taken for granted, even if it is a faculty that is complained about because of its inherent fallibility. However, consider the possibility that because of a neurological condition someone could only live in the present with little appreciation of the past or sense of the future. Such a person would still be able to see, hear and smell, move around, and be capable of language comprehension and speech. They might still be able to learn but – critically – would lack the awareness of having done so. Living solely in the present would reduce their life to one with little personal meaning. Sadly, this has happened to some people –

notably the famous neurological patient Henry Mollaison (HM) first studied by Brenda Milner in Canada in the 1950s – and to other brain-damaged patients who have also been intensively examined for scientific as well as clinical reasons. A common feature, as shown in the structural magnetic resonance image (MRI) of Fig. 1a, is damage to the medial temporal lobe of the brain. In HM's case, this happened deliberately through bilateral surgical resection of the temporal lobe in a successful attempt to relieve intractable epilepsy – the original drawings after the surgery being shown in Fig. 1b. In other patients, such damage may occur naturally through viral encephalitis or ischemic stroke. These patients are then amnesic in the sense of being unable to form new memories although, interestingly, they show dissociations between different types of memory as they can still acquire skills and successfully remember some information from the past.

Numerous relatively simple tests of human memory have been devised, ranging from recognition of familiar pictures through to recall of lists of words or passages of prose. Figure 1 also presents data from HM in which he was shown a complex drawing (Fig. 1c) and asked to copy it. HM's direct copy is in the middle. It is quite good for he could draw well. But asked to draw it again 10 min later and he described having no memory of ever having seen it. Despite this apparently devastating loss of memory, amnesic patients such as HM can still learn, but they do so implicitly without developing any memory of the fact they can. One test was of "mirror drawing" in which HM's initial efforts were clumsy but he gradually learned to transpose right for left, up for down, and so on until he could successfully move a stylus in the alleyway created by a drawing of a star (Fig. 1d). Errors were defined by inappropriate crossings of the line and, while he started poorly, he gradually improved. But HM developed no personal knowledge of having learned. Asked 3 days later to draw in a mirror, he could remember neither the mirror nor his extensive practice of the preceding days. Astonishingly, he then proceeded to execute the skill very well.

Other amnesic patients also show this dissociation between "memory" and "acquired skill." For example, a musician called CW became densely amnesic following viral encephalitis, survived and is still alive, but he has little sense of the passage of time and no memory. But when examined after recovery from the acute effects of the virus, he could still conduct his old choir and play an unseen passage on the piano well. This fascinating dissociation seems to be akin to one that was characterized by the philosopher Gilbert Ryle as the difference between "knowing that" and "knowing how." In modern terminology, this is described as the difference between "declarative" and "non-declarative" memory – where declarative memories are those with truth value whereas non-declarative memory is merely executed (as in a skill). Other amnesic patients show additional dissociations, such as that of a patient called EP who displays a near complete inability to acquire new memories but retains access to very old ones. EP remembers the layout of the town where he used to live but cannot learn or remember where he now lives. These are well-known cases, but many others have been studied to reveal a pattern of dissociations. Many, though not all, show a dissociation between intact very short-term memory (over a

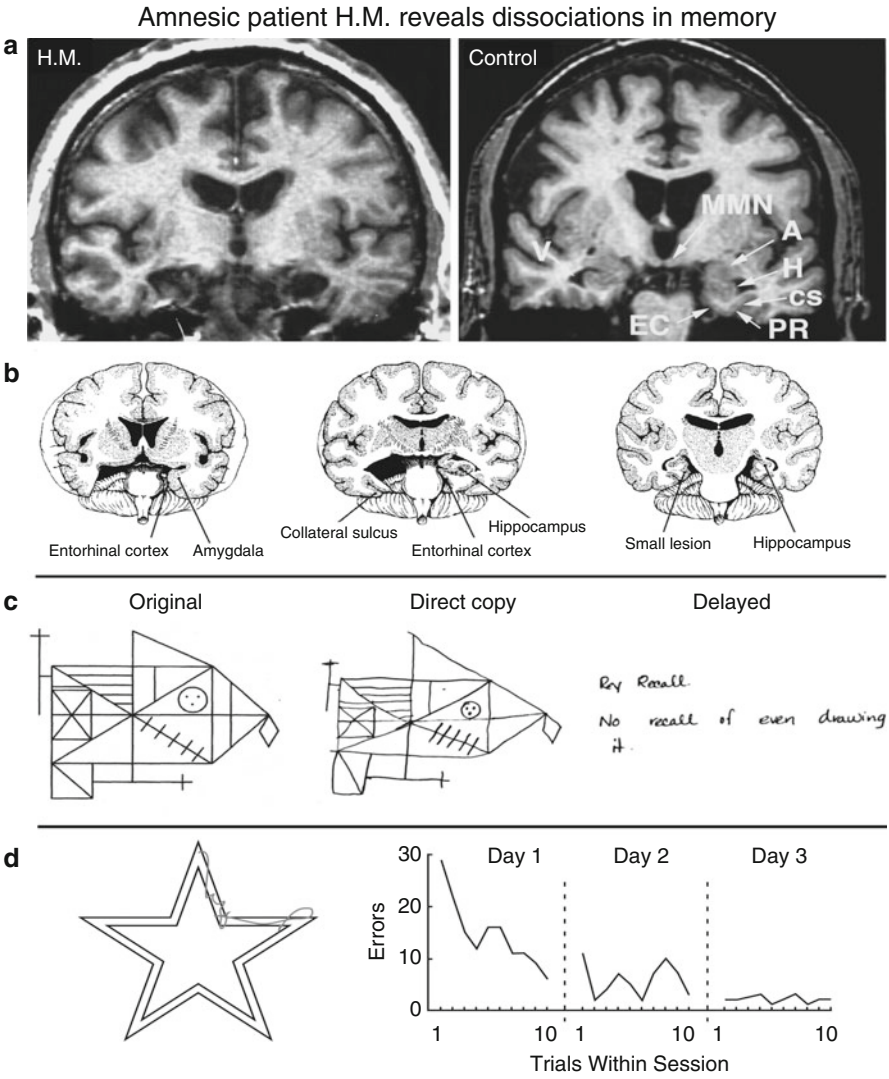

Fig. 1 *Amnesic patients reveal dissociations in memory.* (**a**) Structural MRI scan of the brain of the late patient Henry Mollaison (HM) showing missing amygdala, hippocampus, and adjacent cortical regions compared to an intact control. (**b**) Artist's drawings made at the time of the original surgery; the histology of the brain that has now been sectioned should provide a basis to confirm the damage depicted in these drawings. (**c**) HM's direct copy of the classic "Rey" diagram shows that he had no difficulties in perception or drawing, but he was unable to draw the picture from memory at all. (**d**) HM gradually learned the skill of mirror drawing over several days, much as would a control subject, but he failed to remember explicitly that he had mastered this skill (Modified from the work of Corkin S, MIT, Cambridge, MA, USA)

few minutes) and impaired longer-term memory. Others show subtle distinctions, such as between the successful ability to recognize something as familiar in the absence of any ability to recall any further information with which this thing or person may be associated. Single case and small group studies such as these continue to help us build up a picture of the "fracture lines" of human cognition in the domain of memory. Beyond a natural concern for the plight of these patients, much has been learned about the psychological and anatomical organization of memory systems. This is not only helpful with respect to developing a scientific understanding of the organization of memory, but can also help clinicians in cases of memory loss associated with neurodegenerative conditions such as Alzheimer's Disease where these dissociations emerge in a gradual, insidious manner as the disease progresses.

Distinct Brain Systems Mediate Different Types of Learning and Memory

Research with patients has been the starting point for identifying a number of different brain systems of memory with a "taxonomy" that is shown in Fig. 2. The taxonomic categories are based on other evidence as well – such as qualitative psychological distinctions between different types of memory, evolutionary considerations, the results of experimental lesion studies in both nonhuman primates and laboratory rodents and, more recently, the activation patterns in noninvasive brain-imaging studies. There are several points to note.

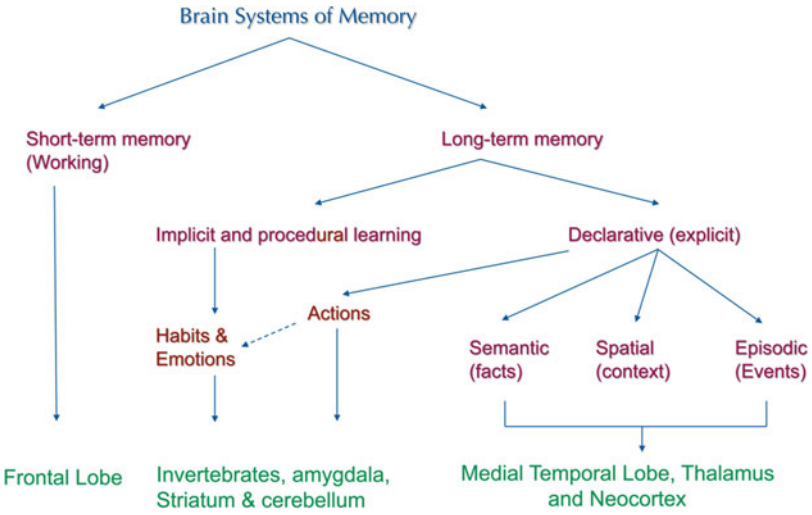

Fig. 2 *A taxonomy of brain systems of memory.* Memory systems are distinguished on the basis of various criteria – duration, type of information processed, nature of the information encoding rule, and brain regions/networks thought to be responsible for encoding or storage. Note ambiguous status of "actions" whose learning entails declarative components but whose expression can become automated into habits (Modified from the original proposal of Squire LR, UCSD, San Diego, USA)

First, there is the distinction between "short-term" or "working memory," on the one hand, and "long-term memory" on the other. Working-memory – the type you use when trying to remember a telephone number or to do mental arithmetic – is thought to depend on the sustained neural activity that is triggered by a particular stimulus and is then thought to "represent" that stimulus in the brain after the stimulus has gone away. Provided there is no interruption or distraction, this representation can be sustained for a period (seconds, possibly minutes, but no longer). Working memory consists, in humans, of several different subcomponents, including a system for holding onto spoken information (the "phonological store"), to visual information (a "visuo-spatial sketchpad"), and for coordinating diverse types of information about the recent past (the "central executive"). To illustrate an aspect of working-memory to yourself, think about what you would need to do if someone tells you to add up a number they have just said to you with one that they have written down on a piece of paper – you will combine these two bits of information using these three components of working-memory. Complementary studies in nonhuman primates, albeit without language, have revealed a population of neurons in the frontal lobe called "delay-activity-neurons" that have properties that would be desirable in a neural short-term memory system. They continue firing after the triggering stimulus has ended. Different neurons fire for different stimuli, such as locations in space on a computer screen or different colors, and even for distinct stimulus associations. As depicted in Fig. 2, working memory is generally held to be mediated by structures of the frontal lobe and be a component of a larger "executive system" there which mediates aspects of decision making and the planning of actions. However, there are now grounds for appreciating that this is an oversimplification as new functional imaging studies in humans reveal a more dispersed picture of the neuroanatomical networks mediating working-memory.

In contrast, long-term memory is thought to involve both an active encoding phase characterized by specific spatiotemporal patterns of neural activity and a long passive (or quiescent) stage when the "traces" of memory, most likely structural changes at synapses, are not actively influencing current neural firing patterns. Stored representations may then later be "retrieved" from long-term memory, as it happens if you are presented with a "retrieval cue" in the form of some reminding stimulus which then triggers a reactivation of appropriate neural activity. This neural firing interacts with the now changed synaptic weights in circuits and networks of multiple neurons that collectively re-create a representation in the brain of something that has happened in the past. Thus, short- and long-term memory differ *both* with respect to the duration for which they hold information *and* the physical manner in which they do so. It is considered that long-term memory traces are stored in an associative manner in regions associated with the relevant sensory processing of the information represented (visual, auditory, somatosensory, etc.), but the creation of such memories can involve other circuitry, notably the hippocampus. This is the area damaged in HM's brain. How an allocortical structure like the hippocampus interacts with diverse neocortical structures to store specific memory traces is still not well understood, though various theories of how this might work have been proposed in the context of memory consolidation.

Second, within the domain of long-term memory, there is a major distinction between "declarative" memories and "non-declarative" skills. The non-declarative component is a group of learning and memory systems that include perceptual and motor skills, emotional learning, and various phenomena generally called "priming." These non-declarative components were preserved in patient HM. One important aspect of learning a skill, such as mirror drawing, is that it may take many "trials" to reach perfection. Skill learning is often said to involve "trial-and-error" as the motor components gradually zero in on the desired movements. This gradual reduction of error involves reward. Emotional or "value" learning also involves an association between an initially neutral stimulus and a second event of biological significance – food, water, or something painful in the case of fear conditioning. In the latter case, this results in learned fear or anxiety that occurs prior to the aversive stimulus which follows, and this anticipatory emotion develops by a process called "classical conditioning" (and sometimes called Pavlovian Conditioning after its discoverer the Russian physiologist Ivan Pavlov who studied alimentary and aversive conditioning of salivary responses in dogs). The initially neutral stimulus now evokes fear/anxiety in the animal or human subject in which these pairings have taken place. Skill learning and emotional learning are mediated by structures of the striatum and amygdala, respectively.

One subtle but important aspect of skill learning that Fig. 2 attempts to capture is the evolving difference between the performance of an "action" and that of a "habit." Initially, the execution of a movement is intentional. An action is a movement performed in a declarative manner to achieve a particular outcome (such as a specific reward). That is, a person learns to rotate the wheel on a car appropriately *in order to* steer the car; and a rat learns to press a lever in an operant chamber *in order to* get food. The critical association that is encoded and will be stored once learning has occurred is between the specific action (A_j) and its specific outcome (O_k) – an A_j-O_k association where j, k are variables that can represent a variety of actions ($j = 1, 2, 3\ldots$) and outcomes ($k = 1, 2, 3\ldots$). The subject in which this learning takes place will then intentionally execute the appropriate movement to secure the desired outcome. However, actions can gradually become habits. Habits are quite different and not necessarily intentional. Over time, and repeated successful execution, an action may gradually become automatized and, as it does so, will necessarily and repeatedly occur in the presence of some particular context or stimulus (S_i). The stage is thus set for an entirely separate association to develop between an apparently precipitating stimulus and the action (S_i-A_j), such that the action will eventually be "run-off" without regard to the intended outcome. Once S_i occurs, the action A_j follows. At this point, the action is said to have become a "habit" and so become more of a "response" than an action, with the S_i-A_j being somewhat independent of the O_k outcome used to achieve conditioning. This is the infamous S-R learning of conditioning theory, first developed by the American psychologist Edward Thorndike at Columbia University, in which the execution of the response is sustained by the reward that follows, but this reward (e.g., O_k) may no longer be neurally represented in the association.

This apparently subtle difference between actions and habits is important. Stimulus–response habits have long been held to be "stamped-in" by the reward that follows their occurrence, and then sustained by it, but the response is literally a learned reflex in that it now occurs automatically given the precipitating stimulus, rather than intentionally. Habits are notoriously difficult to change – as is well known – including such dangerous habits as drug-seeking behavior. Although learned with a view to seeking pleasure (the "high" of an illicit drug), they become automatized in the face of the increasingly negative consequences of drug taking. Recent studies have revealed that the distinct subdivisions of the striatum are involved in this gradual transfer from actions to habits – with the dorsomedial striatum being involved in mediating actions and the dorsolateral striatum engaged in mediating the sensorimotor connections that underpin habits. During initial learning, the neurotransmitter dopamine, released at terminals in the striatum of projection neurons located elsewhere in the substantia nigra and the ventral tegmental area, serves as the error-correcting signal of mismatch between expected and actual reward. That is, dopamine neuron firing can be a neural representation of the O_k of action learning. However, once an action has become a habit, the learned response is much less sensitive to changes in the reward that follows its execution. Study of the modulation of action learning by catecholamines and the neural circuitry underlying actions and habits are both active areas of research. The tasks for studying these are reasonably straightforward and there is a real opportunity for small labs with limited equipment and infrastructure to make a valuable contribution.

The *third* point concerning different types of memory relates to declarative memory. Specifically, the taxonomy of Fig. 2 documents the existence of distinct forms of memory called *semantic memory*, *spatial memory*, and *episodic memory* that are each declarative because they have "truth value." You can declare knowledge – of what something is, of where it is, or of an event having happened – in the way that you cannot of a skill. A skill is expressed in movement rather than declared, a subtle distinction that deserves further explanation. Clearly, humans can declare that they *possess* a skill – such as being able to play a musical instrument – and that declaration can be true or false. However, that declaration is *not itself the skill* – that is the actual playing of the piano, guitar, or other instrument. Thus, whereas non-declarative learning is expressed in performance, declarative knowledge is labile and essentially private, with its existence expressed in language (in humans) or in some other aspect of behavior (humans and animals). In recognition memory, a person can tell us explicitly that he (or she) has seen a particular person or picture before, that he remembers an event from yesterday, that he knows the shape or color of a particular object that he has seen before, and so on. Similarly, an animal can acquire knowledge of what it can eat, of where it can find water, and of events/episodes that have happened to it. This knowledge is often associative in nature (such as where water is to be found), but differs from emotional or value learning in being declarative association that enables one stimulus to evoke the memory of another rather than change the "valence" of a neutral stimulus through association with reward. For this reason, paired-associate learning, as it is often called, is

considered to be different from classical Pavlovian conditioning. The acquisition of declarative knowledge is mediated by structures of the medial temporal lobe, such as the hippocampal formation and adjacent structures, though there are outstanding debates about whether the resulting memory traces are stored there or elsewhere in cortex.

The simplest example of declarative knowledge is event memory in which (a) something happens at a particular point in space and time, and (b) a person or animal later remembers that this happened. Humans can later recall such an event, as in "I saw a car-crash today," and then describe what happened, where and when. Some believe this type of episodic memory is uniquely human, but there are several grounds for doubting this exclusivity to humans as higher vertebrates (mammals and birds) do possess a rudimentary system for recalling events. For example, food-caching birds return without difficulty to the exact location where they stored food to retrieve it when they are hungry, indicating they have formed a food-location association. These types of "paired-association" are now being followed up in laboratory animals as well. However, that birds can do this does not necessarily mean that they can remember *when* they cached the food – indeed they may not need to remember when they did it. Thus, the tripartite basis of "what, where, when" that characterizes the occurrence of an event (or episode) does not commit the memory system to encoding all aspects of that event within the memory association. Ingenious experiments with particularly "clever" birds such as corvids indicates that, under certain training regimes, they can remember the time when an event took place (e.g., whether it was recent or remote in time). However, the generality of this memory pattern is unclear with similar experiments in laboratory rodents not showing successful temporal memory. Rather, rats are good at remembering the context where an event has happened – a "what-where-which" association. This is interesting as the ability to remember when something happened is often turned, through the act of recall, into a memory of "where was I" in the struggle to travel back in time to the moment when something happened. Once the context where something happened is recalled, recall of the events that happened there become possible. Experiments to follow up the structure of associative memory, and so reveal its neural basis, do not require expensive elaborate equipment so much as access to appropriate animals. New thinking to investigate aspects of the evolution of declarative memory in animals would be timely.

However, irrespective of the nature of declarative representations, two further and very important points about episodic memory are as follows. One is that memory encoding can occur at a single moment or – in the jargon of neuroscience – in a "single trial." Whereas skill learning is "trial-and-error" over multiple trials and thus error-correcting, there is no opportunity or necessity for episodic or episodic-like memory encoding to occur over multiple trials; events are generally unique and cannot be repeated. Thus, in true declarative memory, the system must be able to form, store, and retain information about a single event that has occurred at a discrete moment in time. Certain forms of declarative memory, such as semantic and spatial memory, do benefit from multiple "trials" – but such memories can also be successfully acquired at a single moment. Declarative learning is sometimes said to be "hebbian" rather than error-correcting, so named after the Canadian neuropsychologist Donald Hebb, who

proposed that changes in the strength of synaptic connections between neurons occurred quite simply when there were conjunctions in the firing of such neurons. A second intriguing feature of episodic memory relates to the qualitative character of retrieval – that recall involves a representation of an event but one that carries with it the sense of having happened in the past. In a successful act of recollection, there is no confusion with the here-and-now (the perceptual present). How this attribute of memory arises during recall is not well understood, but is an important topic in the domain of studies of "false memory" in humans.

The Existence of Different Memory Processes is Orthogonal to the Concept of Distinct Types of Memory

In addition to different *types* of memory, all forms of long-term memory can be considered as potentially passing through a number of distinct stages of processing. These are *encoding, storage, consolidation*, and *retrieval* (Fig. 3). The concept of a memory "process" is orthogonal to the "type" of memory, with these four features being in common across all types (Fig. 3a) and with these four processes occurring across a period of time (Fig. 3b).

In *encoding*, be it episodic memory or a motor skill, information that is to be remembered is represented in the brain as a pattern of neural activity and, in particular brain regions, is encoded in the sense of creating a lasting trace of having happened. Encoding happens over a timescale of milliseconds to seconds when an event or stimulus or action takes place, and this gives rise to *storage* in the form a change in the nervous system that outlasts the stimulus or action to be remembered (left-hand side of Fig. 3). In the case of working memory this change is sustained neural activity, whereas for long-term memory it is a structural/biochemical change in neurons. From a physiological perspective, encoding can be thought of as "induction" while storage is akin to "expression" – terms that also have their place in thinking about activity-dependent synaptic plasticity such as long-term potentiation. Encoding is modulated by attentional processes that are thought to be mediated, in part, by the action of acetylcholine. Cholinergic neurons that exert this modulatory role are found in the nucleus basalis of Meynert and the medial septum, and their activity impinges on nicotinic cholinergic receptors found on glutamatergic excitatory neurons. Their activation affects the membrane potential of the target neurons, and thereby the likelihood that glutamatergic neurons will engage the mechanisms of activity-dependent synaptic plasticity.

However, while information may be stored in long-term memory, there is no guarantee that this storage will be sustained. It may last a short while – an hour or two – and then fade. Or it may last much longer, although not necessarily for a lifetime. What determines whether a "memory trace" will last is a separate set of processes called *consolidation* in which other patterns of neural activity come into play, generally although not always after memory encoding, and these influence the fate of memory traces. There are two distinct aspects of consolidation – *cellular* and *systems* – although a strict distinction between these two is hard to sustain.

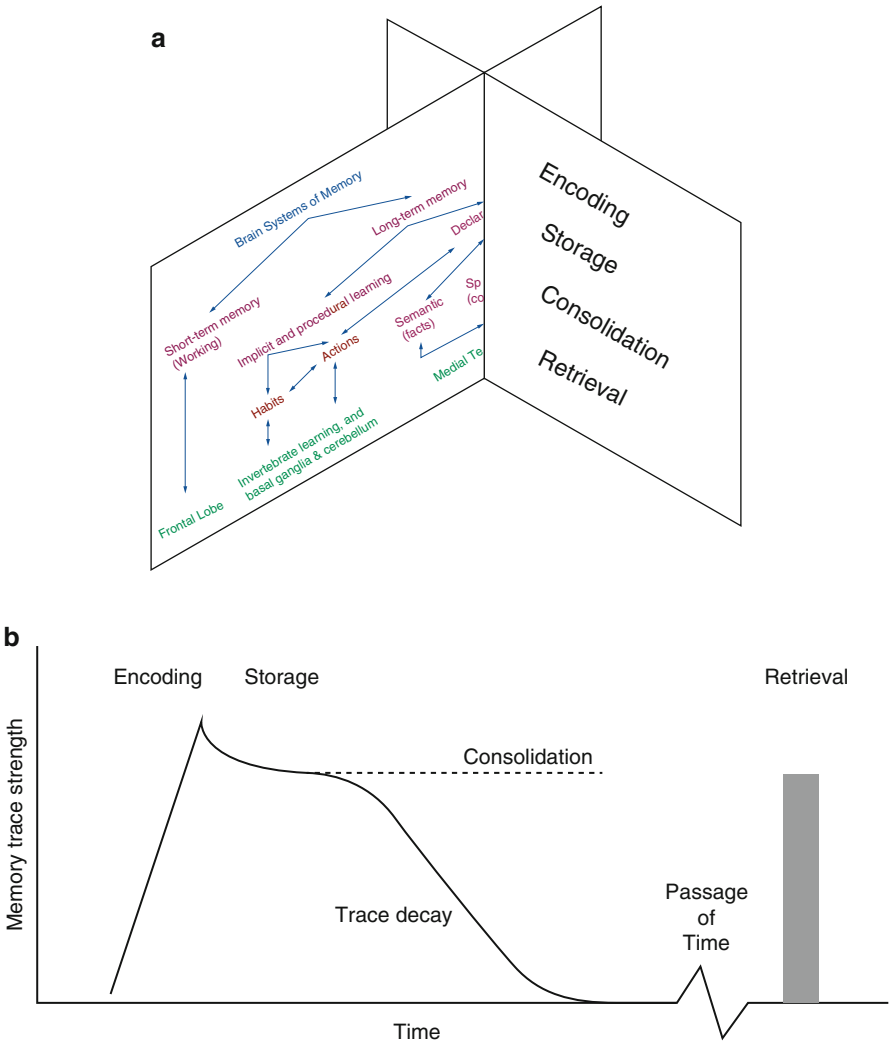

Fig. 3 *Memory trace strength over time*. (**a**) The concept of memory processes is orthogonal to that of different types of memory. (**b**) Notional changes in the strength of a long-term memory from its formation until its later recall. Y-axis shows trace strength and x-axis shows time. Trace strength rises rapidly at the time of encoding and then settles to a level corresponding to the stored state from which it may decay and be lost or may be subject to consolidation that enables stabilization via structural changes at synapses. After an indeterminate passage of time (x-axis), cues triggering retrieval reactivate the memory from its stabilized level of trace strength. This representation does not yet include new ideas about reconsolidation

The essential nature of *cellular consolidation* is the intracellular activation of signal-transduction pathways within neurons that affect the transcription of genes involved in plasticity and/or the translation of their mRNAs to create the plasticity-related proteins needed to change the functional and structural elements of synapses. These signaling systems may be activated by excitatory glutamatergic activation alone, but they are also influenced by the action of other neurotransmitters. Indeed, several theories of consolidation emphasize the role of neuromodulatory transmitters, such as catecholamines, which are released from the axon terminals of neurons that project from mid-brain nuclei such as the ventral tegmental area (dopamine), the locus coeruleus (noradrenalin), and the dorsal raphé nucleus (serotonin), onto neurons where memory encoding has taken place (e.g., in the hippocampus). Other theories point to the critical role of hormones that are released from various locations and circulate in the bloodstream. This apparent complexity in the regulation of consolidation is perplexing, but likely arises because of multiple factors that can influence whether a newly stored memory is to be retained or not (novelty, stress etc.). The common feature of each these different neuromodulatory systems is to activate or modulate diverse intracellular signal transduction pathways. Thus, while their effects are intracellular, there is clearly an interaction between cellular and systems aspects here as these neuromodulatory afferents arise from cell groups located in diverse areas of the brain.

The second form of consolidation, *systems consolidation*, is a process by which different regions of the brain interact to build network connections amongst participating elements and hence stabilize memory traces for lasting preservation. An early systems theory held that activity in the amygdala plays a key role in consolidating information encoded elsewhere (e.g., in the hippocampus or cortex). That is, if the consolidating activity occurs in the amygdala, the stored traces will be retained; but if this does not happen, the memory trace will fade. More recent theories have emphasized interactions, sometimes inappropriately characterized as memory transfer, between the hippocampus and neocortex that occur over the passage of time. The argument has been that detailed information stored in the cortex needs to be connected by an incremental process that builds the appropriate associative interactions that constitute a memory network. Hebb referred to these as "cell assemblies." Located in the neocortex, they enable rapid memory recall in response to retrieval cues that occur later. While distinct from cellular consolidation in terms of underlying neuroanatomy and timescale, the systems process likely uses overlapping neural mechanisms to those mediating the cellular process.

Why has memory consolidation evolved? It seems likely that the memory systems responsible for *encoding* and *storage* of information endeavor to create a veridical representation of the events or other information to be stored, affected only by the direction of an organism's attention. However, evolution has somehow ensured that the "modulatory" systems of *consolidation* that determine the importance and fate of memory traces are kept separate from encoding. Moreover, cellular consolidation might be thought of as acting as a kind of low-pass filter for

information that can then, given yet other factors, be subject to systems consolidation. Like candidates for a job that has been advertised, the brain first draws up a short-list (cellular consolidation) and then endeavors to integrate a subset of these candidates into the existing team (systems consolidation).

The fourth and last process is *retrieval*. The phenomenon of retrieval represents another major difference between short- and long-term memory. In short-term memory, information is completely lost when the sustained activity of delay-period neurons is interrupted. In contrast, in long-term memory, there may be long periods of quiescence when the memory trace is dormant. To illustrate this important difference, think of your last birthday or of some other event in the distant past that is important to you. You can readily re-create some of the events of that day in your mind's eye. But have you been thinking about them throughout the days, weeks, months, or years since that day? Of course not. So how do you have the experience of retrieving and recalling them? Neuroscientists think that retrieval cues (such as unintentionally seeing a triggering stimulus, or being deliberately asked to remember as this writer has just done to you) create specific patterns of neural activity that interact with stored representations to re-create memories in your mind's eye. This re-creative process is called retrieval. Retrieval is an essential part of the equation of distinct memory processes for it is the stage where a labile memory becomes active again and can so be communicated or acted upon. In that sense, retrieval completes the cycle that began at encoding.

However, retrieval can happen many times, interspersed with long periods when the memory is again dormant. In exciting new research, there is growing evidence that each act of retrieval can reinstate some of the conditions of original learning, and thereby the opportunity to erase, add to, or otherwise modify the stored memory trace. This occurs most obviously in an explicit learning context such as school where a child's understanding or memory of something may be pointed out by the teacher to be wrong. In such instances, it is clearly beneficial for the memory traces to be "updated" and the original incorrect traces to be overwritten – even if a lingering memory of having got something wrong in one's mind may persist for a while. In informal situations also, less matters of explicit education than of everyday experience, it is vital that memories be added to or amended. This process associated with retrieval is called *reconsolidation*. Understanding the mechanisms of reconsolidation has become a particular active area of research in the molecular neurobiology of memory.

The Study of Learning and Memory Involves Combining Methods from Numerous Levels of Analysis in Neuroscience

The careful study of neurological patients led to the discovery of psychological dissociations between distinct memory systems and pointed to the critical neuroanatomical structures and circuits necessary for memory. It also catalyzed experimental work aimed at understanding distinct memory processes. But neurobiologists seek a more profound understanding – and for that other methods than experimental observation are needed including invasive techniques that take us into the brain (Fig. 4).

Box 1 The Watermaze

The "watermaze" is a widely used task in behavioral neuroscience laboratories using rodents for studying the psychological processes and neural mechanisms of spatial learning and memory. It was developed at the University of St Andrews in Scotland and first described in two publications in the early 1980s. Place navigation in the watermaze is now often used as a general assay of cognitive function. Studies have included testing the impact on cognition of various disturbances of the nervous system (e.g., animal models of stroke, aging, neurodegenerative disease, and the potential impact of novel therapeutic and cognitive-enhancing drugs. The task has inspired computational neuroscientists and roboticists who are interested in navigation, and has recently been adapted for studies with *Drosophila*. A "virtual" watermaze has been developed for research involving functional brain-imaging in humans.

The Basic Task

The basic task is very simple. Rats or mice are placed in a large circular pool of water and required to escape from water onto a hidden platform whose top surface is hidden below the opaque water surface and so normally identified only using spatial memory (Box Fig. 1). There are no local cues indicating where the platform is located. Conceptually, the task derives from "place cells" that are neurons in the hippocampus which identify or represent points in space in an environment.

Distinction Between the Apparatus and the Tasks Trained within It

The watermaze is sometimes described as if it were a single task. Strictly speaking, the watermaze itself is no more than a pool of water in which a variety of different tasks can be trained. However, the simplest water escape learning task of learning to find a hidden platform in a single fixed location is generally the first step of a series of more involved training and testing protocols for investigating specific issues. Distinct protocols engage different mechanisms of navigation, learning, and memory.

Apparatus

The apparatus consists of a large circular pool, generally 1.5–2 m in diameter, containing water at around 25 °C made opaque by adding milk or another substance (e.g., latex liquid) that helps to hide the submerged platform (Box Fig. 2a). The water in the pool is filled and drained daily via a simple but automated filling and draining system. This choice of water temperature at around 13 °C below body temperature is sufficiently stressful to motivate the animals to escape, but not so stressful as to inhibit learning. There is a mild stress reaction on day 1 of training (indicated by elevated corticosterone), but this habituates over days. If the pool temperature drops to 19 °C, performance

(*continued*)

Box 1 (continued)
improves, but when the water temperature drops to 12 °C, it gets worse – reflecting the inverse U-shaped function relating stress to cognitive function.

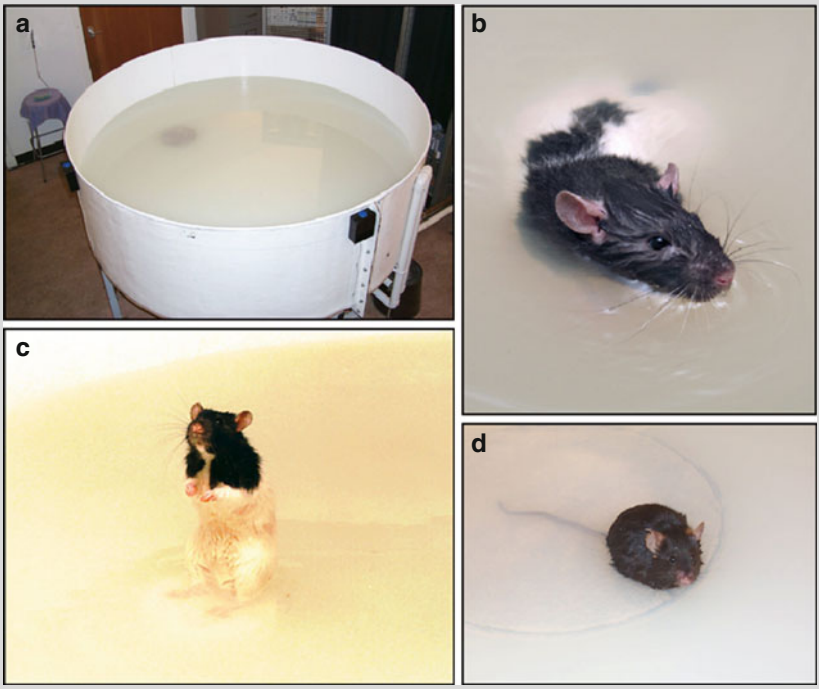

Box Fig. 1 *The watermaze*. (**a**) Photograph of the 2 m diameter watermaze used at the University of Edinburgh in Scotland. (**b**) Lister-hooded rat swimming to find the hidden platform. Hooded rats (e.g., Long Evans) have better vision than widely used white rats (e.g., Sprague Dawley). (**c**) Rats often rear up on the escape platform to inspect distal visual cues. (**d**) A transgenic mouse (PDAPP) on the escape platform. Use of a full 2 m pool is also recommended for mice, but with a large escape platform. It is unwise to make the pool smaller for mice

A digital camera is placed above the center of the pool to capture images of the swimming animal, and this is connected to a DVD recorder, and an online computer system running specialized tracking software. The top surface of the hidden platform, usually about 10–15 cm in diameter and thus between 1/50th and 1/100th of the surface area of the pool, is 1.5 cm below the water surface (Box Fig. 1a, c, d). The pool itself should be located in laboratory room with distinctive 2- and 3-D distal cues that aid orientation, with 3-D cues being particularly helpful. Alternatively, the pool may be surrounded with hanging

(*continued*)

Box 1 (continued)
curtains that occlude fixed room cues, enabling moveable cues to be hung inside that can be rotated relative to the room when this degree of precise experimental control is required. Again, it is best if these are 3-D rather than flat surfaces. The use of a cue-controlled environment has proved particularly helpful in studies of pattern completion.

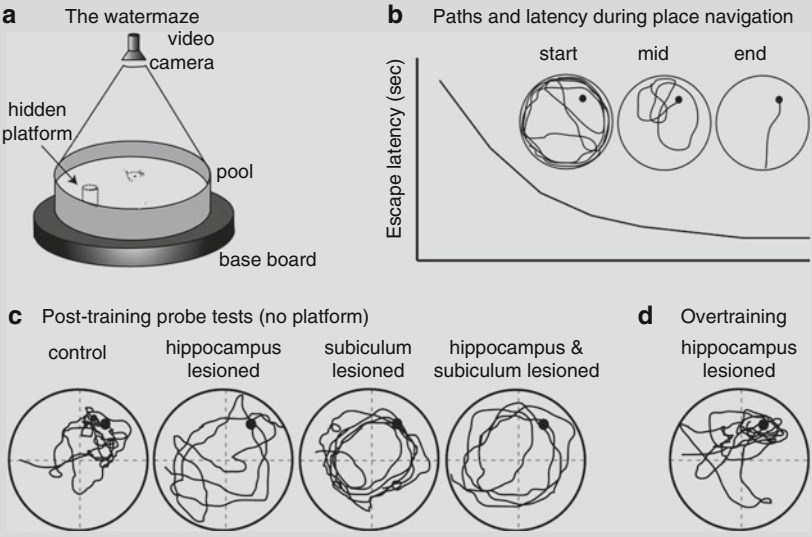

Box Fig. 2 *Basic procedures*. (**a**) Axonometric drawing of a typical watermaze setup with overhead videocamera and rat swimming to find the hidden platform. (**b**) Representative escape latency graph and swim paths across various stages of training – initial swimming at the side walls, then circuitous paths across the area of the pool, and finally directed path-navigation. (**c**) The hidden platform is removed for post-training probe tests. Whereas normal or sham-lesion controls swim to the target quadrant (NE, within *dotted gray lines*), rats with hippocampus, subiculum, or combined lesions do not. (**d**) Overtraining of hippocampus lesioned rats can result in quite focused search patterns in a probe test

Training Protocols

Spatial Reference Memory

Reference memory protocols are widely used in which the platform is in a fixed location relative to the room cues across days. The animals are placed into the water at and facing the sidewalls of the pool, at different start positions across trials, and they quickly learn to swim to the correct location with decreasing escape latencies and more direct swim paths (Box Fig. 2b). The tracking system measures the gradually declining escape latency across trials, and parameters such as path length, swim speed, directionality in relation to

(*continued*)

Box 1 (continued)
platform location, and so on. Observation of the animals reveals that, having climbed onto the escape platform, they often rear up and look around as if trying to identify their location in space. Rearing habituates over trials, but then dishabituates if the hidden platform is moved to a new location.

Probe Trials to Test Memory

During or after training is complete, it is vital that the experimenter conducts a probe trial in which the escape platform is removed from the pool and the animal allowed to swim for 60 s. Typically, a well-trained rat will swim to the target quadrant of the pool and repeatedly across the former location of the platform until starting to search elsewhere (Fig. 2c). This spatial bias, measured in various ways, constitutes evidence for spatial memory. Rats with lesions of the hippocampus and dentate gyrus, subiculum, or combined lesions, generally do poorly in post-training probe tests.

Impact of Overtraining

If rats with hippocampal lesions are given "overtraining" (typically consisting of a large number of trials over many days), their performance in probe tests can be quite good. Even rats with HPC lesions can show quite localized searching (Fig. 2d), particularly if the most septal pole of the longitudinal axis of the hippocampus is spared. These findings have led to the suggestion that spatial memory traces may be stored in cortex rather than hippocampus, but this point remains controversial.

Reversal Learning

Numerous other protocols have been developed to test specific hypotheses. Many involve the experimenter cryptically moving the hidden platform – the appearance of the water surface to the animals remaining unchanged. This might be a "reversal" procedure in which, after one location has been thoroughly trained, the platform is moved to a different quadrant of the pool. Because it is hidden, it is not apparent that anything has changed until the animal fails to find the platform in its usual place. The focus then is on how the animal reacts to this change and how quickly it learns the new location. The relearning that occurs in reversal protocol reflects cognitive flexibility, and has been used in a major genetic "factor analysis" of the determinants of watermaze behavior across different strains of mice.

The Atlantis Platform

As the animals sometimes "bump" into the submerged platform by chance, one useful innovation is an "on-demand" or "Atlantis" platform (named after the lost city) that is initially at the bottom of the pool and only becomes

(*continued*)

Box 1 (continued)
available when the animal swims in its vicinity for some predetermined time (Box Fig. 3a). An automatic release system allows the platform to rise gently to near the surface of the water (but it remains hidden). This procedure results in the acquisition of a highly focused searching strategy focused on the target location during training. Reversible inactivation of the hippocampus with a drug that blocks excitatory neurotransmission *after* training is complete results in animals displaying localized searching at inappropriate places in the pool, indicating that they retain the procedural strategy of searching during hippocampal inactivation but do not know where to search (Box Fig. 3b). In contrast, pharmacological inactivation of fast synaptic transmission *during* training results in a failure to develop this search strategy because the animals cannot learn where to execute the strategy in the pool. The accuracy of searching can also be measured using a zone analysis that measures time spent in a virtual zone around the place where the platform is located.

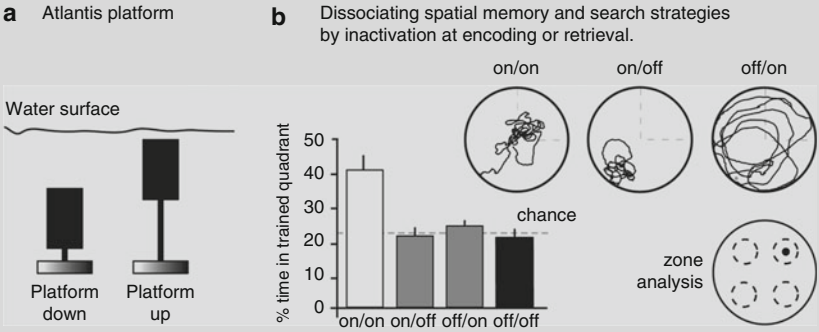

Box Fig. 3 *Modern developments*. (**a**) The Atlantis Platform. The hidden platform is at the bottom of the pool where the swimming rat cannot bump into it by chance. Online automated data capture of swim paths is used to determine whether the rat swims within a virtual zone around the platform's location, raising the platform to within 1.5 cm of the water surface when a criterion is reached. This protocol trains highly focused search patterns. (**b**) Reversible hippocampal inactivation with a glutamate antagonist during training (encoding) or at retention (retrieval) results in poor probe test performance compared to controls that have the hippocampus working continuously. Analysis of the search patterns show that rats trained with the hippocampus "on" during encoding, but "off" at retrieval, display searching at the wrong location in the pool as quantified using a zone-analysis

Delayed Matching-To-Place and One-Trial Learning

In other protocols, sometimes called "working memory" or "delayed matching-to-place" (DMP) protocols, the platform is moved to a new location each day (Box Fig. 4a). In this procedure, the animal can never know where the platform is hidden on trial 1 of each day. However, once it finds the

(*continued*)

Box 1 (continued)
platform in the first training trial (usually after about 60 s of searching), it can generally encode this new location in that one trial. This is shown by the animal finding the platform much faster on trial 2 and subsequent trials of that day. In effect, this procedure enables the study of repeated instances of one-trial learning. The intertrial interval (ITI or memory delay) between trials 1 and 2 can then be systematically varied to explore how well 1-trial spatial memory is remembered, a procedure with some similarities to delayed matching and non-matching tasks used to examine recognition memory. Rats with complete hippocampal lesions never show rapid 1-trial learning required in the DMP task (even after extended training) and are just as poor at a short ITI between trials 1 and 2 as a long one (Box Fig. 4b). In contrast, treatment with an NMDA antagonist such as D-AP5 results in a selective deficit in memory at a long ITI, but the animals can remember over short memory intervals between trials – a finding that links the impact of D-AP5 on behavior to its well-known effect on long-term potentiation. Studies of transgenic mice with regional-specific deletions of the NMDA receptor have revealed a role for area CA3 in the hippocampus in rapid learning using this DMP paradigm.

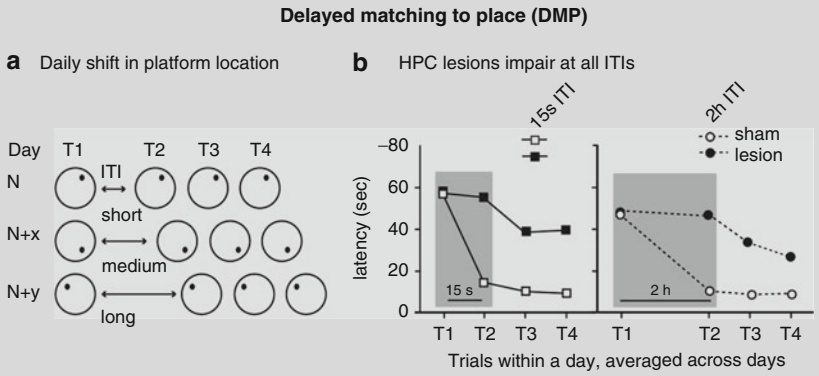

Box Fig. 4 *Delayed matching to place to study one-trial learning and memory.* (**a**) The training protocol involves four trials per day with the location of the hidden platform moved between days. Training can continue indefinitely with this protocol, enabling within-subject drug manipulations throughout the life span and averaging across days. Acquisition typically takes 8–10 days. (**b**) Results averaged across days for each of the four trials of a day. On trial 1 of each day, the animals search for the platform, typically taking 60 s to find it, encode its location and show fast escape latencies on trials 2–4. Hippocampal lesioned rats cannot learn this task irrespective of the intertrial interval (*ITI*) between trials 1 and 2. Shaded zone shows the ITI between T1 and T2 extended to 2 h. Sensitivity to ITI is shown when intrahippocampal infusions of various neurotransmitter antagonists are used (data not shown)

(*continued*)

Box 1 (continued)

Other Protocols

One of the exciting features of using the watermaze is that developing new protocols is quite easy. They can be designed to address specific questions. Procedural variants to date include alterations to the apparatus, such as constraining the path of the swimming animal to minimize navigational demands such as an annular watermaze, or to simply the memory demands as in the radial watermaze. This combines the virtues of the radial maze with the ease of training to escape from water and has proved invaluable in testing transgenic mice expressing familial Alzheimer mutations.

Treatments and Control Procedures

A wide variety of treatments have been explored including lesions, drugs, and molecular genetic alterations. These alter watermaze "performance" in various ways, but experimenters must be cautious about the interpretation of such deficits as such alterations need not be specific to spatial learning or memory processes per se. Lesions or drugs may have a direct effect upon learning mechanisms, and many seem to do so, but they may also affect an animal's ability to see the extramaze cues (a sensory deficit), or their motivation to escape from the water, or to translate knowledge into action, rather than learning per se. Factor analytic studies reveal that many molecular genetic alterations influence the probability of mice to stay at the sidewalls (thigmotaxis) instead of swimming into the center of the pool. These performance effects are statistically independent of the effects on spatial memory as has been shown using factor analysis of different mouse strains.

Accordingly, treatments must be accompanied by relevant control conditions. A common control protocol is to include trials in which the escape platform is made visible, the idea being that treatments which merely affect motivation to escape should impair performance in this task as well as the basic task. It is unclear how sensitive this assay really is, but it does provide a "first-pass" at detecting gross sensorimotor abnormalities. As blind rats have been claimed to do surprisingly well in the watermaze (except in probe trials), more taxing psychophysical techniques have been introduced offering precise control of the spatial frequency of cues that are more sensitive to subtle visual deficits. The use of sham-lesioned animals, vehicle-infusion conditions, "floxed" mice, and other pertinent manipulations has also become widespread to ensure that any alterations in spatial learning in the experimental group are not an unintended by-product of achieving the treatment.

(*continued*)

Box 1 (continued)

Strengths and Limitations of the Watermaze as an Assay or Learning and Memory

As understanding of the impact of various treatments on diverse aspects of cognitive function has developed, for example, executive function, the watermaze has been subsumed into larger test batteries for investigating diverse aspects of brain function. A clear virtue of the task is that the various protocols are so sensitive to manipulations of normal brain function in many brain areas, not just the hippocampus, that these can be used almost like a "litmus test" of the "normality" of cognitive function. This brings behavioral observations of function into fields of neuroscience that have historically relied exclusively on endocrine measures (studies of stress hormones), neuropathology (stroke research), biochemical analyses (Alzheimer's disease), or electrophysiology (development of cognitive-enhancing drugs). The limitation, or analytical weakness, is that watermaze tasks which are affected by such a wide variety of treatments are gradually being revealed as having less "specificity" than was once believed. Notwithstanding these limitations, the watermaze remains a still widely used task in behavioral neuroscience.

Further details of the "water maze" with the citations of relevant articles may be found at the online at www.Scholarpedia.org

It is important for readers of this text to appreciate that new techniques are often the driving force of discovery in science – in neuroscience no less than in other disciplines. Think of the telescope in astronomy and of the worlds beyond this planet that this invention opened up. However, to be really useful for science, the application of new methods must be motivated by specific scientific questions for which existing techniques fail to provide the definitive evidence that science seeks. Over the past 50 years, key questions for understanding memory have included: *Are there specific firing patterns of neurons during learning?* This question led to the application of intracerebral single-cell recording in behaving animals, and thence to the discovery of delay-dependent neurons in the frontal lobe of nonhuman primates (short-term memory). Key findings in similar research on long-term memory include place cells, head-direction units, and grid cells of the medial temporal lobe in rodents (spatial memory) all of which are now typically recorded using "tetrodes" (Fig. 4a). These enable the identification of individual cells using extracellular recording techniques. While first introduced many years ago, the full potential of tetrode recording is only now being realized in studies of place cells (Fig. 4b) and other types of neurons. Another question has been: *Can synapses change in strength?* And the corollary of this has been: *What are the circumstances and mechanisms responsible, and what role do such changes play in storing information?* These questions led to the discovery of long-term potentiation in whole animals in vivo, the development of brain slice

(*continued*)

Box 1 (continued)
methodologies and in vitro physiological paradigms, and from these to a range of techniques from computational models of information storage to the detailed analysis of the receptors responsible for mediating and altering synaptic strength and associated intracellular signal transduction and trafficking pathways. The generic framework for thinking about such changes is the glutamatergic synapse (Fig. 4c) that is discussed in detail elsewhere in ▶ Chap. 3, "Ionotropic Glutamate Receptors" by Rolf Sprengel. More recently, neurobiologists have become interested in asking: *"Can such changes in neurons be seen as they happen in real time?"* This interest has been instrumental in the application of in vitro and now in vivo confocal microscopy to synaptic plasticity, leading to visualization of the dynamic changes in dendritic branching, filopodial, and dendritic spine growth, and to the use of an exciting range of fluorescent and other markers that can be engineered to be activated by specific molecular promoters. Very recently, using elegant multiphoton confocal imaging techniques, it has become possible to record dynamic changes in the shape and size of individual dendritic spines during different patterns of LTP-inducing stimulation (Fig. 4d). While the field is still at the frontiers in such work, it is to be hoped that visualizing the various stages of encoding and storage will become possible.

Particularly effective for progress in this area of science is to try to combine these elegant physiological, pharmacological, and optical techniques with behavioral paradigms in which living animals are trained on analytically useful learning and memory tasks. This may seem daunting, and to those in certain laboratories of the developing world, may be too difficult or expensive to contemplate. In practice, such labs should not be daunted, for relatively simple techniques still have a very important place if used wisely in pursuit of a well-constructed question. Moreover, combining techniques is also an exciting challenge at the frontier of contemporary neuroscience. One example involves the use of optogenetics (Fig. 4e) in which light at various wavelengths (e.g., blue light) can be directed via an optical fiber into specific areas of the brain where a virus (e.g., an adenovirus) has been used to express a gene encoding light-sensitive "channelrhodopsin." The light thereby activates sodium channels that in turn result in cell firing. This approach, and its inhibitory companion involving green light and halorhodopsin, is now starting to be used to ask causal questions about the role of place cells in spatial navigation about neuromodulation and about spatiotemporal dynamics of the neural representations of memory in the brain.

These physiological, optical, and interventional techniques can be used to learn more about neural mechanisms, but the resulting findings have then to be tested or "reverse-engineered" back to the behavioral tasks that actually reveal learning and memory. Many paradigms are in use in animal studies to study

(*continued*)

Box 1 (continued)
recognition memory, spatial memory, episodic-like memory relational learning, and the formation of paired-associates, motor-skill learning, and so on. Figure 4f, g show two new paired-associate tasks. In one, rats are trained to learn that stimulus A is associated with stimulus B, and that B is associated with C, the question of interest being whether these two bits of information can be put together such that the animal can recognize a link from A to C. In the second, rats are trained to associate a particular flavor of food with a spatial location in a test arena with six separate paired associates being trained concurrently. They are cued before entering the arena on any trial with one of the six flavors, and this indicates that this particular flavor is available on that trial and none of the others. These and other related paradigms are opening up research into the organization of information in long-term storage after systems consolidation has taken place. Often simple techniques, including a number of behavioral techniques such as the watermaze (see Box 1), can be put together with recording and optogenetic approaches to address interesting new questions in the neurobiology of memory.

Making Declarative Memories
Whereas short-term memories are held in reverberating patterns of neural activity that have been shown using functional imaging and single-cell recording in awake primates, long-term memory is thought to involve structural changes to the nervous system that alter the strength of synapses and the connectivity of networks. This applies equally to semantic (fact memory), episodic memory (memory for events, including sequences of events), and spatial memory (locations and navigation). A common feature of these is that information appears to be encoded initially in parallel in both the hippocampal formation and neocortex, but eventually retained as long-term memory traces or "engrams" in the neocortex.

There Are Three Types of Declarative Memory: Semantic Memory, Episodic Memory, and Spatial Memory
By *semantic memory* is meant the organized knowledge of the world that humans possess within language, in everyday life, and various domains of expertise. While arguments prevail about the organization of such information, certain models imply some kind of tree structure in which we, for example, distinguish animate from inanimate, mammals from birds, flying birds from flightless, and so on. This kind of information is acquired slowly through childhood or, in the case of professional knowledge, through painstaking scholarship over time. Semantic memory is believed to have an associative structure with memory traces distributed across wide regions of the cerebral cortex, although the learning machinery for acquiring semantic associations is

(*continued*)

Box 1 (continued)
thought to involve the hippocampal formation in the medial temporal lobe. *Episodic memory* is different. This pertains to events that by their very nature are unique. What you did last Christmas, the births of your children, and so on. Importantly, not only can such unique events be remembered, but their very character is such that the memory system has to be able to do it in one trial. You cannot arrange for your children's births to happen a second time if you failed to remember it very well the first time! Similarly, watching a famous goal scored in an important football match, or the final winning putt of a golfing championship, or the moment when a car crash happens are all examples of events that can only happen once. As the French photographer Henri Cartier-Bresson put it so well, photography is sometimes about *". . ..the decisive moment."* Interestingly, the way photography is used, as tourists or within one's family, is precisely to capture events. Like semantic memory, episodic memory has an associative structure, but one that links unique events to the spatiotemporal context in which they occurred. *Spatial memory*, not always considered a separate entity, refers to the ability to learn the layout of the various objects and landmarks of the world around us, and to represent their locations relative to each other in an allocentric ("object-centred") framework. This framework is clearly different from a semantic representation (associations have to do with location rather than meaning) and different from event memory as the representation will, even if learned rapidly, be timeless. Importantly, spatial representations in memory are also vital both for navigation – getting from A to B – and, separately, for providing a context with which to remember events. Thus, while the episodic memory system will encode and store a specific event, it needs to tie into spatial memory to form an association between the event and the context where it happened. Such associations appear to be formed automatically – you do not remember an event as if it occurred in isolation.

Numerous protocols exist for analyzing semantic, episodic, and spatial memory in humans but their study in animals is more problematic as the "semantic-like" and "episodic-like" analogs have to be carefully validated. For a long period, work using nonhuman primates was conducted using a system called the Wisconsin General Testing Apparatus (WGTA) in which various protocols were developed for examining the neural basis of discrimination learning, spatial learning, and recognition memory. By means of the technique of experimental lesions, numerous discoveries were made including the critical role of the perirhinal cortex in recognition memory and the differential reliance of egocentric and allocentric spatial memory codes in distinct brain regions. The WGTA apparatus has now been replaced by newer digital touch screen methods controlled by computers with virtual "object-in-place" learning claimed as a model of episodic-like memory.

(*continued*)

Box 1 (continued)

Similarly, work using rats, mice, and flies has been dedicated to discovering the role of different brain regions in different forms of memory, although the extent to which rodents (still less flies) can be said to possess semantic or episodic memory is very much a matter of debate. One popular task involves the spontaneous exploration of novel objects, as a test of recognition memory (a component of episodic memory). Another involves context fear conditioning (in which learning takes place in one trial), but this is more likely a form of very rapid value learning in which the animal learns that a specific context is a place in which to be afraid by virtue of its association with shock. This need not entail relational "what-where-when" learning. The discovery of place cells in the hippocampus led to spatial memory tasks becoming very popular – and as a field of investigation in its own right. In laboratory rodents, spatial memory is investigated using T-mazes, the radial maze, open-field arenas, and the watermaze. A heat-maze analog of the watermaze has recently been developed to study learning in *Drosophila* in which the fly has to find and remember the "cool-spot." In Box 1, the apparatus and various protocols for using the "watermaze" are described in more detail. This is an apparatus that is easy to construct, and the various protocols are straightforward to follow and relatively inexpensive to run.

Physiological, Pharmacological, and Molecular Engineering Approaches to the Study of Memory

Mechanisms of Activity-Dependent Synaptic Plasticity Are Involved in Encoding and Storing Information into Long-Term Spatial Memory

In ▶ Chap. 3, "Ionotropic Glutamate Receptors" by Rolf Sprengel, excitatory glutamate receptor-dependent synaptic transmission is described in detail. Fast synaptic transmission is mediated by AMPA receptors whose associated ion channel allows sodium (Na^+) ions to enter the postsynaptic neuron and so mediate an excitatory postsynaptic current. Another subtype of glutamate receptor is the *N*-methyl-D-aspartate (NMDA) receptor (see Figs. 4c and 5b). It is a long molecule that loops in and out of the lipid bilayer of the postsynaptic membrane within the postsynaptic density, and each NMDA receptor includes an ion channel through which charged molecules can travel. It is now well established that, in specific brain regions, the ionotropic NMDA receptor is critical for memory encoding that is achieved by exploiting an intriguing feature of its biophysics (NMDA receptors mediate distinct functions in other circuits). Embedded in the ion channel is the divalent cation Magnesium (Mg^{2+}) that ordinarily blocks the channel. However, if

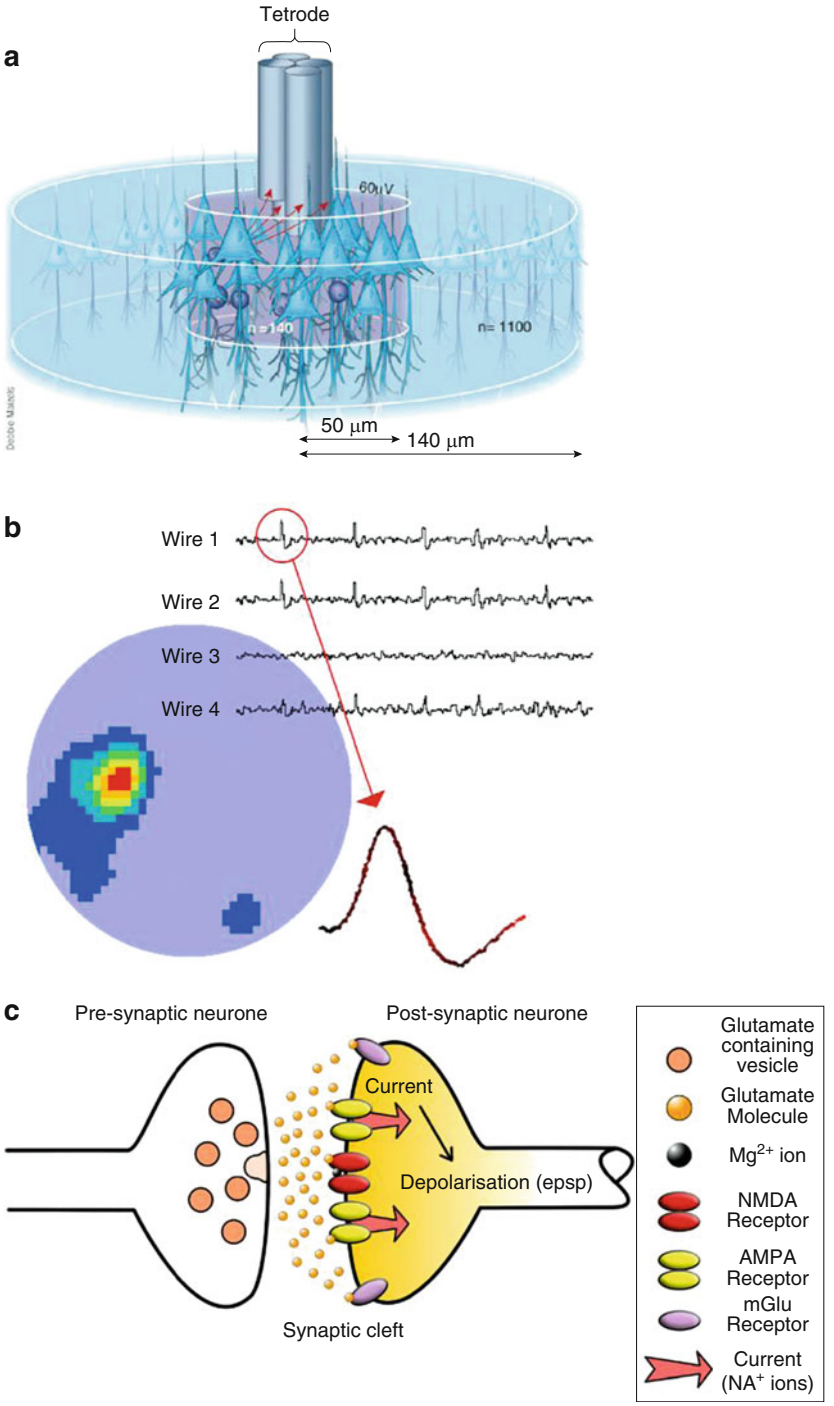

Fig. 4 (continued)

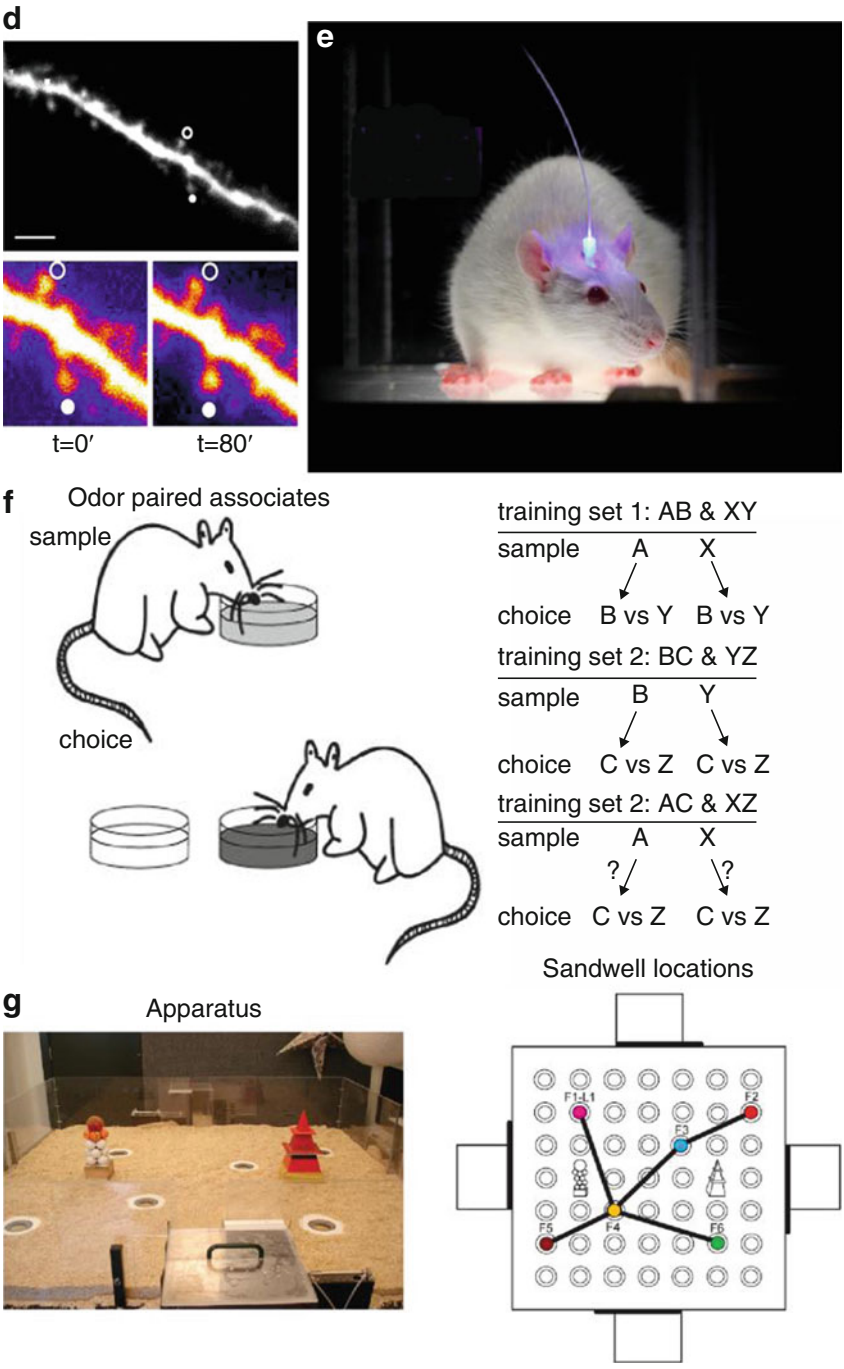

Fig. 4 (continued)

the presynaptic neuron and the postsynaptic neuron fire at the same time, the voltage on the postsynaptic side changes so much that Mg^{2+} is ejected from the channel. Glutamate is released from the presynaptic side, whereupon it binds to its so-called ligand binding site postsynaptically; provided the Mg^{2+} has been ejected; this binding of glutamate now has the effect of successfully opening up the blocked ion channel. Calcium then flows though the NMDA channel and, once inside the postsynaptic neuron, selectively activates various intracellular signal transduction cascades. These can do several different things. One of these is to change the strength of the synapse when it is used in future. It may increase the strength of the connection between neurons as in long-term potentiation, or may decrease it as in long-term depression. These are the classical examples of the phenomenon of *activity-dependent synaptic plasticity* and many neuroscientists believe these processes are essential for learning and memory.

An illustrative and critical experiment implicating NMDA-receptor-dependent plasticity in learning involved blocking the receptor in vivo with a drug called D-2-amino-phosphonopentanoic acid (D-AP5). This drug binds in a competitive way to the ligand-binding site for glutamate, such that glutamate released by the axon terminals of the presynaptic neuron is unable to activate NMDA receptors. When this drug was infused into the hippocampus of rats, the animals were unable, or at least had great difficulty, in learning a reference memory task in the watermaze and later showing localized searching in the probe test (Fig. 5a). This is an intriguing finding for three reasons. First, the drug achieves a more subtle manipulation of the nervous system than an explicit lesion (which damages brain tissue) or even of pharmacological blockade of normal synaptic transmission as would occur with an AMPA receptor antagonist. In the presence of D-AP5, the AMPA receptors that mediate fast synaptic transmission work in an essentially normal manner. Second,

Fig. 4 *Contemporary techniques in the neurobiology of learning and memory.* (**a**, **b**) Introduced some years ago, extracellular tetrode recording makes it possible to record the activity of individual closely packed cells by comparing the neural activity of each of four closely spaced fine wires. This is giving greater precision to learning about the neural codes with which memories are represented (Source: (**a**) From Fig. 1 of Buzsaki (2004), (**b**) Source: By kind permission of de Hoz L, Wood E, Univ. Edinburgh; (**c**) The glutamatergic synapse, described in detail elsewhere in this book, including two ionotropic and one metabotropic receptor. Glutamate receptors are to be found in the dendritic spines of excitatory neurons (Source: By kind permission of Collingridge G, Doherty A, Univ. Bristol). (**d**) Multiphoton confocal microscopy will be used increasingly to visualize structural and functional changes at synapses. Here is an example of the imaging of spine volume in a study of synaptic tagging and capture (Source: Is taken from Fig. 1a of Govindarajan et al. 2011). (**e**) Optogenetic stimulation of neurons deep in the brain will hopefully reveal the causal role of specific patterns of activity (Is taken from the work of Deisseroth K, Gradinaru V, Carnett J, Stanford University). (**f**) Behavioral studies are vital in relation to these elegant physiological, biochemical, and anatomical studies. Learning the relational association of odors to reward has provided evidence that the episodic-like memory system of rats can function in an inferential manner (Source: Is taken from "The Hippocampus Book" published by OUP). (**g**) The event arena in which flavor-place associates can be learned, with the layout of the different paired associates arranged in a schema. Such studies may help us learn about the organization of semantic-like memory systems

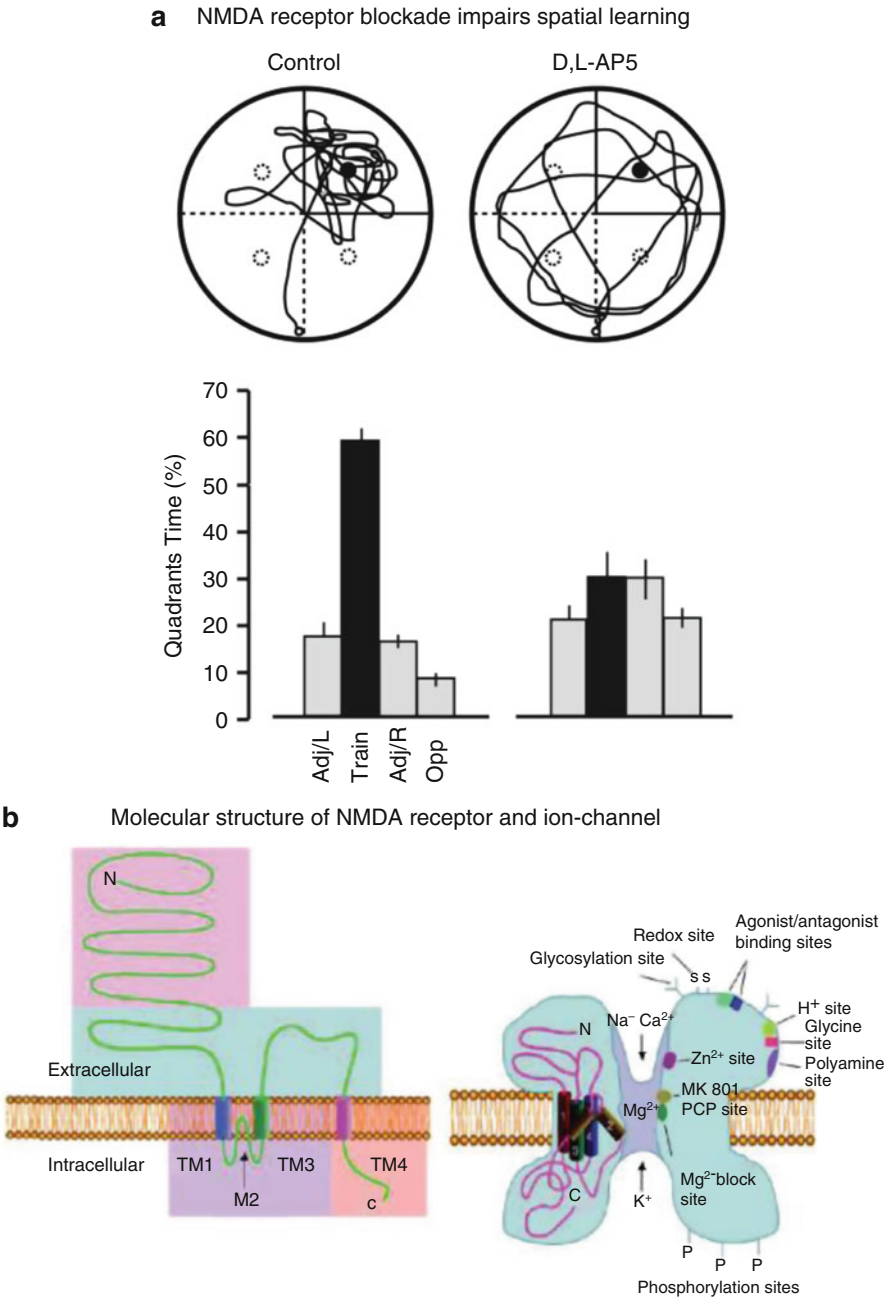

Fig. 5 (continued)

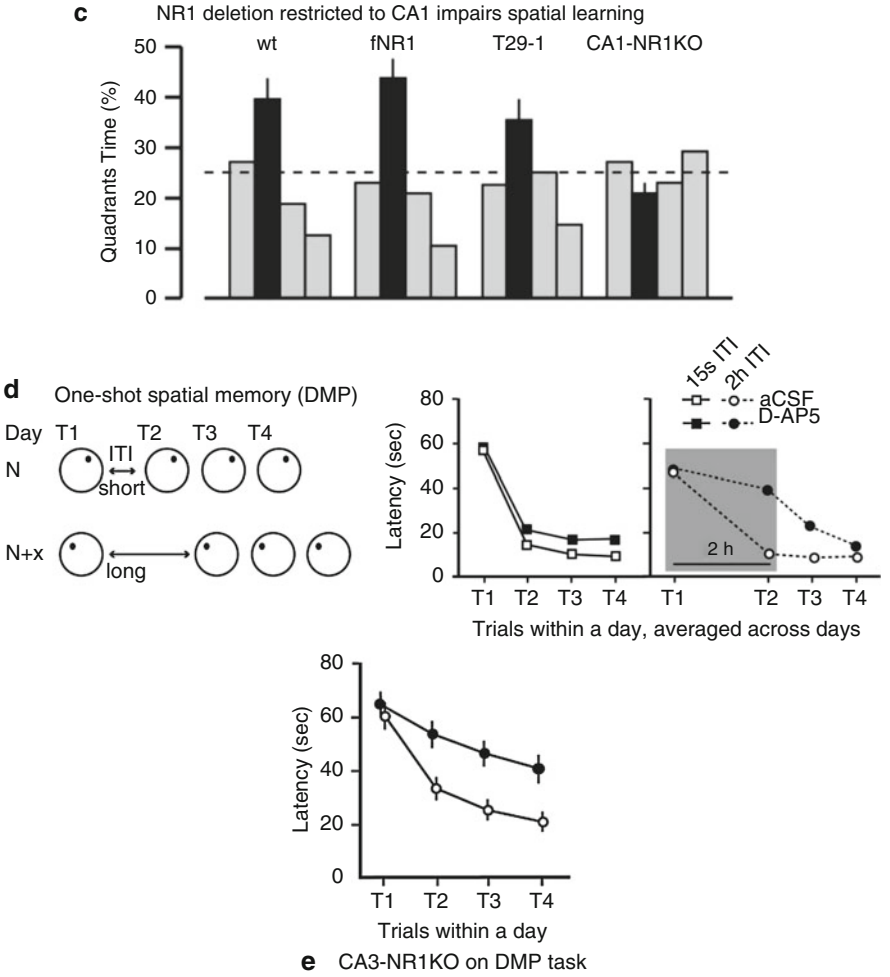

Fig. 5 *The role of hippocampal NMDA receptors in learning and memory.* (**a**) Microinfusion of the NMDA receptor antagonist AP5 impairs the acquisition of spatial memory in the watermaze. Note different paths of control and AP5 treated groups during the probe test. (**b**) The molecular structure of the NMDA receptor and its associated ion channel. (**c**) Selective cell-type and regional deletion of the NMDA receptor in area CA1 of the hippocampus is sufficient to impair spatial learning in the watermaze. (**d**) A different protocol for the watermaze called the "one-shot" or delayed matching-to-place task (described in Box 1). The escape platform is moved each day. Performance is characterized by long escape latencies on trial 1 of each day, when the platform location is not known, but faster escape thereafter. AP5 only affects performance when memory of trial 1 has to be maintained for a long time (e.g., 2 h). (**e**) Selective deletion of the NMDA receptor in area CA3 results in slow daily acquisition of the one-shot task

while this behavioral deficit in the watermaze looks like a learning deficit, it could be argued that it is merely masquerading, as a learning deficit with the animals functionally blind rather than functionally amnesic. The experimenters involved in doing this work were concerned about this because NMDA receptors are not only present in the hippocampus, but also in the visual cortex. Accordingly, a visual discrimination learning modification of the watermaze was devised that involved visibly distinct rigid and floating platforms that could occupy different locations in the pool. The rigid platform (painted gray) provided escape from the water whereas the floating platform (painted with black and white stripes) did not. The drug D-AP5 had no effect on learning to discriminate these visibly distinct platforms – indicating that the animals were not blind. Third, using a classical pharmacological dose–response analysis, it was established that the behavioral deficit induced by D-AP5 occurred at cerebral doses of the drug that blocked synaptic plasticity as measured in the physiological assay of long-term potentiation. This study illustrates several points made earlier in this chapter – the value of approaching a question using multiple techniques and levels of analysis, the need for control conditions, and the wisdom of a skeptical approach to discovery.

Modern neurobiological studies of learning and memory have moved on in an exciting way. Pharmacological experiments are, in truth, a bit messy because the spatiotemporal spread of drug action is very hard to control. Nor is it clear what range of cell types are affected. Nowadays, neurobiologists tend to prefer additional molecular evidence before being persuaded. The advent of transgenic and knockout animals could provide just such evidence because this technology offers the opportunity to target the NMDA receptor in a very selective cell-type and regionally specific manner. The temporal precision of these newer techniques is not as good, but the data that has been secured so far complements and extends the pharmacological evidence.

Unfortunately, knocking out the NMDA receptor is tricky because, if done all over the nervous system, the offspring are embryologically lethal. This is because the receptor has diverse functions including a critical role in building neural circuits in the brain. Accordingly, the laboratory of Susumu Tonegawa at MIT set about a different strategy using a procedure called Cre-Lox. This involves creating two lines of mice, one harboring a construct called Cre, which is a kind of molecular scissors, and another line of mice in which the gene of interest had been “floxed” – a molecular marking with sequences at which the scissors are supposed to act. Crossbreeding these two lines of mice yielded some lines of mice in which the NMDA receptor was selectively removed. Specifically, by placing the gene that acts as scissors downstream of a promoter that is expressed only in adult mice and expressed particularly strongly in the hippocampus in certain lines of these mice, it was possible to target NMDA receptors in just one tiny region of the hippocampus – area CA1. They went on to show that this very unusual line of CA1-NMDA receptor-knockout mice could not learn the watermaze very well (Fig. 5c). They also showed abnormalities in the firing of hippocampal place cells that took the form of a kind of “fuzziness” of the spatial representation. The cells are, together with head-direction and grid cells, critical components of the navigational system of the

brain and thus central to accurate spatial memory. Other molecular genetic experiments that have targeted different aspects of hippocampal circuitry have provided evidence for differential roles of areas CA1, CA3, and the dentate gyrus with a cell and regional specificity that could not be achieved through lesion or pharmacological experiments alone. Not withstanding these differences of technique, findings using selective pharmacological blockade of the hippocampus and knockout of the NMDA receptor in area CA3 in the variant of the watermaze called delayed matching-to-place (see Box 1) were remarkably similar (Fig. 5d, e).

Synaptic Plasticity is Critical for Storing Information in Specific Neural Networks

The idea that synaptic plasticity mediates the storage of information for memory goes back at least a century – the great Spanish anatomist Ramon y Cajal discussed the idea of a growth in dendritic arborizations in his Croonian Lecture at the Royal Society in London in 1894. The Polish neurophysiologist Jerzy Konorski and the neuropsychologist Donald Hebb in Canada discussed its importance with reference to the physiological conditions in which such anatomical changes might take place. Later, the English mathematician David Marr developed the first computational models of how the intrinsic architecture of the cerebellum and the allocortex could, when coupled to changes in synaptic weights between neurons, store information. For the cerebellum, his model related to the timing and precision of fast learned movements; for allocortex, it was a model of memory for the events of a day. Soon after Marr's early computational models, the psychiatrist and neuroscientist at Columbia University, Eric Kandel, led a team that provided the first concrete evidence that changes in synaptic weights occur during simple forms of learning in the sea-slug *Aplysia*. Thus, anatomy, physiology, and computational models come together in thinking about the storage of memory traces.

There are many reasons why changes in the strength of synapses are thought to be so important. These include *persistence, input specificity*, and *associativity* – each being properties that are also seen in the expression of long-term potentiation in the hippocampus. *Persistence* implies the possibility that structural/biochemical changes at synapses can persist over a long time, despite protein turnover. *Input specificity* relates to the synaptic nature of the change, rather than a change in membrane excitability that may occur all over a neuron or at least throughout one or more dendritic branches. The presence of large numbers of synapses in individual cortical cells (possibly in excess of 10,000 per neuron) implies considerable storage capacity. *Associativity* is important as it implies the possibility of connecting one piece of information with another – as in A is associated with B, and B with C, and so on. In relation to the earlier point about the importance of memory retrieval, associativity indicates a retrieval cue may be able to retrieve a representation of information with which it was associated at the time of encoding, but is now no longer present. In that sense, memory reflects a process of bringing things back to mind.

Neuroscientists would like to know more about the underlying molecular and cell-biological mechanisms of synaptic plasticity and memory storage. However, before turning to this important topic of "The Structural and Functional Expression of Memory Traces" section below, it is helpful to illustrate how merely changing the strength of connections in a network can store an association between two events. There are many different forms of "distributed associative networks," of the kind to be described, and broadly speaking they all operate according to similar principles. They differ in terms of their underlying neural architecture (e.g., whether there are feedback connections or not), how information is represented (e.g., the presence of what is called sparse coding), and the type of learning rule used to change connectivity (e.g., whether the rule is Hebbian or involves error correction). Such differences are, of course, important and have been the subject of intense study for over 40 years when these kinds of networks have been examined. However, what they all have in common is that all distributed associative memory networks rely on changes in synaptic connectivity to store information, and they all illustrate the conceptual principles of persistence, input specificity, and associativity.

As shown in Fig. 6a, imagine an arrangement of four wires running horizontally and four running vertically. These cross over at points that correspond to "switches." The switches can be off or on, and they all start in the off state. The role of these switches is to switch a small amount of current from a horizontal going wire to a vertical one, but they only do this if the switch is in the on state. Such an arrangement may seem artificial when viewed from a neuroscientist's perspective, but one can begin the "validity-test" by imagining that these wires correspond to the axons of an afferent pathway that synapses *en passage* (horizontal wires) passing through a dendritic tree of the neurons to which they are connected (vertical wires). Stimulus information is represented in a binary manner, reflecting the All-or-Nothing principle of neuronal communication. These features are arguably realistic.

Let us suppose that stimulus A (0011) occurs at the time that the postsynaptic neurons are firing in the B pattern, with 1010, indicating that the first and third neurons are firing. If this network stores information according to a Hebbian

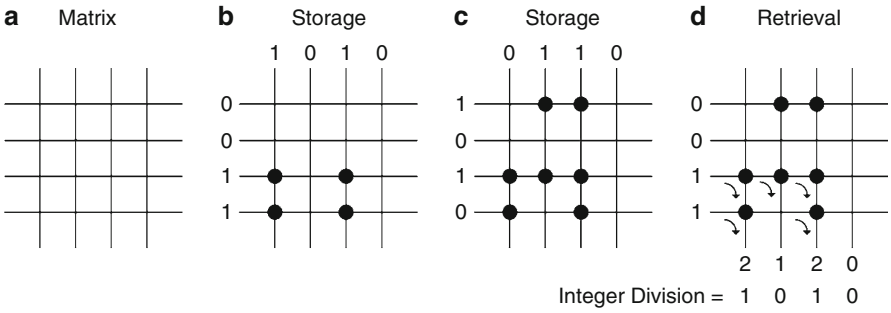

Fig. 6 *Distributed associative memory.* As described in the text, this system of wires and switches laid out in a network and changing in line with the rules of "hebbian" synapses can store information in an overlapping and distributed manner

principle, the growth process of which Cajal and Hebb both wrote will occur if and only if the presynaptic neuron and the postsynaptic neuron are firing together. Thus, no change in a switch will occur when a 0 meets a 0, nor when either 0 meets a 1 or a 1 meets a 0, but a switch will turn on when a 1 meets a 1. As shown in Fig. 6b, the state of the switches after this encoding event will be a spatial pattern of states that correspond to the storage of the association between A and B. This is a "distributed" system in that the memory trace is not any one location, but at multiple locations throughout this simple network.

Now let us suppose some further information comes in, such as that the input B is associated with a firing pattern corresponding to the stimulus C. This reflects 1010 (Input B) meeting the C firing pattern of postsynaptic neurons (0110). Again, an association is formed with some additional switches turned to the on state (Fig. 6c). As this is a distributed associative matrix, both associative memories are stored within this single network, taking advantage of the input-specificity principle referred to earlier.

The last step by way of illustration is to retrieve information from the network. Suppose it is in a quiescent state and the retrieval cue A (0011) is applied to the matrix in its final state (Fig. 6d). Current is now switched from the horizontal wires to the vertical ones in proportion to the number of afferent fibers with activity meeting switches in the on state – giving rise to a dendritic input that corresponds to 2120. Clearly this is different from B. However, the difference is deceptive because, if this number is divided by the sum of the number of active afferents in the retrieval cue (2) using integer division, the result will be 1010 – which is stimulus B. Thus, provided this integer division operation can be performed, retrieval cue A will evoke a memory of stimulus B.

David Marr was the first to suggest that one function of inhibition within distributed associative networks was to perform this division function. He reasoned this on the basis that some inhibitory interneurons had dendritic branches that extended vertically throughout the region where an afferent pathway was active, and had axonal ramifications that could affect the firing of multiple principal cells. It follows that the relationship between the number of active afferents and the extent of feed-forward inhibition will be monotonic. This is an exciting but speculative idea.

Distributed associative memory systems are intriguing and have fascinating "emergent properties." Their analysis raises a number of questions about how biologically realistic they are of which the first is: *Where is the memory system?* Like the archetypal foreign tourist being taken around the English city of Cambridge, there eventually comes a moment when they ask "But where is the University" to which – of course – the answer is "everywhere" (Cambridge lacks a campus and is literally spread around the town). So also in a distributed memory system – the engram is instantiated by memory traces that are throughout the network. Second: *Does such a "model" system scale up when there are hundreds of neurons and thousands of synapses?* Yes, they do, and in fact they work much better when scaled up – there being much less opportunity for interference. Third: *What is the engram?* This is a subject of much current study – for neuroscientists want to know in precise

molecular detail the nature of the structural/biochemical changes at synapses and the molecular changes that mediate them. Central to this question is whether the changes last over time and so reflect the consolidation process. This is the final topic of this chapter.

The Structural and Functional Expression of Memory Traces

The Structural and Biochemical Changes at Synapses Associated with Long-Term Potentiation May Reflect the Various Stages of Memory Consolidation

The idea that memory traces might gradually stabilize over time is an old one, but the notion that it might be possible to understand that process in terms of structural and biochemical changes at synapses is gradually emerging as a tractable goal of contemporary neuroscience.

That a time period is required for memories to be consolidated was first proposed by the German scientists Müller and Pilzecker over a century ago. Retroactive inhibition paradigms in humans later provided evidence that new traces are subject to modification during the consolidation period. Parallel animal models also offered interventional evidence for causal links between brain and memory. In the middle of the last century, it was discovered that application of electroconvulsive shock (ECS) to rodents after training induced experimental amnesia. Experimental amnesia was also demonstrated by using drugs that inhibited protein synthesis. Control studies were conducted to rule out the possible confounding of physical sickness that might be induced by the drugs inhibiting protein synthesis, and otherwise validate its role in memory consolidation. This was an important advance because it suggested that cerebral protein synthesis was more than "mere" housekeeping; it was a specific biological process involved in the stabilization of new memories.

As discussed above, memory traces in the mammalian brain might be encoded as distributed patterns of synaptic weights within specific types of neural networks. But how might the phenomenon of synaptic consolidation be brought into the picture? A suitable physiological model of such changes, with changes in synaptic strength lasting for variable durations of time, is long-term potentiation. A protein synthesis-*independent* form of LTP, often called Early-LTP (E-LTP), lasts at most for 3–4 h. Protein synthesis-*dependent* late-LTP (L-LTP) lasts longer – both in vivo and in vitro. The reason why both forms of LTP evolved may be related to the reasons why both cellular and systems consolidation mechanisms are required for LTM formation. The difference between E-LTP and L-LTP also reflects a key difference between STM and LTM – that de novo protein synthesis is required for a short-lasting trace to be converted into a long-lasting one. It draws upon experimental work in *Drosophila*, *Aplysia*, early learning in birds, and mammalian memory.

Making Lasting Changes at Synapses Involves Them Being Tagged and the Tags Then Sequestering Plasticity-Related Proteins

An exciting new perspective on synaptic memory consolidation is the Synaptic Tagging and Capture (STC) hypothesis of memory trace formation. This hypothesis goes along with the standard view that plasticity-related proteins (PRPs) are critical for the persistence of synaptic memory traces, but argues that de novo synthesis of PRPs is not necessarily triggered by neural activity associated with the actual events to be remembered. New PRPs are still required, but their synthesis may be regulated in other ways and over a longer time-window. According to this idea, the creation of long-term memory traces is a *dual process*. In one step, the potential for a long-term memory is triggered by glutamatergic activation of NMDA and AMPA receptors and established locally at synapses in the form of rapidly decaying early LTP accompanied by the setting of a "synaptic tag." In the other step, a series of biochemical interactions including activation of various signal transduction pathways and protein–protein interactions, converts this synaptic potentiation into a stabilized trace at those synapses at which "tags" have been set. Tags capture PRPs, and thus the theory describing this cellular consolidation process has come to be called the "synaptic tagging and capture" (STC) theory. The events that lead to these interactions can be set in motion shortly before the event to be remembered, at the same time (as in most behavioral and in vitro brain slice experiments to date), or shortly afterward. This leads to the interesting idea that the persistence of memory does not have to be determined at the exact moment of initial memory encoding. A key finding that led up to this proposal was the discovery that the patterns of stimulation that normally trigger decaying E-LTP actually result in L-LTP if they occur within a short window of time of L-LTP being induced on independent synapses in the same population of neurons.

Thus, for the purposes of discovering the structural and biochemical changes that are thought to take place during memory formation, the following steps need to be considered:

1. Synaptic potentiation itself
2. The setting of a local synaptic tag
3. The synthesis and distribution of plasticity-related proteins (PRPs)
4. The capture of these molecules to replenish tagged synapses, *and*
5. The stabilization of synaptic strength

Taking these steps in turn, it has already been noted that *synaptic potentiation* involves an alteration in the number of AMPA receptors within the postsynaptic density (PSD) of individual dendritic spines. It may also involve changes in transmitter release. The *setting of a synaptic tag* is now thought likely to involve a temporary alteration in the actin cytoskeleton of a spine. This alteration of spine architecture is both permissive and necessary for the remodeling of the PSD that is

essential for the stabilization of LTP. The tag is unlikely to be a single molecule, and may even be the altered state (e.g., its phosphorylation) or several molecules. Some molecules required for the tagging of synapse architecture are now known to be specific to the direction of the synaptic change (CaMKII for potentiation; PP2B/Calcineurin for depression). However, there are still a number of candidate processes or molecules including protein degradation, scaffolding molecules such as Homer1a, and cell-adhesion molecules. Based on the current knowledge of molecular interactions, tagging may be best seen as a permissive "unlocking" process without which the novel synthesis and supply of PRPs is incapable of stabilizing plasticity. The *synthesis and distribution of plasticity-related proteins* likely involves both somatic and dendritic compartments. One key transcription factor is CREB whose activation is a key step in the synthesis of PRPs. Genetic manipulations of CREB in transgenic mice block (or enhance) the creation of long-term memory. Exciting new research indicates the dendritic localization of mRNAs whose translation may be regulated by activation of the MEK and mTOR signal transduction pathways. The tag-mediated *capture of plasticity-related proteins* is the next step in the stabilization of both the functional and structural alterations to a dendritic spine. The molecular identity of all the PRPs involved is unknown, but likely includes GluR1, homer1A, PKMzeta, and Arc. A recently outlined model suggests that additional Lisman "slots" (or scaffolding molecules) are inserted into an enlarged PSD (in association with complementary changes on the presynaptic side of the synaptic cleft). The number of AMPA receptors per slot settles back to the original level, but now with a greater number of slots. AMPA receptor trafficking into and out of these slots continues dynamically to realize a new and sustained state of potentiation that is characterized as L-LTP. Without these PRPs, the altered structural "scaffold" that constitutes the tagged state will gradually revert to an untagged state again as the activity of kinases responsible for synaptic tagging fades. The end result of a *stabilized state of synaptic strength* is an increase (L-LTP) or decrease (late-LTD) in the number of "slots" available for AMPA receptors realized by the remodeling of spine structure. AMPA receptors endocytose to and exocytose from these slots in this dynamic steady state; likewise glutamate transmitter molecules to their vesicular release sites on the presynaptic side. Figure 7 offers a graphical illustration of these sequential steps in synaptic potentiation taken from a recent model of synaptic tagging and capture.

Parallel computational models of STC also display these five steps, explain the induction of protein synthesis-dependent late-LTP (L-LTP) on a weakly tetanized pathway after strong tetanization of an independent pathway, and make novel predictions about the statistical variability of EPSPs after LTP induction. Such models will be particularly valuable in developing novel predictions about the "dynamic stability" of synapses involved in memory storage.

Consolidation Need Not Happen Only after Memory Encoding

The STC hypothesis represents a quite radical departure from previous thinking about memory consolidation. Whereas past models have asserted that consolidation

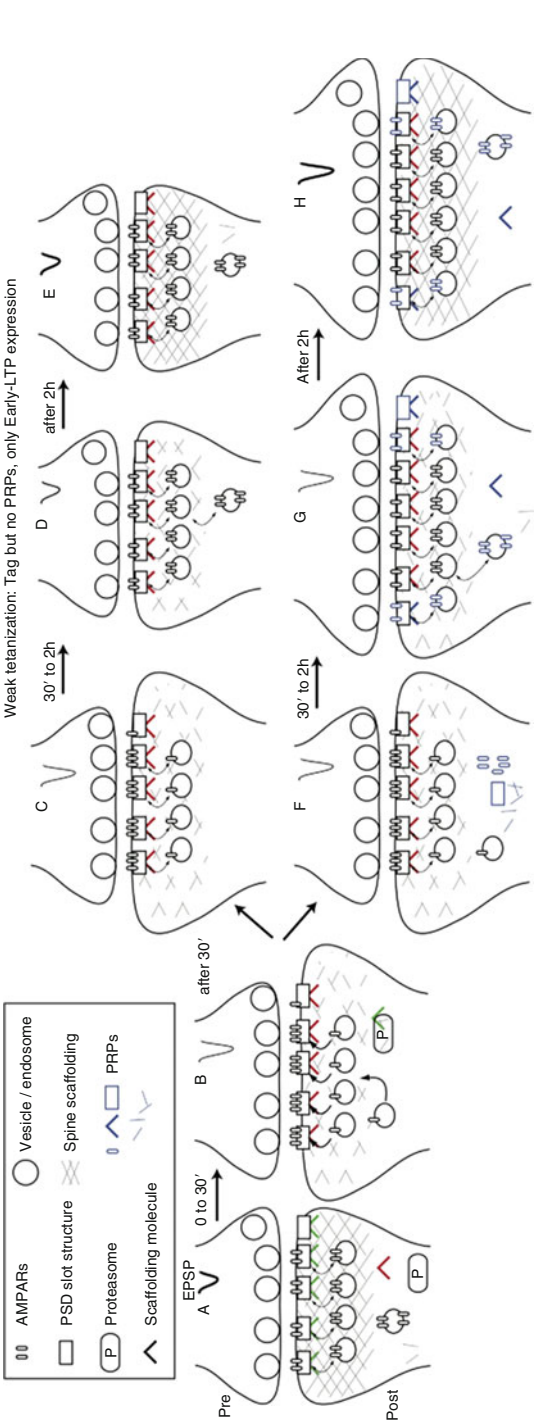

Fig. 7 *A molecular model of synaptic tagging and capture*. The model depicts the location and movement of candidate receptors, tag-molecules, and plasticity-related proteins in response to weak (E-LTP) and strong (L-LTP) inducing stimulation. (**a**, **b**) Additional AMPA receptors become embedded in slots in the postsynaptic density over the initial period of LTP induction and likely mediate the enhanced EPSP that is observed. In addition, a structural dis-assembly of the actin cytoskeleton occurs. (**c**, **d**, **e**) Following weak tetanization, the lack of PRPs results in the actin cytoskeleton reverting to its initially stable state and the enhanced number of AMPA receptors recedes as these are effectively removed by exo-cytosis. (**f**, **g**, **h**) Following strong tetanization, the supply of PRPs enables the altered actin cytoskeletal framework to be retained, additional slots are transported into the postsynaptic density, and the total number of AMPA receptors remains elevated. The potentiated synapse is sustained (Source: A black and white version of Fig. 4 from Redondo and Morris 2011)

is, by definition, a set of processes that must happen *after* memory encoding, the STC framework looks upon this as an oversimplification. This survey of the physiological and cell-biological mechanisms of memory formation reveal a myriad of ongoing processes that collectively achieve trace stability in the face of the constantly changing milieu of neuronal activity. Different components of "consolidation" may occur before or after memory encoding.

Specifically, if the prior history of the neuron has upregulated the availability of PRPs in a particular dendritic compartment, they may be rapidly captured by local synaptic tags, and so complete the stabilization of the synaptic component of a new memory trace soon after it is encoded. Critical components of this "consolidation" will have occurred before encoding. Conversely, if neural activity inducing PRP synthesis does not happen until some time after potentiation and tag setting, stabilization of the otherwise temporary synaptic potentiation will occur much later, with the temporal duration of the tag being the main parameter determining whether stabilization will occur at all. Thus, according to the STC hypothesis, the time course of consolidation is a malleable entity that varies as a function of what the neural network has been doing or may yet do in the near future. This theoretical framework is conceptually distinct from others that have previously asserted there is a "grace-period" or "window of time" *after* encoding when consolidation can be interrupted. The view taken here is that heterosynaptic events *before* and *after* encoding can also determine the fate of memory traces.

Outlook

This chapter has outlined a framework for thinking about the neurobiology of learning and memory that has included a journey from amnesic patients to the molecules that may be involved in forming synaptic memory traces. Along the way, it has attempted to explain that the psychological concept of "memory" reflects a family of different types, all of which involve a number of stages of their induction, expression, and stabilization. A particular focus was on three different subtypes of declarative memory called semantic, episodic, and spatial – with the example of spatial memory considered in most detail. It also discussed the manner in which a neuronal network can store associations given only the prerequisites of synaptic modification and a specific distributed neuronal architecture. This led on to asking questions about the molecular basis of the changes that occur at synapses within such parallel distributed networks.

With respect to the intriguing puzzle of memory consolidation, it was pointed out that some memories fade while others last. This is entirely desirable, because a memory system that retained everything would rapidly saturate to a point where information could not be retrieved. If systems consolidation is the set of processes that determine what memory traces last in neocortical networks, cellular consolidation can be thought of as the "low-pass" filter that determines the sub-set of information which can be subject to systems consolidation. The biophysical mechanisms of synaptic tagging and capture, mediated by interactions between structural

and functional changes at excitatory synapses associated with plasticity, provide a biologically beautiful way of extending the time course for which the induction of lasting memory traces is determined. This time-travel is both backward and forward in time. Specific biophysical devices, such as the NMDA receptor, help to instantiate the computational appropriate algorithms for inducing short or lasting changes in synaptic efficacy, with this and other molecular players engaged by the stimulus and behavioral events that occur during the everyday events that characterize animal and human life. A grand challenge for the future of the neuroscience of memory is to better understand the neural circuits and patterns of neural activity that intersect between the mechanisms of cellular and systems consolidation.

References

Buzsaki G (2004) Large scale recording of neuronal ensembles. Nat Neurosci 7:447

Corkin S, Amaral DG, Gilberto Gonzalez R, Johnson KA, Hyman BT (1997) H. M'.s medial temporal lobe lesion: findings from magnetic resonance imaging. J Neurosci 17:3964–3979

Deisseroth K (2011) Optogenetics. Nat Methods 8:26–29

Dudai Y (1996) Consolidation: fragility on the road to the engram. Neuron 17:367–370

Eichenbaum H, Fortin NJ (2009) The neurobiology of memory-based predictions. Philos Trans R Soc Lond B Biol Sci 364:1183–1191

Frey U, Morris RGM (1998) Synaptic tagging: implications for late maintenance of hippocampal long- term potentiation. Trends Neurosci 21:181–188

Govindarajan A, Israely I, Huang SY, Tonegawa S (2011) The dendritic branch is the preferred integrative unit for protein synthesis-dependent LTP. Neuron 69(1):134

Kandel ER (2001) The molecular biology of memory storage: a dialogue between genes and synapses. Science 294:1030–1038

Mishkin M, Suzuki WA, Gadian DG, Vargha_Khadem F (1997) Hierarchical organization of cognitive memory. Philos Trans R Soc Lond B Biol Sci 352:1461–1467

Morris RGM (2006) Elements of a neurobiological theory of hippocampal function: the role of synaptic plasticity, synaptic tagging and schemas. Eur J Neurosci 23:2829–2846

Moser E, Kropff E, Moser M (2008) Place cells, grid cells, and the brain's spatial representation system. Annu Rev Neurosci 31:69–89

Nakazawa K, McHugh TJ, Wilson MA, Tonegawa S (2004) NMDA receptors, place cells and hippocampal spatial memory. Nat Rev Neurosci 5:361–372

O'Keefe J, Nadel L (1978) The hippocampus as a cognitive map. Clarendon, Oxford

Redondo RL, Morris RGM (2011) Making memories last: the synaptic tagging and capture hypothesis. Nat Rev Neurosci 12:17–30

Squire LR, Stark CE, Clark RE (2004) The medial temporal lobe. Annu Rev Neurosci 27:279–306