Charles R. Legéndy*

On the 'data stirring' role of the dentate gyrus of the hippocampus

DOI 10.1515/revneuro-2016-0080
Received December 6, 2016; accepted February 18, 2017; previously published online April 7, 2017

Abstract: Understanding hippocampal (HC) function, as it is presently known, includes exploring the HC role in episodic memory storage. As pointed out by Teyler and DiScenna in the 1980s, the apparatus needed for recalling a stored episode, and awakening all its components in a coordinated manner, by necessity includes a triggering device able to reach each of the mental entities that must be awakened. In the context of neuronal networks, the triggering device in question takes the form of a large cell assembly, a separate one made for every new episode stored. The present paper deals with the creation and the properties of these cell assemblies ('pointer groups'). To perform the function of episodic memory retrieval, each of these must possess the information capacity (entropy) enabling it to single out an episode and the network connections enabling it to reach all components of it; further, to deal with the unpredictability of the memory items it has to address, it must have its member neurons well distributed through the length of the network (the HC). The requirements imply that the creation of a pointer group must include a randomizing step analogous to 'stirring'. It is argued that many of the known peculiarities of granule cells in the dentate gyrus arise as solutions to the practical problems presented by the creation of the pointer groups and the details of 'stirring', and so do a series of other features of the HC network, some of them only discovered in the last few years.

Keywords: cell assembly; dentate gyrus; hippocampus; random code; stirring.

Introduction

The role of the hippocampal (HC) complex in memory storage was first revealed to the world with the publication of the case of 'patient H.M.' (Scoville and Milner, 1957; Milner et al., 1968; Squire, 2009). Patient H.M. (Henry Molaison, 1926–2008) had had intractable epileptic seizures, dating back to a childhood bicycle accident, and had a large part of his medial temporal lobe surgically removed in the attempt to relieve his condition. His seizures became much less severe, and he retained his intelligence, his pleasant personality, and the memory of most events before surgery, but he could not remember new events, kept asking where he was, and often lost his way, even in his usual surroundings.

At first, it was not clear how much of this was due to loss of the HC itself, because the surgery of patient H.M. included several areas outside the HC (Squire, 2009), but subsequently a series of case studies became available where the damage was reliably confined to the HC (Zola-Morgan et al., 1986; Victor and Agamanolis, 1990; Vargha-Khadem et al., 1997), and through them, it became possible to isolate many of the deficits attributable to HC loss. In general, three statements emerged as valid for HC-dependent memories.

First, it was found that only memories accessible to conscious recall were affected. Language skills seemed unaffected; in fact, two of the patients of Vargha-Khadem et al. (1997), whose HC damage occurred around the time of birth, were able to learn to speak, read, and write, and were able to go to mainstream schools. It was found, in addition, that a long list of memory tasks along the lines of skill learning and conditioning were immune to HC damage. The prominent role of consciously recallable memories led to characterizing the HC-dependent part of memory as declarative (or explicit) memory (Squire and Zola-Morgan, 1991; Eichenbaum, 2001).

Second, the memories affected were ones dealing with time-confined events and sequences (episodes). This led to the description of HC-dependent memory as episodic memory (Tulving, 1972; Tulving and Markowitsch, 1998). The latter term emphasizes the ability of the HC to grasp information quickly without requiring a second chance. The 'episodic' designation (adopted in the present paper) assigns a special role to memories acquired in the course of one-trial learning, which in this case is understood to include sessions lasting many minutes, as when the subject 'plays' with an object for a while, typically in multicomponent sensorimotor interaction with it.

Third, it was found that HC damage mainly affected the recording of new memories – those memories stored

***Corresponding author: Charles R. Legéndy,** Department of Psychology, Columbia University, New York, NY 10027, USA, e-mail: crl2122@columbia.edu

before the damage seemed largely unaffected. This led to the conclusion that after a while memories underwent consolidation and were passed down from the HC to the cortex, which could then retrieve them without help from the HC. (For simplicity, the neocortex will be referred to as the 'cortex' below).

It must be pointed out that, on the general subject of 'what the HC does', there are actually two schools of thought – the 'other' one insists that the HC does not deal with memory but with spatial location. Of course, the evidence of a HC role in memory is ironclad, but the fact still remains that when the HC pyramidal cells (PCs) are studied individually they are all found to be 'place cells' (O'Keefe and Dostrovsky, 1971; O'Keefe and Nadel, 1978; O'Keefe, 1999), and that HC patients, along with their memory problems, invariably show a profound deficit in spatial orientation. The present paper offers a resolution to the paradox and points out that the HC role in episodic memory, as described below, is in fact compatible with the role in spatial localization (see Discussion).

It will be noted that acquiring the knowledge of current location satisfies the 'one-trial' and 'episodic' criterion, because it requires retaining the results of a time-limited effort of studying a specific set of navigational cues along with noting the sensorimotor sequence leading up to the current position; and this is a one-time affair, even when it is identical to similar earlier evaluations at the same location (Nadel and Moscovitch, 1997). The locations at the other earlier times are of no interest; the objective is limited to localization at the present unique moment.

The special position of single-session learning in HC function was dramatically highlighted during the attempts to develop an animal model of HC-dependent memory loss. When the original operation done on patient H.M. was repeated on monkeys (Correll and Scoville, 1965), in the hopes of reproducing the same deficits, the experiments failed – the monkeys were apparently able to learn what patient H.M. was not. The explanation, as we now know, was that the early monkey experiments sought to test memory via the standard learning tasks using simple conditioning that involved hundreds of trials, and these depend on the basal ganglia rather than the HC (Squire, 2009); they make use of 'nondeclarative' memory. Subsequently, success was achieved in tests using single encounters that involved the voluntary exploration of objects (Mishkin, 1978).

The successful experiments made use of the 'delayed non-match-to-sample learning' paradigm (Berlyne, 1950; Bevins and Besheer, 2006), which uses the fact that the novelty of an object increases the length of time spent exploring it, and when an already familiar object is, in a subsequent session, placed next to an object not previously explored, the familiar object is largely ignored in favor of the new object. When, after HC trauma, an object is not ignored under such circumstances, it is justifiable to conclude that the machinery of episodic memory storage has been damaged. In fact, in animals, the non-match-to-sample task appears to be the test of choice for detecting strictly localized HC damage. Other memory-sensitive tasks, not relying on single-encounter learning, have also been described and found successful when HC damage extended to neighboring structures (Zola-Morgan et al., 1989) but not when the HC alone was disabled (Squire and Zola-Morgan, 1991).

Turning next to the aims of the present paper, the first step is to point out that when a piece of episodic memory is stored the manner of storage must ensure that the memory can also be retrieved afterwards when needed. There are many episodes stored in the course of a lifetime, and the act of accessing any one of them involves retrieving its specific set of details, together, and in their original form of conjunction. By the nature of episodic memory, the stored form is determined at the time of storage; it is not a continually evolving affair as in a skill. This in turn implies the presence of a fixed retrieval-initiating device for each episode to dictate the time when the coordinated recall should happen.

The device, to perform its function, must have the selectivity (entropy) enabling it to single out the episode recalled and must have network properties enabling it both to reach all components of the episode and to be reached by all potential clues indicating when it should be invoked. These requirements suggest that the retrieval-initiating device be in the form of a large ignitable neuron group ('pointer group'), the kind meeting the definition of a cell assembly (Hebb, 1949), playing a role along the lines described by Teyler and DiScenna (1986), Milner (1989), Nadel and Moscovitch (1997), and Buzsáki (2006).

The version of cell assemblies used here (Legéndy, 1967, 2009, 2016; Scott, 1977; Palm, 1981; Wickelgren, 1999; Buzsáki, 2010) describes them as randomly selected sets of neurons, linked together only by having the synaptic contacts between them reinforced. The sets are allowed to overlap and, in the present case, extend over entire cortical or HC areas. Their connectivity permits cell assemblies to 'ignite' (Rapoport, 1952; Legéndy, 1967, 2009; Palm, 1981, 1982; Wickelgren, 1999), in the sense that it enables firing to spread rapidly within a cell assembly to all of its members and, when the system is properly configured, not to spread to the other cells in the area.

In the case of a pointer group, the requirements of entropy and reaching are optimally satisfied when the

group members are randomly and uniformly distributed through their area. It will be argued below that the random and uniform distribution is essential to function, so that when the HC creates a new pointer group it must perform what amounts to active 'stirring' of the population to spread out and randomize the cell assembly membership. The role in creating uniformly spread-out neuron groups matches the anatomy of the HC, which, as Buzsáki (2006) put it, 'is a giant cortical module'.

The developing cell assembly itself is built in CA3, and in its final form, after some iterations of randomization, it assumes an ignitable configuration with its members linked together via CA3-to-CA3 synapses. During the weeks or months afterwards, the new group is exported (via CA1) to the cortex, where local representative groups are built in multiple cortical areas, and these continue to act as a pointers to local pieces of the same episodic memory.

What the 'stirred' configuration accomplishes is, among other things, that the cell assemblies can be recognized by samples of their output taken from any arbitrary location. In the HC complex, this enables the cell assemblies representing the details of the episode to recognize the pointer group regardless of their location, and at the cortical level (as will be argued), the same mechanism accounts for the ability of visual areas to 'ignore' the location of the retinal (and cortical) image of an object. Another corollary is the smooth interplay between many sense modalities and their corresponding motor sequences, although they are implemented in separate and often distant cortical regions. [The present paper uses the term 'HC' to include the dentate gyrus (DG), the HC proper, and the subiculum, and the term 'HC complex' to include the HC, the amygdala, the entorhinal cortex (EC), the perirhinal cortex, and the parahippocampal cortex].

In support of the picture of memory storage outlined, a series of known HC features will be listed below and reinterpreted as solutions to various challenges presented by creating the randomized cell assemblies that are to serve as pointer groups. Granule cells (GCs) are not interconnected. Their synapses are not NMDA dependent. Their synaptic linkages are strong enough to enable single GCs, without help from others, to fire their postsynaptic cells. Long-term potentiation (LTP) in these synapses is presynaptically determined and can be induced by single presynaptic GCs. There is a significant turnover of the GC population, continuing into adulthood, with neurogenesis reaching a level of 1% per day. And the cells spreading the GC outputs through the DG, the mossy cells (MCs), spread them out onto most of the longitudinal extent of the DG but leave out the first 1–2 mm near the cell bodies. The randomizing 'stirring' action to be described also appears to be served by the incendiary properties of the semilunar GCs (SGCs), the effect of theta activity on the firing of GCs, the connectivity of the hilus and of CA3, and the back-projection from CA3 to the DG.

The data stirring hypothesis

Random neuron groups in the role of 'code words' for mental entities

Before turning to details specific to the HC, it is desirable to say a few words about the role assigned to cell assemblies in the present paper (Legéndy, 2009, 2016).

It is said that fingerprints are unique and that no two people have the same fingerprint. But in fact fingerprints are completely random; their uniqueness is merely probabilistic. The source of the 'uniqueness' is that there are so many degrees of freedom in the way ridge patterns can form on fingers that there are many more distinguishable fingerprints than there are people.

The concept of a 'random code', in its rudimentary form, may be illustrated by the example of a collection ('ensemble') of N sequences of binary bits ('code words'), each of them random and n bits long, obtained by flipping a coin n times. [The general concept goes under the heading of Shannon's random code ensemble (Shannon, 1957; Mézard and Montanari, 2009, p. 107)]. It is noted that the number of n-bit binary numbers is $N = 2^n$, which is the same as saying that the smallest possible nonrandom word length for an ensemble of N code words, where all the bits hand-picked, is $n = \log_2 N$.

It is easy to show that, when the digits are not hand-picked but are random, one only needs to double the number of binary bits per code word compared to the smallest possible word length (and make code words of $2n$ bits rather than n bits) to keep the probability below 0.5 that no two code words are the same; and further that each added bit doubles the safety factor against coincidence. For instance, if a million code words are formed, with each one made of 60 random binary bits, then (noting that 1,000,000 is approximately 2^{20}) the probability that any two turn out to be the same will be less than 1/2,000,000.

The use of randomly generated patterns to identify entries in databases, in a role where ordinary serial numbers would be the more obvious choice, is quite common in large systems (Hamming, 1950), particularly

in situations where serial numbers (or other deterministic identifiers) cannot be used, because there is no way to check, before assigning an ID number to an item, whether it has already been assigned to some other item.

Mental entities within the brain are an example, because mental entities tend to arise from different and unrelated sources, and there is no way to compare notes between the sources. Accordingly, it is proposed here (on a number of grounds; Legéndy, 2009, 2016) that cell assemblies play a role of giant 'code words' in the brain, acting as the unique pointers by which individual mental entities (for instance, episodic memories) can be retrieved on demand. From an abstract point of view, cell assemblies form a special kind of Shannon-style random code ensemble – one that may be called the 'super-redundant random ensemble' (Legéndy, 2016), where the code words are so large that small pieces of them are already fairly secure against duplicate assignment (in the same way as are partial fingerprints, when large enough).

In the aforementioned binary example, if each of the code words were made up of a thousand random bits instead of 60, then arbitrary 60-bit segments taken from the larger code words, taken at a fixed locations assigned on the 1000-bit sequences, would also form a reliable ensemble of a million different code words, and the probability that any two would turn out to be the same would still be about 1/2,000,000.

The power of the super-redundant ensembles is illustrated in Figure 1, which shows how a piece of dendrite can, under proper conditions, receive enough information from an igniting cell assembly to support recognition of the cell assembly.

However, there is one catch: recognition based on samples can only be reliable when the samples are representative. In the present case, this means that (ideally) the samples must contain enough elements from every cell assembly of the ensemble to make them recognizable, which in turn means that the samples, wherever they are taken, must contain enough information (entropy) to permit recognition to an acceptable level of certainty.

This is where the importance of the 'data stirring' enters the picture. If a cell assembly is not well 'stirred', its member neurons and their processes may be concentrated near a few locations within the area and then other locations would not get any of them; the cell assembly would be invisible at those other locations.

The idea that a small sample from a cell assembly's output can be unique enough to permit recognition is illustrated in Figure 1. It shows three copies of a segment of dendrite, each drawn to carry about 160 viable synapses, with about 5% of them receiving input from an igniting cell assembly (Losonczy and Magee, 2006), for instance, 8 synapses in the case of the 'pliers'. The thing to note is that the number of ways in which 8 synapses can be selected from among 160 is about 10^{13}, which means that the small piece of dendrite shown receives enough information that the cell assembly can, in principle, be recognized from it. [The issue of using the combinatorial diversity of synapse choices in neuronal response selection is addressed elsewhere (Legéndy, 2016)].

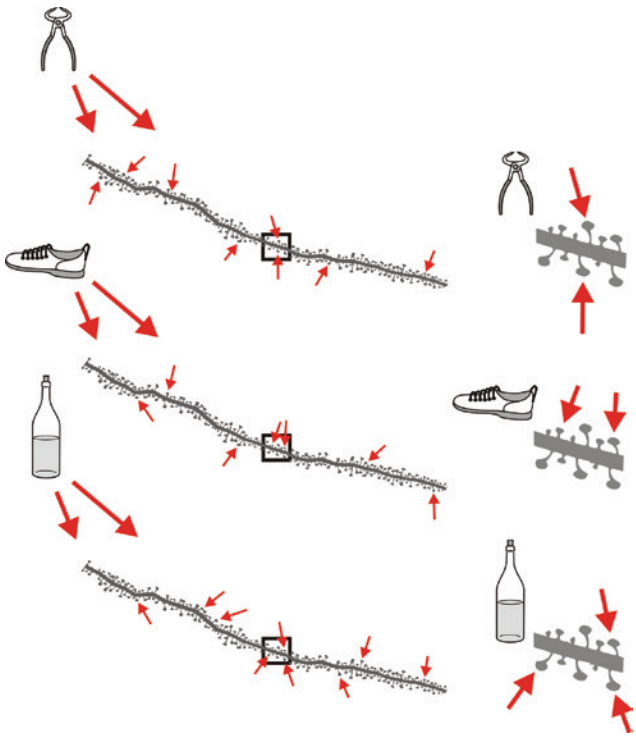

Figure 1: Recognition of cell assemblies based on small samples of their output.
Area-wide signal fields (igniting cell assemblies) send multisynaptic input to a piece of dendrite. The icons of 'pliers', 'shoe', and 'bottle' are chosen to be examples of objects participating in episodes for which the HC may form 'pointer groups', for instance when these are the objects with which the experimental subject in a 'delayed non-match-to-sample' session is becoming familiar. (The object shown in each drawing is only one of the elements of the episode designated by the 'pointer group' in question, and it is only used here to denote the episode as a whole for the convenience of illustration. It is noted that the same object can appear in many episodes). A segment of a dendrite, shown here as having about 160 viable synapses on it, is reproduced in three copies, each receiving a volley of spikes arriving to a subset of its synapses (arrows) from presynaptic elements (not shown). Modified from Legéndy (2016).
Note: It will be appreciated that episodes centered on simple tangible objects only constitute a small subclass of all possible themes of episodic storage. They are chosen here because they are especially well suited for illustration; and, significantly, non-match-to-sample learning is verifiably a function of the HC.

Two non-obvious feats of brain processing attributable to the 'stirring'

One notable aspect of brain function is the impeccable coordination between different sensory and motor modalities, undisturbed by the fact that the data streams of different modalities travel over separate pathways and are processed in separate cortical areas. During everyday sensorimotor activity, a sensory quest brought up through one sense organ is often pursued, without delay, by motor action that serves other sense organs. It is hard to escape the impression that the different sensory and motor components are subject to unified central control — and yet a central controlling organ does not always exist.

It does exist in the context of a HC-centered task, like the initial exploration of a new object in a non-match-to-sample experiment, as the HC complex is known to communicate, directly or indirectly, with all the cortical areas serving the sensors and the muscles. Accordingly, it is not unexpected, for instance, that during the initial familiarization with an experimental object, which often consists of 'playing' with the object, the sensorimotor streams should be well coordinated. But what is astonishing is that some weeks or months later, when the memory is successfully consolidated in the cortex and the HC does not necessarily play a role any longer, the coordination between the diverse data streams remains as perfect as before.

The solution to the enigma, as proposed here, centers on the 'pointer group' and the way it is reproduced in the participating cortical areas. One corollary of the super-redundancy principle is that a pointer group can be cut into pieces and each piece can independently pinpoint the same stored episode, which is what happens when a pointer group is formed in the HC and is subsequently recreated in the EC (via CA1). The recently formed HC cell assembly is effectively cut into pieces, with each piece communicating its contents to a different anatomically distinct portion of the EC (Suzuki and Amaral, 1994; Kerr et al., 2007), selected essentially by convenience of connectivity.

When the pieces are subsequently passed further down into the cortex and repeatedly cut into further pieces, the resulting local versions of the pointer group will continue to be stable and to correctly pinpoint (with probability close to 1.0) the same piece of episodic memory. By

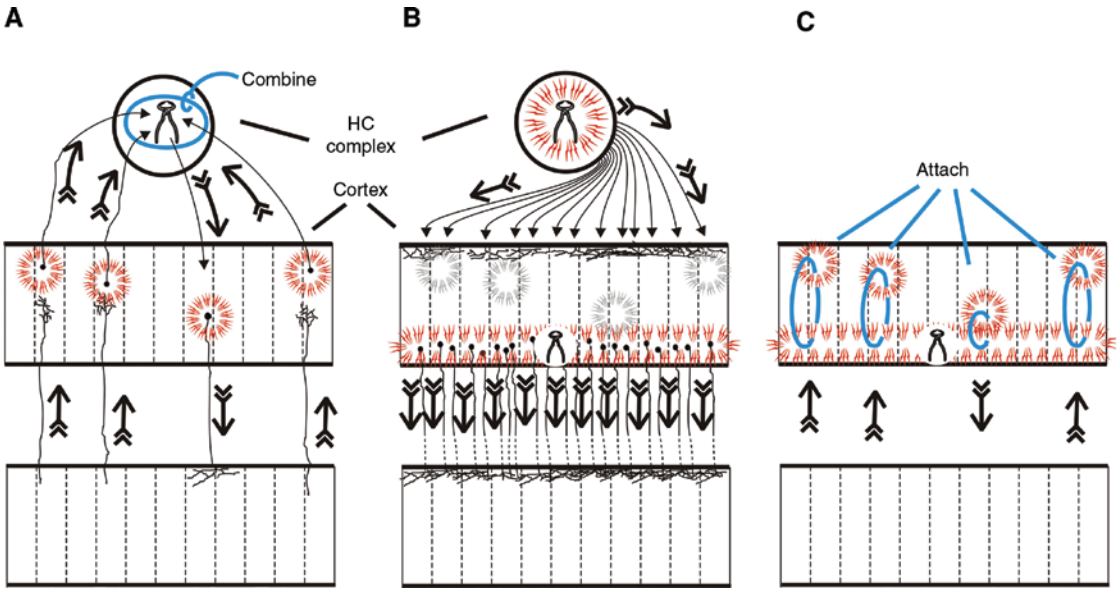

Figure 2: Hippocampal recording and cortical consolidation of an episode.
Idealized sketch of data flow between the cortex (shown as flattened) and the HC complex. The component ignitions of the recorded episode, actually occurring in separate cortical areas, are shown side-by-side for simplicity; also, the 'pointer group' ('pliers' icon), actually envisioned as recreated in many distinct cortical areas, is shown as if continuous. The data flow includes sensory as well as motor components (upward and downward arrows), as expected in motor-assisted sensory exploration. (A) Several components of an episode in the cortex make their way to and from the HC complex. In the HC, a pointer group is created for the episode ('pliers') and subsequently linked up (circling in blue) to the locally accessible cell assemblies whose ignitions represent the components of the episode. (B) Repeated ignitions of the pointer group in the HC establish corresponding pointer groups in the cortex. Like the HC version, they are spread out horizontally through entire cortical areas. (C) Consolidation: Ignitions in the cortex, initiated from the HC complex, create locally detectable coactivity between the newly created cortical pointer groups and the local component assemblies, which had originally transmitted the episode components to and from the HC. The temporal linkage is used, via Hebbian plasticity, to create synaptic linkages in each of the cortical areas involved (loops in blue, forming 'links' to the igniting groups).

pinpointing it, each local pointer assembly is in a position to place its own cortical area into an operating mode corresponding with the referenced episode, independently of the other cortical areas, and invoke the internal structures to govern the coordination.

The cortical processes making this possible are initiated at the time of the original episodic recording, when the cortical components contributing the episode, carried to and from the HC complex (Figure 2A), are active and are visible to local mechanisms tasked with marking them as being potentially of interest. As the decision is reached to record the episode, and a pointer group is created for it in the HC, the cortical regions containing the marked components receive local clones of the new pointer group (Figure 2B). Subsequently, the same components are linked to the new local pointer groups through Hebbian training (Figure 2C).

The coordination of diverse sensory and motor modalities is the first of two non-obvious feats attributable to the randomization of cell assemblies; the second one is the

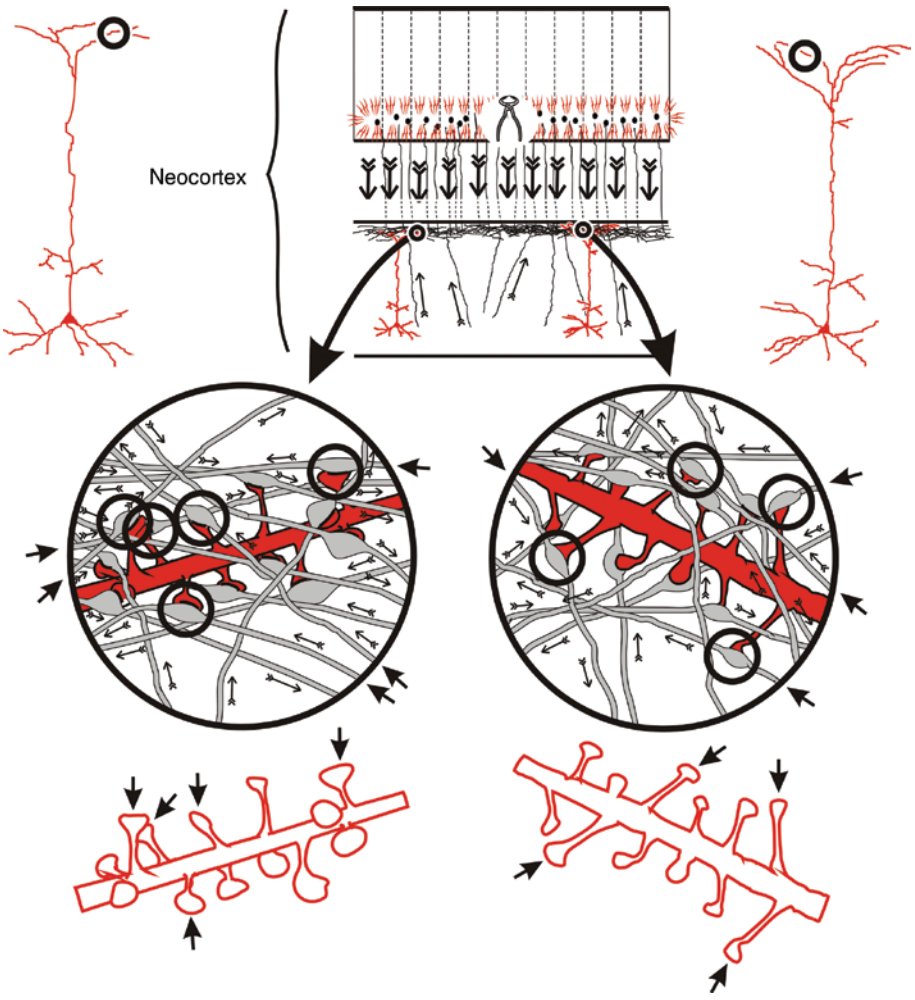

Figure 3: The 'pliers' pointer group is recognizable at both places.
When an earlier-recorded episode is recalled, each component of the episode must be awakened by a common signal (ignition of a pointer group) specific to the episode. The components themselves are represented by cell assemblies, and awakening them implies that the common signal must be recognizable to each of the cell assemblies. In fact, it must be recognizable by their individual neurons, wherever they happen to be in their respective areas (see Discussion). (A) The local pointer group of an episode ('pliers') ignites in one cortical area, and its volleys are received in another cortical area, below it. They arrive to the distal processes of the apical dendrites, where they can induce dendritic spikes when recognized. In the lower cortical area, two PCs are shown (red, magnified in separate drawings at the sides). A short segment of an apical dendritic branch from each PC is circled and shown in insets. In each inset, the piece of dendrite of interest (red) is surrounded by a network of axons, some of them synapsing on it. Arrows at the periphery point to the axons, which bring input from the 'pliers' (at the synapses circled). Separate sketches (bottom) show the dendrite segments with the surrounding axons removed and arrows pointing to the synapses receiving input (comparable to the arrows in Figure 1).

ability of cell assemblies to exert their specific effect on local cortical circuits regardless of cortical location. It is what (indirectly) enables the visual apparatus to 'ignore' the retinal location of objects seen and in fact disables the perception of retinal location altogether. The location-neutrality feature, which is not confined to vision but extends to all modalities, hinges on the fact that the cell assemblies can be recognized at any locality within the same cortical area.

Location-neutral processing is a corollary of the property of super-redundant random cell assemblies that they can be recognized (with a probability close to 1.0) based on relatively small samples of their firing output, as pointed out in connection with Figure 1. The cells that recognize the pointer group in this context (Figure 3) are members of the local 'episode component' cell assemblies, which, by responding to the pointer assembly of an episode, enable whatever local modes of activity are consistent with revisiting the episode. In the case of visual areas, this means the activation of visual components of the underlying object and of the spatial relations between them (Legéndy, 2009).

Figures 2 and 3, incidentally, suggest a way to 'explain' the peculiar architecture of the cortex and the dominant role of PCs in it. The PCs are enabled by their geometry to receive rich input from single layers of the cortex, remote from the cell body, and particularly from layer 1, the layer most richly innervated from higher cortical areas.

Although distant from the cell bodies, these inputs, through their ability to initiate dendritic spikes when recognized (Svoboda et al., 1997; Kamondi et al., 1998; Gasparini et al., 2004; Losonczy and Magee, 2006), and thanks to the robust conductive properties of the apical dendrites, can readily travel down to the cell bodies and influence the output. As a result, widespread cell assemblies in higher areas can send volleys to layer 1 of the areas below them, where they can be recognized anywhere on the cortical sheet and make themselves felt; this means that the PCs can pull together multiple cortical areas into coordinated platforms of mental effort.

Heuristic remarks suggesting that the DG is suitable for 'data stirring'

The first indication that the DG might perform more than its often-stated 'preprocessing' function came from the discovery of the curious synaptic targeting of the mossy fiber axons in the inner molecular layer (IML; Amaral and Witter, 1989). For no apparent reason, these axons, reaching extremely far along the longitudinal HC axis, land their influence preferentially at places a certain distance away from the laminar locations containing the cell bodies of the originating MCs and create (in the rat) a roughly 2 mm gap with almost no synaptic terminals (Figure 4). The MCs catapult their excitatory influence to places away from their source of excitation, where they then add their firing to whatever activity exists at the new locations.

At the new localities, as far as can be assessed from the data (see below), the MC volleys initiate parallel excitatory and inhibitory barrages in the hilus, briefly competing for control of the GCs and adding a large random component to the mechanisms causing some of the GCs to emit intense volleys of spikes (see below). The occasionally arising volleys impinge on a new rank of MCs and give rise to a wave of MC firing starting from the new locations and taking the excitation to yet another set of new locations, again some distance away, in a self-repeating iterative process.

Vague analogies come to mind. A deck of cards can be shuffled by splitting the deck, moving one part a little distance away from its original place, and dropping its cards against the cards in the new place, allowing them to arbitrarily interleave where they happen to fall, and then repeating the process a few times. Milk is sometimes stirred into coffee by pushing a spoon through it, with its face at right angles to its motion, so that the liquid escapes at a relatively high speed sideways from the moving surface, sending eddies deep into the surrounding medium at the two sides, mixing the newly arriving fluid

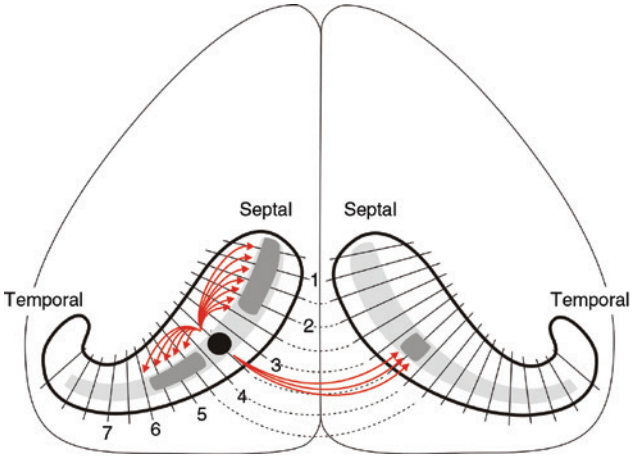

Figure 4: MC associational output leaves a 2 mm gap near its origin.
One of the phaseolus vulgaris leukoagglutinin (PHA-L) experiments described by Amaral and Witter (1989), sketched with septotemporal distance markers, referring to the rat brain. The part of interest for the present purposes is the ipsilateral projection.

Figure 5: Spread of the randomized region through repeated theta cycles.
Through the action of the ipsilateral associational projection of MC axons and the events initiated by them (Figure 6), the region containing randomly scattered tagged CA3 PCs (darker gray) spreads by large leaps to ever-newer locations until the process covers the whole HC length several times over, and the tagged PCs become well distributed.

with the stationary pool, and then, once again, adding a few iterations of the same action.

The analogies, of course, do not have persuasive value, but they are mentioned here to call attention to four elements shared by the mechanisms described below: (1) sending out unmixed samples of a population to places some distance away from where they originally were; (2) making the samples impact upon the surrounding matter in an unpredictable way, thereby creating random results at the new places; (3) repeating the process multiple times; and (4) permitting a strong overlap between the regions affected by the different iterations (Figure 5), thereby letting the process repeatedly revisit all locations.

Solitary granule cells can mark CA3 cells; they do not need to be synchronized

The GCs of the DG are special. First, the mossy terminals, delivering their output to the CA3 PCs and to the MCs (Chicurel and Harris, 1992), are so powerful that a strong burst from a single GC can reliably fire the CA3 PCs as well as the MCs postsynaptic to it (Acsády et al., 1998; Henze et al., 2002). This means that GCs can exert their effect independently of each other; it is not necessary for them to fire in concert with other GCs to fire the CA3 PCs and the MCs, as long their firing is in the form of sufficiently intense bursts. [One GC sends its output to about 20 CA3 PCs (Buzsáki, 2006) and to about half that many MCs (Ribak et al., 1985; Acsády et al., 1998)].

Equally significant, the GCs can induce LTP on the synapses they make on the CA3 PCs and MCs, without having to coordinate their firing with other GCs; because LTP in the mossy terminals is NMDAR independent (Harris and Cotman, 1986; Lysetskiy et al., 2005) and is presynaptically determined (Langdon et al., 1995). Solitary GCs can induce LTP merely by firing a number of times in quick succession.

The fact that GCs do not require the help of other GCs to initiate LTP is consistent with the connectivity of GCs, which does not support cell-assembly-like synchronization between them: there are essentially no direct granule-to-granule synaptic contacts. The linkage between GCs (except in seizure-related situations; Tauck and Nadler, 1985; Sutula et al., 1989) is indirect; it operates through MCs, SGCs (Williams et al., 2007), and interneurons. This means that the usual synchronizing mechanism, active throughout the cortex and the HC proper, and known to be especially prominent in CA3, is absent in the DG. When there is synchrony between GCs, it tends to come either from the effect of EEG or from synchrony present in the input from the EC.

As seen, the GCs themselves cannot form ignitable cell groups; their role in the scheme is to determine the cell selection for the CA3 cell assemblies by selecting (via burst firing) those CA3 cells that are to become members. The thing to note here is that in this role it is advantageous for the GCs to be free to assume as many diverse burst-emitting patterns as possible and thereby to be free to imprint an equally large number of diverse cell assemblies on the CA3. In view of such considerations, it will be appreciated that the independent and self-sufficient ways of GCs are highly desirable, as they increase the diversity of possible imprintable patterns. If GCs were constrained to act in synchrony, and if CA3 cells required correlated GC input to undergo LTP, the constraints would narrow down the choice of selectable combinations.

In the imprinting protocols described below, shorter-lasting LTP of GC-to-CA3 contacts, resulting from GC bursts, is followed by more permanent Hebbian LTP of the CA3-to-CA3 synapses between pairs of CA3 PCs. The latter arises because both members of such pairs give off repeated firing in the process of being potentiated, and they both fire within the same portion of a theta cycle (see below). In the process, after some iterations, the reinforced CA3-to-CA3 synapses combine the coactive cells into what becomes the pointer group.

Clustering of GC firing during theta activity

The GCs of the DG are relatively silent most of the time (Jung and McNaughton, 1993; Chawla et al., 2005), and even during times of elevated firing rate, for instance inside their place fields (Leutgeb et al., 2007), they often emit single spikes rather than bursts. However, the desirable effects of GCs described in the last section arise as a consequence of intense bursts emitted by the GCs; therefore, it is of interest to explore the conditions under which these occur.

The present paper concentrates on the situation as it exists during theta activity; the reason being that burst firing by GCs appears to be encouraged during theta (Buzsáki et al., 1983). In general, (most of) GCs are considered to be 'theta cells' (Bland et al., 1980; Buzsáki et al., 1983; Rose et al., 1983) in the sense that theta activity is demonstrably linked to increased intensity of GC firing. It has also been found that the induction of LTP is phase-linked to the theta waves (Pavlides et al., 1988).

In generating the burst events themselves, input from the medial septum and the diagonal band of Broca (MS-DBB) appears to play a crucial role. The MS-DBB input is known to have a driving effect on theta waves throughout the HC (Petsche et al., 1962; Donovick, 1968; Tóth et al., 1993; Buzsáki, 2002) through widespread excitatory (cholinergic) and inhibitory (GABAergic) input (Freund and Antal, 1988; Buzsáki, 1984; Hangya et al., 2009). The MS-DBB input is not the sole controller of theta activity, as the control of theta is strongly influenced by the local current generators and by multiple interacting theta generators, including the CA3 network itself (Buzsáki, 2002). All the same, the controlling influence of the MS-DBB is universally present, and it is probably safe to say that its excitatory and disinhibitory volleys, directly and indirectly, account for the profound change in GC firing during theta activity.

The excitatory component is relatively wide in its time duration, whereas the inhibitory volleys are relatively sharp (Vertes and Kocsis, 1997). The longer-lasting excitatory bombardment has a strong synaptic component on the inhibitory interneurons, and through them it will, for a while, suppress the excitatory effect of the direct MS-DBB projection onto the GCs. Then, upon arrival of the sharp inhibitory volley from the MS-DBB (which only acts upon the inhibitory interneurons; Freund and Antal, 1988), the inhibitory cells are suddenly silenced, and their silencing, helped by the accumulated effect of excitatory input to the GCs, gives rise to significant postinhibitory rebound (Cobb et al., 1995; Paulsen and Moser, 1998), which encourages burst production in GCs.

The postinhibitory rebound tends to cause bursts where its effect is allowed to prevail, but there is also a second mechanism present, in parallel, which disturbs the underlying GC inputs. It is the action of a little-studied population of cells in the IML, named SGCs by Cajal (1893, 1995), which appear to contribute a strong irregular component to the GC input.

As described by Williams et al. (2007), the SGCs exert intense excitatory influence in the hilus both on the MCs and on the interneurons. In fact, the SGCs appear to have an incendiary effect, on full display in slice preparations, where the theta waves and MS-DBB input are absent. Under these conditions, perforant path stimulation causes intense barrages of firing throughout the hilus both by the MCs and by the interneurons (Larimer and Strowbridge, 2010). The barrages, lasting roughly 10 s, cause the GCs to be simultaneously bombarded with competing excitatory and inhibitory input. (In the slice preparation the inhibition wins out and silences the GCs during the hilar 'up states').

During theta activity, the SGC effect is truncated by the time discipline imposed by the theta, which periodically silences most firing, but the strong excitatory synapses from the SGCs are still present, and when the SGCs emit volleys (see below), those will still send strong excitation both to the MCs and to the hilar interneurons, which, in the brief periods permitted, send mutually opposing inputs to the GCs.

The result is a biological version of what in electronics would be called a 'race condition' – something avoided by engineers because it leads to unreliable (or partly unreliable) outcomes. Although in the DG the 'race' only extends to a small portion of each theta cycle, its brief appearance is of interest here because it has a randomizing effect.

As noted, the generation of bursts is encouraged by postinhibitory rebound, which is highly sensitive to the precise time relations between excitation and inhibition. The crude impact of competing input volleys is enough to ruin the timing in some places, while leaving it or improving it in others, in an unpredictable way, and by doing so, it adds a considerable variance to the distribution function of GC burst events. This, in turn, contributes to the random outcomes at the impact edge of the 'stirring' action, as described in the next section.

Step-by-step spread of the nascent pointer group to the whole length of the HC

The mossy terminals at the output end of GCs do not only contact the CA3 PCs but also the hilar MCs, and their MC synapses have the same intense effect as their synapses on the CA3 PCs (Lysetskiy et al., 2005). Accordingly, the

subpopulation of GCs emitting bursts intense enough to fire and potentiate their postsynaptic CA3 PCs (in what follows, they will be called the 'hot' GCs for short) will also fire and potentiate the MCs they contact.

The MCs, in turn, through their longitudinal associational fibers, project to far-away locations (Amaral and Witter, 1989) and send on the excitation from the hot GCs to those locations.

At the distal end, the MCs make synapses in the IML (Amaral and Witter, 1989), which is the site of the SGC cells, and in fact, the MCs make synapses on the SGCs (Williams et al., 2007). The MC volleys, coming from distant hot GCs, are expected to induce corresponding volleys by the SGCs at the new locations. As noted, the SGCs have an incendiary effect; their volleys send strong excitation into the hilus, both to the local MC and to the local interneurons, and tend to induce the aforementioned 'race condition', adding randomness to the theta-synchronized bursting of GCs there.

In this way, the activity arriving from distant locations impacts on the local GC activity in an unpredictable manner (providing the step analogous to the local mixing in each iteration of the 'coffee stirring' or the 'card shuffling'). The randomly arising hot GCs then give rise to new volleys by local MCs; thus, they initiate another iteration in the randomizing action by sending their volleys to yet-another set of distant SGCs (Figure 6), and the process continues until it has covered the whole HC length several times over (Figure 5), at which point other influences stop it (see below).

The spreading described so far only relies on mechanisms inside the DG, but because of the dual output of the GCs the excitation also spreads out to the CA3 region (Figure 6). The hot GCs fire and potentiate the CA3 PCs they reach ('tag' them, as can be said), causing the set of randomly tagged CA3 cells eventually to extend over the whole length of the HC (Figure 5).

Through the tagging, the events inside the DG encourage the creation of a cell assembly in CA3. The CA3-to-CA3 synapses are NMDAR dependent (Nakazawa et al., 2002); therefore, properly timed coactivity will tend to

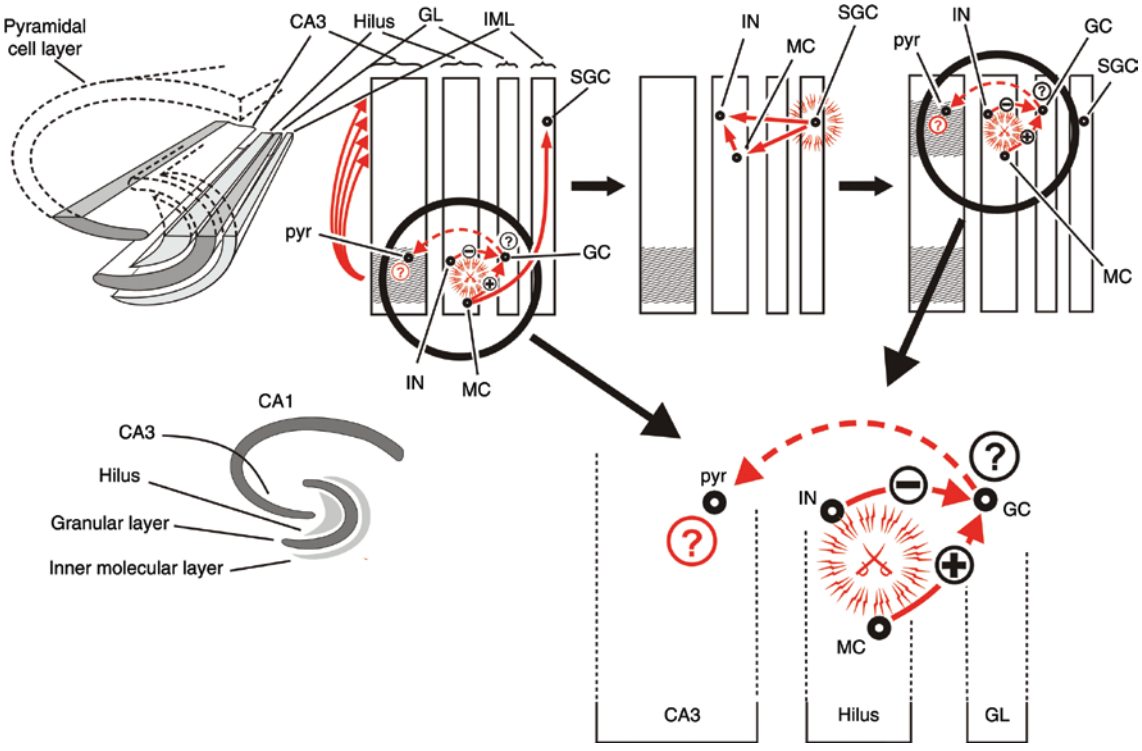

Figure 6: One step in the 'stirring', showing the locations of the cells involved.
Left: The principal HC cell layers of interest are shown in cross-section (bottom) and with suggestion of their 3D septotemporal extension (top). Only the inner one-third of the DG IML is included. Top right: Phases of a step in the longitudinal spreading. Bottom right: Magnified illustration of the brief 'race condition' (crossed swords), occurring in every theta cycle, due to the opposing assaults from MCs and hilar interneurons (IN), both initiated by volleys from the SGCs and both impinging on the local GCs. The unpredictable dynamics of simultaneous excitatory and inhibitory volleys (emphasized by question mark) add a random component to the occurrence of the GC burst firing during theta activity. The bursts, where they occur, induce plastic change via mossy terminals on the CA3 PCs (pyr). Also, the same GC bursts fire off a pool of MCs, which send out their effect to distant SGCs (see first group on top), and start the same process at new locations.

interlink the coactive cells via Hebbian reinforcement. Since the CA3 PCs, when tagged, emit a volley of spikes, and since the tagging, like the underlying burst activity of GCs, is linked to the theta activity, which confines it to relatively narrow time intervals within each theta cycle, the synchrony due to theta will create a situation favorable to inducing LTP (Pavlides et al., 1988).

The CA3-to-CA3 synapses, which are the hardware substrate for the cell assembly formation of interest, cause the network of CA3 PCs to be a good approximation to a randomly connected neuronal network (Buzsáki, 2006; Wittner et al., 2007). They are ideally suited for converting an arbitrarily selected subset of CA3 PCs into a permanently linked cooperative cell group, by causing them to fire in synchrony, and letting Hebbian LTP take its course (Figure 7A and B). In this manner, the gradually spreading GC activity leaves behind a gradually spreading ignitable cell assembly in CA3, which after some iterations would solidify into a 'pointer group'. Once ready, the pointer group becomes attached to afferents from the EC, via direct EC-to-CA3 synapses (which are NMDAR dependent; Martinez et al., 2002; Figure 7C).

Challenges of implementation and the way the HC addresses them

The following paragraphs propose to interpret a number of known features of the HC as solutions to specific problems of function and control, arising in connection with the protocols outlined above.

GCs 'hot' in one theta cycle are given preference thereafter

The burst firing of GCs is largely silenced at some point during each theta cycle, setting the stage for a new race of random selection for hot GCs (and tagged CA3 PCs). The problem is that if the selection is allowed each time to restart from scratch, without any memory of the previous cycle, the selections made in each cycle are simply replaced by the unrelated results of subsequent cycles, and any lasting snapshot effect is erased.

The problem calls for a mechanism to give the GCs and PCs, once selected, a competitive advantage in the selection process in subsequent theta cycles. In fact, there is evidence that a potentially suitable mechanism exists for this, as the synapses from MCs to GCs are NMDAR dependent (Jackson and Scharfman, 1996), and these synapses offer a viable substrate for giving the GCs the required advantage through classic Hebbian LTP.

Because a very large proportion of the MCs is expected to give off firing during the active phase of each cycle (Buzsáki, 1984; Freund and Antal, 1988; Scharfman, 1994b), it is expected that most of the hot GCs will receive MC input over many of their synapses, which means that the multisynaptic Hebbian LTP will tend to strengthen these synapses. The result will be a selective

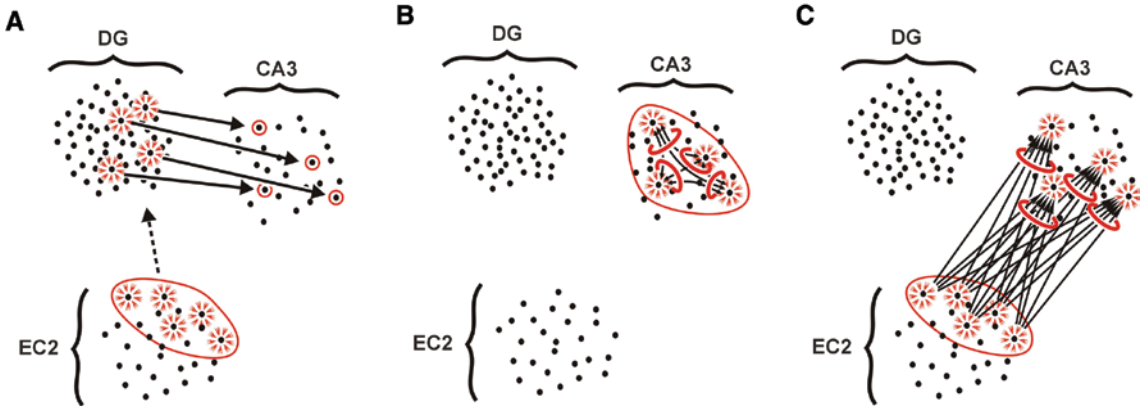

Figure 7: Effect of the DG events on CA3: the tagged cells link up into a cell assembly.
(A) In rough overview: Activity in EC (layer 2) combines with the effect of events local to the DG and theta drivers to cause a set of GCs to become 'hot' (shown by 'sparks'), and these in turn 'tag' a set of CA3 PCs (circled). (B) The pool of tagged PCs, co-firing under theta synchronization, contains enough sets of properly connected and properly timed cells (converging arrows circled) to cause their synapses, which are NMDAR dependent, to undergo highly selective Hebbian LTP and link up to form an ignitable cell assembly (a new pointer group). (C) Subsequently, the EC cells, also through NMDAR-dependent LTP, link up to the newly formed CA3 cell assembly (converging arrows circled), enabling the EC to selectively activate it. Through the EC cells, whose firing identifies the episode being recorded and its location, the new links enable later activation of the pointer group and the initiation of recall.

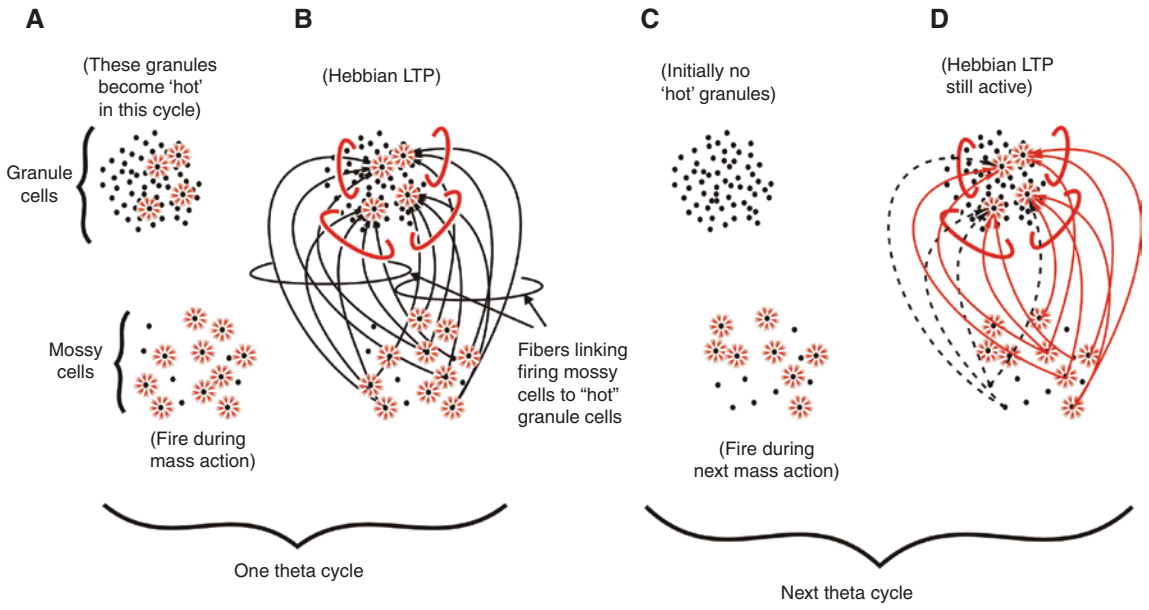

Figure 8: GCs hot in one theta cycle receive extra MC excitation in the next one.
The synapses from MCs to GCs are NMDAR dependent and accordingly require multisynaptic convergence for inducing LTP. Because the LTP requires coincident firing, which is highly selective, it creates correspondingly selective responsiveness in the postsynaptic neurons. (A and B) As a large proportion of MCs fire during most theta cycles, many of the hot GCs ('sparks') will have converging MC synapses terminating on them and will undergo Hebbian multisynaptic LTP (converging arrows circled). (C and D) In the next theta cycles, not all the previously active MCs will fire (see broken-line arrows), but enough of them will to enable the potentiated synapses to convey selective (although somewhat less selective) excitation to the same GCs that were hot in the previous cycle. The result is an enhanced likelihood that these GCs will become hot again.

reinforcement of these synapses, sending the hot GCs extra excitation in the subsequent theta cycles – excitation that is added to their randomly determined input and makes the GCs more likely to become hot again (Figure 8).

Keeping the cell assembly sparse despite repeated additions to it

Next, the problem arises that if the hot GCs are made to remain hot in all the subsequent theta cycles, and then new hot GCs are added in each of these cycles, the expected result is that the fraction of the hot GCs will keep increasing. If the set of tagged CA3 cells is to remain sparse, there must exist a mechanism for limiting the steady increase in the fraction of hot GCs.

The problem calls for a system of negative feedback in some form. In fact, the data on 'back-projection' from CA3 PCs to the DG (Scharfman, 1994a, 2007), particularly the back-projection to the hilar basket cells acting upon the GCs (Scharfman, 1994b; Kneisler and Dingledine, 1995), appear to offer a suitable feedback mechanism. As the number of tagged CA3 PCs increases, their summed inhibitory effect on the GCs will become stronger, until equilibrium is reached between the excitation and basket cell inhibition arriving to the GCs (Figure 9A–C). In the process of reaching an equilibrium situation in this way, it is expected that all septotemporal locations are visited by the process multiple times, leading to a 'well-stirred' cell population throughout the length of the HC (Figure 5).

At the end: undoing the changes made for preserving the hot GCs

The next issue is that after a cell assembly is completed and linked up, the network must be 'wiped clean', so as to be ready for creating the next cell assembly (to serve as pointer group for the next recorded episode). The problem is that, if the synaptic biases created above (Figure 8) are left in place, they will make the next pointer group similar to the one just created; because these biases were configured to preserve the pattern of hot GCs from one iteration to the next. However, this is just the opposite of what is needed when the iteration for a new pointer group is started. Clearly, if the pointer groups are to function as a 'random code ensemble', they must be uncorrelated to one another.

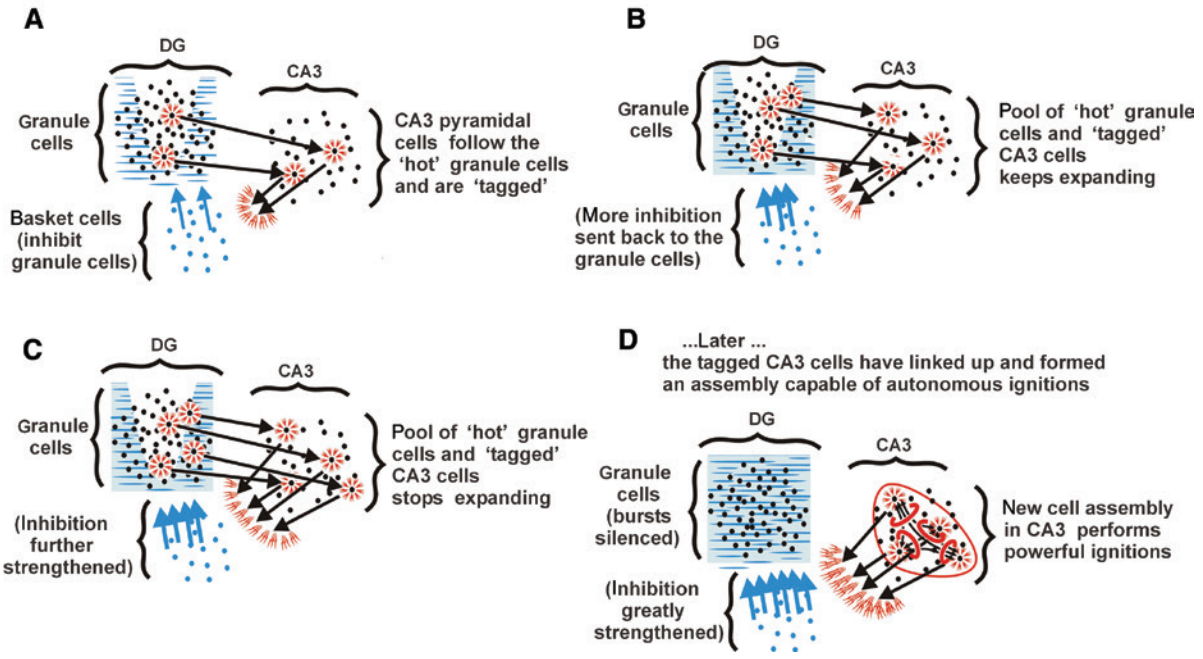

Figure 9: CA3 back-projection limits and eventually stops granule burst activity.
(A) CA3 PCs tagged by GCs ('sparks'), through back-projection, send excitation to basket cells, which inhibit the GCs (see 'cold bath' immersion of GCs). (B) In subsequent theta cycles, the same GCs are (ideally) still hot and their corresponding CA3 PCs remain tagged, but new ones are also added. Accordingly, the number of CA3 cells sending excitation to the DG basket cells is increased (the 'cold bath' gets colder), and the growth of the hot GC pool is slowed. (C) Back-projection-caused inhibition reaches a point where the density of hot GCs stops increasing, and an equilibrium is reached between the hot GCs and the tagged CA3 PCs. (D) The dynamics change after the CA3 PCs have linked up into an ignitable cell assembly (circles around arrows converging on cells) capable of performing ignitions autonomously and with EC input (see Figure 7); which means that CA3 firing at this point no longer requires input from the hot GCs. The autonomous CA3 ignitions are more powerful than the earlier firing, which had passively followed the GC bursts, and the result is that they now suppress the GCs sufficiently that all GC burst firing is stopped (there are no more hot GCs).

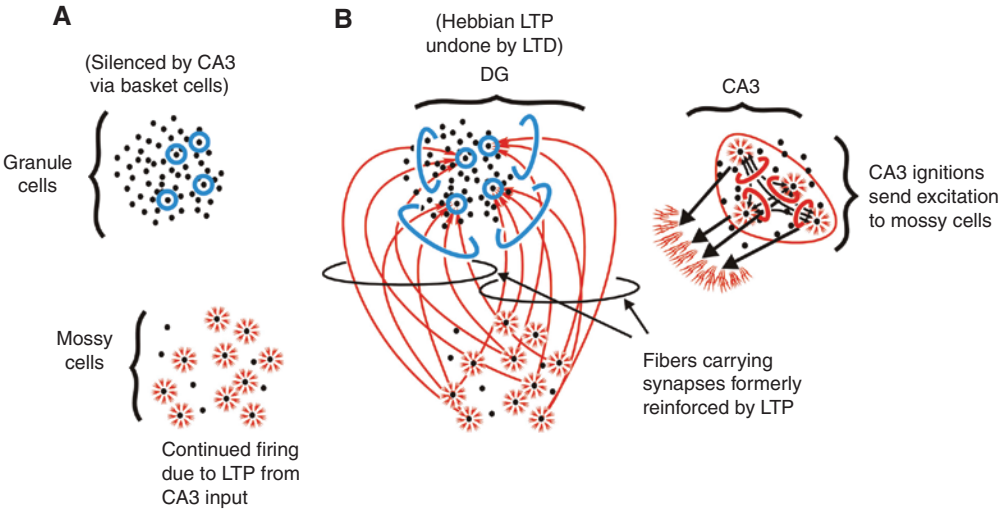

Figure 10: Switch from LTP to LTD to undo the bias favoring hot GCs.
The back-projection from CA3 to DG, in addition to suppressing the GCs via basket cells (Figure 9), also sends excitation to the MCs. After a newly formed CA3 cell assembly (pointer group) starts performing autonomous ignitions, its back-projection imposes a new situation on the DG: it acts to silence the GCs (Figure 9D), while sending excitation to the MCs and thereby causing them to send out firing, even in the absence of GC input. The result is that the MC-GC synapses, which previously underwent LTP (Figure 8), now undergo LTD through the combination of presynaptic firing and postsynaptic silence (see blue circling around the red arrows converging on the silent GCs).

The problem calls for a way to undo the aforementioned biases in some way. In fact, the well-documented machinery of long-term depression (LTD; Collingridge et al., 2010) fits the requirement, provided that LTD can be induced on the same MC-to-GC synapses that have previously undergone LTP.

The desired form of LTD will arise, at least approximately, if the GCs are forcibly silenced by some mechanism, while excitatory stimulation still continues to arrive to them from the same MCs that were presynaptic to the earlier LTP (Dudek and Bear, 1992). The required silencing mechanism is expected to be supplied (Figure 9D) by the intense volleys from the newly formed cell assemblies in the CA3 network and the consequent intense basket cell inhibition of the GCs (Scharfman, 1994b; Kneisler and Dingledine, 1995).

The second half of the requirement for LTD, the maintained firing by the MCs, involves the second component of the back-projection from the CA3. It will be noted that, with the GC activity suppressed, much of the dentate excitation received by the MCs is also suppressed. However, the back-projection from CA3 has a component sending excitatory input from the CA3 PCs to the MCs (Scharfman, 2007); these have the potential to help maintain the MC firing needed to support the LTD in the MC-to-GC synapses (Figure 10).

Granule neurogenesis: replacing the GCs degraded by residual bias

This would in fact undo the bias if the effect of LTD could precisely undo the earlier LTP, but in practice all steps of LTP and LTD are subject to some error and overlap, and as a result, only an approximation can be achieved to erasing the correlations created by the earlier LTP.

It is expected therefore that the machinery for creating uncorrelated cell assemblies is gradually degraded. The safest way to clear out the biases on GC input synapses is to replace the GCs altogether by brand-new ones.

Such considerations may explain why there is a significant turnover of GCs and why it continues through adulthood. In fact, the rate of adult neurogenesis in the rat has been reported to be as high as 1% per day (Cameron and McKay, 2001). A steady supply of fresh GCs is continually added to the mix, whereas old GCs are taken out of circulation is almost equal numbers, thereby maintaining a large pool of unbiased or almost unbiased GCs.

Significantly, a recent study by Scharfman and Bernstein (2015) found that a large fraction within the population of MCs was making its new synapses on adult-born GCs. The finding is consistent with the concept that when new pointer cell assemblies are formed, they tend to use the new and 'unspoiled' GCs in preference to the old ones.

Discussion

'Place cell dilemma'

The HC has been linked both to the storage of new memories and to spatial localization – two concepts that appear to be mutually contradictory, which are however both supported by incontrovertible evidence. The place cell data (O'Keefe and Dostrovsky, 1971; O'Keefe and Nadel, 1978; O'Keefe, 1999) and the data supporting nonspatial roles played by the HC (Cohen and Eichenbaum, 1991; Vargha-Khadem et al., 1997; Eichenbaum et al., 1999) are both extensive.

It is a corollary of the present paper that the two concepts are in fact compatible — because messages in the brain are not transmitted by neurons but by groups of neurons.

When the principal cells of the HC are studied individually, they are all found to be place cells. However, as argued here, when the HC records a piece of episodic memory, it assigns a cell assembly to it to serve as a pointer group; the assembly is presumably made up entirely of place cells appropriate to the location where the episode was recorded. Because there are many more appropriate place cells than are needed for any one pointer group, there is a wide variety of acceptable neuron choices, and the combinatorial diversity allows for a vast number of distinguishable pointer groups, enough to supply a lifetime's worth of episode identifiers.

There is a clear-cut functional rationale for storing location information side-by-side with the memorized details of an episode. It becomes apparent once it is noted that the pointer group of an episode is like the 'key' needed to unlock its memory contents, which means that when a specific item from an episode needs to be recalled (for instance the name of a person in the episode), it is first necessary to 'find the key'; in this case, it means locating the pointer group needed to 'unlock' the episode. Arguably, location information is among the most effective clues available for doing this.

It will be noted that when an important event, for instance an event involving danger, occurs in a particular place, it has survival value to be able to remember the event when the same place is visited again. In general, everyday human experience shows that when we 'lose a thought', returning to the place where the thought

was conceived is an effective way to 'jog our memory'. (In addition, people who recall a history-changing event from the past, like Pearl Harbor or the Kennedy assassination, tend to remember exactly *where* they were when they heard the news).

What is said of spatial information, in its capacity of leading the flow of consciousness to the pointer group of an episode, can also be said of time information. It is noted that in listing the deficits shared among HC patients, Vargha-Khadem et al. (1997) mentioned the inability of these patients to find their way in familiar surroundings side-by-side with their inability to orient themselves in date and time.

How does a cell assembly ignite a cell assembly?

The present model implicitly relies on the assumption that it is possible for one cell assembly to ignite another and that it is possible, further, to create the required addressable linkage between cell assemblies in an adaptive manner – all of which is not as simple as it may sound.

When it becomes necessary to induce a set of cell assemblies to perform coordinated ignitions, as when an episode is recalled (sometimes after many years), the awakened cell assemblies must first be 'restored' to being viable. This means that the synapses between cell assembly members must first be restored to the parameters they originally had at the time right after the assemblies were formed.

The reason this step is necessary is that, over time, the synaptic weights are usually overwritten by subsequent layers of unrelated LTP and LTD, and according to time-lapse microscopy data, the synapses disappear altogether after a few months as part of synaptic turnover (Lendvai et al., 2000; Trachtenberg et al., 2002; Stettler et al., 2006). Therefore, recreating an old cell assembly means individually addressing a vast number of synaptic locations for the purpose of recreating the synapses (Legéndy, 2016).

After a cell assembly (in this case, a pointer group) is made viable again, it must, in the brief time before being ignited, be brought close enough to instability to be easily ignitable. This involves dealing with the neurons as a whole rather than with the synapses, and bringing each neuron of a cell assembly close to threshold. Addressable induction of controlled instability in a set of neurons is expected to be possible, through the phenomenon of 'slow inward currents' (SICs), found to be governed from nearby astrocytes (Araque et al., 1998; Fellin et al., 2004; for reviews, see Fellin, 2009; Halassa and Haydon, 2010).

The two steps, recreating preassigned cell assemblies and sensitizing them to input, have been taken up in a separate paper (Legéndy, 2016), devoted to the microscopic mechanisms of long-term to memory, with emphasis on the logistics of information flow, and on setting up the local macromolecular systems in dendrites and astrocytes.

Acknowledgments: The author would like to thank Prof. Norma Graham for providing him with the research environment that made this work possible.

References

Acsády, L., Kamondi, A., Sik, A., Freund, T., and Buzsáki, G. (1998). GABAergic cells are the major postsynaptic targets of mossy fibers in the rat hippocampus. J. Neurosci. *18*, 3386–3403.

Amaral, D.G. and Witter, M.P. (1989). The three-dimensional organization of the hippocampal formation: a review of anatomical data. Neuroscience *31*, 571–591.

Araque, A., Parpura, V., Sanzgiri, R.P., and Haydon, P.G. (1998). Glutamate-dependent astrocyte modulation of synaptic transmission between cultured hippocampal neurons. Eur. J. Neurosci. *10*, 2129–2142.

Berlyne, D.E. (1950). Novelty and curiosity as determinants of exploratory behavior. Br. J. Psychol. *41*, 68–80.

Bevins, R.A. and Besheer, J. (2006). Object recognition in rats and mice: a one-trial non-matching-to-sample learning task to study "recognition memory". Nat. Prot. *1*, 1306–1311.

Bland, B.H., Andersen, P., Ganes, T., and Sveen, O. (1980). Automated analysis of rhythmicity of physiologically identified hippocampal formation neurons. Exp. Brain Res. *38*, 205–219.

Buzsáki, G. (1984). Feed-forward inhibition in the hippocampal formation. Prog. Neurobiol. *22*, 131–153.

Buzsáki, G. (2002). Theta oscillations in the hippocampus. Neuron *33*, 325–340.

Buzsáki, G. (2006). Rhythms of the Brain (Oxford, UK: Oxford University Press).

Buzsáki, G. (2010). Neural syntax: cell assemblies, synapsembles, and readers. Neuron *68*, 362–385.

Buzsáki, G., Leung, L.-W.S., and Vanderwolf, C.H. (1983). Cellular bases of hippocampal EEG in the behaving rat. Brain Res. Rev. *6*, 139–171.

Cajal, S.R. (1893). Estructura de1 asta de Ammon y fascia dentata (Madrid: Imprenta de Fortanet).

Cajal, S.R. (1995). Histology of the Nervous System (Oxford, UK: Oxford University Press).

Cameron, H.A. and McKay, R.D.G. (2001). Adult neurogenesis produces a large pool of new GCs in the dentate gyrus. J. Comp. Neurol. *435*, 406–417.

Chawla, M.K., Guzowski, J.F., Ramirez-Amaya, V., Lipa, P., Hoffman, K.L., Marriott, L.K., Worley, P.F., McNaughton, B.L., and Barnes, C.A. (2005). Sparse, environmentally selective expression of *Arc* RNA in the upper blade of the rodent fascia dentate by brief spatial experience. Hippocampus *15*, 579–586.

Chicurel, M.E. and Harris, K.M. (1992). Three-dimensional analysis of the structure and composition of the CA3 branched dendritic spines and their synaptic relationships with mossy fiber boutons in the rat hippocampus. J. Comp. Neurol. *325*, 169–182.

Cobb, S.R., Buhl, E.H., Halasy, K., Paulsen, O., and Somogyi, P. (1995). Synchronization of neuronal activity in hippocampus by individual GABAergic interneurons. Nature *378*, 75–78.

Cohen, N.J. and Eichenbaum, H. (1991). The theory that wouldn't die: a critical look at the spatial mapping theory of hippocampal function. Hippocampus *1*, 265–268.

Collingridge, G.L., Peineau, S., Howland, J.G., and Wang, Y.T. (2010). Long-term depression in the CNS. Nat. Rev. Neurosci. *11*, 459–473.

Correll, R.E. and Scoville, W.B. (1965). Effects of medial temporal lesions on visual discrimination performance. J. Comp. Physiol. Psychol. *60*, 175–181.

Donovick, P.J. (1968). Effects of localized septal lesions on hippocampal EEG activity and behavior in rats. J. Comp. Physiol. Psychol. *66*, 569–578.

Dudek, S.M. and Bear, M.F. (1992). Homosynaptic long-term depression in area CA1 of hippocampus and effects of N-methyl-D-aspartate receptor blockade. Proc. Natl. Acad. Sci. USA. *89*, 4363–4367.

Eichenbaum, H. (2001). The hippocampus and declarative memory: cognitive mechanisms and neural codes. Behav. Brain Res. *127*, 199–207.

Eichenbaum, H., Dudchenko, P., Wood, E., Shapiro, M., and Tanila, H. (1999). The hippocampus, memory, and place cells: is it spatial memory or a memory space? Neuron *23*, 209–226.

Fellin, T. (2009). Communication between neurons and astrocytes: relevance to the modulation of synaptic and network activity. J. Neurochem. *108*, 533–544.

Fellin, T., Pascual, O., Gobbo, S., Pozzan, T., Haydon, P.G., and Carmignoto, G. (2004). Neuroal synchrony mediated by astrocytic glutamate through activation of extrasynaptic NMDA receptors. Neuron *43*, 729–743.

Freund, T.F. and Antal, M. (1988). GABA-containing neurons in the septum control inhibitory interneurons in the hippocampus. Nature *336*, 170–173.

Gasparini, S., Migliore, M., and Magee, J.C. (2004). On the initiation and propagation of dendritic spikes in CA1 pyramidal neurons. J. Neurosci. *24*, 11046–11056.

Halassa, M.M. and Haydon, P.G. (2010). Integrated brain circuits: astrocytic networks modulate neuronal activity and behavior. Annu. Rev. Physiol. *72*, 335–355.

Hamming, R.W. (1950). Error detecting and error correcting codes. Bell Syst. Technol. J. *29*, 147–160.

Hangya, B., Borhegyi, Z., Szilágyi, N., Freund, T.F., and Varga, V. (2009). GABAergic neurons of the medial septum lead the hippocampal network during theta activity. J. Neurosci. *29*, 8094–8102.

Harris, E.W. and Cotman, C.W. (1986). Long-term potentiation of guinea pig mossy fiber responses is not blocked by N-methyl D-aspartate antagonists. Neurosci. Lett. *70*, 132–137.

Hebb, D.O. (1949). The Organization of Behavior (New York: Wiley).

Henze, D.A., Wittner, L., and Buzsáki, G. (2002). Single GCs reliably discharge targets in the hippocampal CA3 *in vivo*. Nat. Neurosci. *5*, 790–795.

Jackson, M.B. and Scharfman, H.E. (1996). Positive feedback from hilar mossy cells to granule cells in the dentate gyrus revealed by voltage-sensitive dye and microelectrode recording. J. Neurophysiol. *76*, 601–616.

Jung, M.W. and McNaughton, B.L. (1993). Spatial selectivity of unit activity in the hippocampal granular layer. Hippocampus *3*, 165–182.

Kamondi, A., Acsády, L., and Buzsáki, G. (1998). Dendritic spikes are enhanced by cooperative network activity in the intact hippocampus. J. Neurosci. *18*, 3919–3928.

Kerr, K.M., Agster, K.L., Furtak, S.C., and Burwell, R.D. (2007). Functional neuroanatomy of the parahippocampal region: the lateral and medial entorhinal areas. Hippocampus *17*, 697–708.

Kneisler, T.B. and Dingledine, R. (1995). Synaptic input from CA3 pyramidal cells to dentate basket cells in rat hippocampus. J. Physiol. *487*, 125–146.

Langdon, R.B., Johnson, J.W., and Barrionuevo, G. (1995). Posttetanic potentiation and presynaptically induced long-term potentiation at the mossy fiber synapse in rat hippocampus. J. Neurobiol. *26*, 370–385.

Larimer, P. and Strowbridge, B.W. (2010). Representing information in cell assemblies: persistent activity mediated by semilunar GCs. Nat. Neurosci. *13*, 213–222.

Legéndy, C.R. (1967). On the scheme by which the human brain stores information. Math. Biosci. *1*, 555–597.

Legéndy, C.R. (2009). Circuits in the Brain: A Model of Shape Processing in the Primary Visual Cortex (Berlin, Heidelberg, New York: Springer).

Legéndy, C.R. (2016). Synaptic and extrasynaptic traces of long-term memory: the ID molecule theory. Rev. Neurosci. *27*, 575–598.

Lendvai, B., Stern, E.A., Chen, B., and Svoboda, K. (2000). Experience-dependent plasticity of dendritic spines in the developing rat barrel cortex *in vivo*. Nature *404*, 876–881.

Leutgeb, J.K., Leutgeb, S., Moser, M.-B., and Moser, E.I. (2007). Pattern separation in the dentate gyrus and CA3 of the hippocampus. Science *315*, 961–966.

Losonczy, A. and Magee, J.C. (2006). Integrative properties of radial oblique dendrites in hippocampal CA1 pyramidal neurons. Neuron *50*, 291–307.

Lysetskiy, M., Földy, C., and Soltesz, I. (2005). Long- and short-term plasticity at mossy fiber synapses of mossy cells in the rat dentate gyrus. Hippocampus *15*, 691–696.

Martinez, C.O., Do, V.H., Martinez, J.L., Jr., and Derrick, B.E. (2002). Associative long-term potentiation (LTP) among extrinsic afferents of the hippocampal CA3 region *in vivo*. Brain Res. *940*, 86–94.

Mézard, M. and Montanari, A. (2009). Information, Physics, and Computation (Oxford, UK: Oxford University Press).

Milner, P.M. (1989). A cell assembly theory of hippocampal amnesia. Neuropsychologia *27*, 23–30.

Milner, B., Corkin, S., and Teuber, H.-L. (1968). Further analysis of the hippocampal amnesic syndrome: 14-year follow-up study of H. M. Neuropsychologia *6*, 215–234.

Mishkin, M. (1978). Memory in monkeys severely impaired by combined but not by separate removal of amygdala and hippocampus. Nature *273*, 297–298.

Nadel, L. and Moscovitch, M. (1997). Memory consolidation, retrograde amnesia and the hippocampal complex. Curr. Opin. Neurobiol. *17*, 217–227.

Nakazawa, K., Quirk, M.C., Chitwood, R.A., Watanabe, M., Yecke, M.F., Sun, L.D., Kato, A., Carr, C.A., Johnston, D., Wilson, M.A.,

et al. (2002). Requirement for hippocampal NMDA receptors in associative memory recall. Science *297*, 211–218.

O'Keefe, J. (1999). Do hippocampal pyramidal cells signal non-spatial as well as spatial information? Hippocampus *9*, 352–364.

O'Keefe, J. and Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence form unit activity in the freely-moving rat. Brain Res. *34*, 171–175.

O'Keefe, J. and Nadel, L. (1978). The Hippocampus as a Cognitive Map (Oxford, UK: Oxford University Press).

Palm, G. (1981). Towards a theory of cell assemblies. Biol. Cybern. *39*, 181–194.

Palm, G. (1982). Neural Assemblies. (Berlin, Heidelberg, New York: Springer-Verlag).

Paulsen, O. and Moser, E. (1998). A model of hippocampal memory encoding and retrieval: GABAergic control of synaptic plasticity. Trends Neurosci. *21*, 273–278.

Pavlides, C., Greenstein, Y.J., Grudman, M., and Winson, J. (1988). Long-term potentiation in the dentate gyrus is induced preferentially on the positive phase of Θ-rhythm. Brain Res. *439*, 383–387.

Petsche, H., Stumpf, C., and Gogolak, G. (1962). The significance of the rabbit's septum as a relay station between midbrain and the hippocampus: I. The control of hippocampus arousal activity by septum cells. Electroencephalogr. Clin. Neurophys. *14*, 202–211.

Rapoport, A. (1952). "Ignition" phenomena in random nets. Bull. Math. Biophys. *14*, 35–44.

Ribak, C.E., Seress, L., and Amaral, D.G. (1985). The development, ultrastructure and synaptic connections of the mossy cells in the dentate gyrus. J. Neurocytol. *14*, 835–857.

Rose, G., Diamond, D., and Lynch, G.S. (1983). Dentate GCs in the rat hippocampal formation have the behavioral characteristics of theta neurons. Brain Res. *266*, 29–37.

Scharfman, H.E. (1994a). Evidence from simultaneous intracellular recordings in rat hippocampal slices that area CA3 pyramidal cells innervate dentate hilar mossy cells. J. Neurophysiol. *72*, 2167–2180.

Scharfman, H.E. (1994b). EPSPs of dentate gyrus GCs during epileptiform bursts of dentate hilar "mossy" cells and area CA3 pyramidal cells in disinhibited rat hippocampal slices. J. Neurosci. *14*, 6041–6057.

Scharfman, H.E. (2007). The CA3 "backprojection" to the dentate gyrus. Prog. Brain Res. *163*, 627–637.

Scharfman, H.I. and Bernstein, H.L. (2015). Potential implications of a monosynaptic pathway from mossy cells to adult-born granule cells in the dentate gyrus. Front. Syst. Neurosci. *9*, 112.

Scott, A.C. (1977). Neurodynamics: a critical survey. J. Math. Psychol. *15*, 1–45. The article, as cited by Hebb (1976), was available in preprint form 2 years earlier, under the heading of "Neurodynamics (A Critical Survey). MRC Technical Summary Report No. 1548 (Mathematics Research Center, University of Wisconsin-Madison, WI, USA), October 1975".

Scoville, W.B. and Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. J. Neurol. Neurosurg. Psychiatry *20*, 11–21.

Shannon, C.E. (1957). Certain results in coding theory for noisy channels. Inf. Control *1*, 6–25.

Squire, L.R. (2009). The legacy of patient H. M. for Neuroscience. Neuron *61*, 1–7.

Squire, L.R. and Zola-Morgan, S. (1991). The medial temporal lobe memory system. Science *253*, 1380–1386.

Stettler, D.D., Yamahachi, H., Li, W., Denk, W., and Gilbert, C.D. (2006). Axons and synaptic boutons are highly dynamic in adult visual cortex. Neuron *49*, 877–887.

Sutula, T., Cascino, G., Cavazos, J., Parada, I., and Ramirez, L. (1989). Mossy fiber synaptic reorganization in the epileptic human temporal lobe. Ann. Neurol. *26*, 321–330.

Suzuki, W.A. and Amaral, D.G. (1994). Topographic organization of the reciprocal connections between the monkey entorhinal cortex and the perirhinal and parahippocampal cortices. J. Neurosci. *14*, 1856–1877.

Svoboda, K., Denk, W., Kleinfeld, D., and Tank, D.W. (1997). *In vivo* dendritic calcium dynamics in neocortical pyramidal neurons. Nature *385*, 161–165.

Tauck, D.L. and Nadler, J.V. (1985). Evidence of functional mossy fiber sprouting in hippocampal formation of kainic acid-treated rats. J. Neurosci. *5*. 1016–1022.

Teyler, T.J. and DiScenna, P. (1986). The hippocampal memory indexing theory. Behav. Neurosci. *100*, 147–154.

Tóth, K., Borhegyi, Z., and Freund, T.F. (1993). Postsynaptic targets of GABAergic hippocampal neurons in the medial septum-diagonal band of Broca complex. J. Neurosci. *13*, 3712–3724.

Trachtenberg, J.T., Chen, B.E., Knott, G.W., Feng, G., Sanes, J.R., Welker, E., and Svoboda, K. (2002). Long-term *in vivo* imaging of experience-dependent synaptic plasticity in adult cortex. Nature *420*, 788–794.

Tulving, E. (1972). Episodic and Semantic Memory. Organization of Memory. E. Tulving and W. Donaldson, eds. (New York/ Academic Press).

Tulving, E. and Markowitsch, H.J. (1998). Episodic and declarative memory: role of the hippocampus. Hippocampus *8*, 198–204.

Vargha-Khadem, F., Gadian, D.G., Watkins, K.E., Connelly, A., Van Paesschen, W., and Mishkin, M. (1997). Differential effects of early hippocampal pathology on episodic and semantic memory. Science *277*, 376–380.

Vertes, R.P. and Kocsis, B. (1997). Brainstem-diencephalon-septohippocampal systems controlling the theta rhythm of the hippocampus. Neuroscience *81*, 893–926.

Victor, M. and Agamanolis, D. (1990). Amnesia due to lesions confined to the hippocampus: a clinical-pathologic study. J. Cogn. Neurosci. *2*, 246–257.

Wickelgren, W.A. (1999). Webs, cell assemblies, and chunking in neural nets: introduction. Can. J. Exp. Psychol. *53*, 118–131.

Williams, P.A., Larimer, P., Gao, Y., and Strowbridge, B.W. (2007). Semilunar GCs: glutamatergic neurons in the rat dentate gyrus with axon collaterals in the inner molecular layer. J. Neurosci. *27*, 13756–13761.

Wittner, L., Henze, D.A., Záborszky, L., and Buzsáki, G. (2007). Three-dimensional reconstruction of the axon arbor of a CA3 pyramidal cell recorded and filled *in vivo*. Brain Struct. Funct. *212*, 75–83.

Zola-Morgan, S., Squire, L.R., and Amaral, D.G. (1986). Human amnesia and the medial temporal region: enduring memory impairment following a bilateral lesion limited to the CA1 field of the hippocampus. J. Neurosci. *6*, 2950–2967.

Zola-Morgan, S., Squire, L.R., and Amaral, D.G. (1989). Lesions of the amygdala that spare adjacent cortical regions do not impair memory or exacerbate the impairment following lesions of the hippocampal formation. J. Neurosci. *9*, 1922–1936.