Neural computation in the mammalian hippocampus

Timon Kunze*[1] and Alessandro Treves†[1, 2]

[1]SISSA, Trieste, Italy
[2]University of Agder, Kristiansand & Grimstad, Norway

November 2024

Contents

*https://orcid.org/0000-0002-4316-6946
†https://orcid.org/0000-0001-7246-5673

Keywords
mammalian hippocampus, attractor networks, computational modeling, memory, place cells, brain evolution, history of neuroscience, David Marr, CA3, dentate gyrus

Key Points

- At the core of the theoretical perspective inspired by David Marr is the hypothesis that the CA3 field of the hippocampus functions as an autoassociative recurrent network.
- The Dentate Gyrus, which precedes it, appears like a mammalian invention, useful to alleviate a storage-retrieval conflict common to all such networks.
- Ideas about the contribution of the CA1 field are more vague, despite the wealth of experimental data.
- While many details remain to be clarified, recent observations are sowing doubts about consolidated conclusions as well.

Abstract

Lesion studies both in mammals (in particular rodents and humans) and in non-mammals show that the hippocampus and its homologue structures are critical for spatial or otherwise complex memory. The internal organization with its subdivision into the three main fields DG, CA3 and CA1 is, however, highly specific to mammals and common to all mammalian species.

The pioneering efforts of David Marr aimed at understanding this organization as the one required to serve its memory function. Successive mathematical and computational modeling has largely validated his general perspective, but brought it much closer to experimental data, especially with rodents.

Now that important questions appear to have been clarified, new findings are casting doubt on the validity of some idealized assumptions, and on the fact that our theoretical understanding has been based almost exclusively on data obtained in the lab rather than outside.

1 Introduction

The hippocampus has been understood to play a crucial role in the formation of our autobiographical and episodic memories (Scoville & Milner, 1957) and is thus intimately related to our sense of individual identity, to 'who we are', as distinct and mutually irreplaceable persons. In a wider sense, however—zooming out from a time scale of decades to one of a hundred million years—the hippocampus is also intimately connected to our collective mammalian identity. Our mammalian hippocampus is not the sole element, but probably a critical one, in making us mammals. This is evident if one contemplates a cross section of the hippocampus as it appears in different mammalian species, and readily sees the same design, with the same sub-fields labeled as DG, CA3,

CA1 (Figure 1 with illustrations by Golgi, 1885 and Gloor, 1997)—despite obvious differences in overall scale and less obvious but nonetheless important differences even among closely related species (Amrein, Slomianka, & Lipp, 2004). Comparative neuroanatomy tells us that such mammalian design is different from those observed in our cousins, reptiles and birds (Figure 2; Striedter, 2016). Even though calling it a 'design' might attribute too much intentionality to the random stuttering of evolutionary processes, it is hard to escape the impression that the highly specific organization of the mammalian hippocampus is there for a purpose—as a neural 'chip' that is particularly efficient at implementing a certain function. We would like to validate such an impression, and better understand to what extent the function determines or at least favors the particular structure we observe. But what is this function?

Here come the two major non-linearities in the development of hippocampal science over the last few decades. First, two narratives have been laying competing claims to capture the core of what the hippocampus does: episodic memory vs. spatial computations. Second, however one chooses to describe the function in abstract terms, those terms seem to apply also to what the hippocampus does in non-mammals, e.g. in birds, whose hippocampus has a different internal structure. Therefore, to be able to speculate convincingly about the computational link between structure and function, how function determines structure—the grand aim of the research program of the young David Marr (1971)—it has been necessary both to characterize better what function we consider, and to accept that such determination of the structure by the function may not be single-valued: the same function might favor two or more structural 'solutions'.

The 'alternative' view, that the hippocampus operates primarily as a spatial computer, has been variously presented as just a difference in emphasis all the way down to a radically distinct hypothesis. In the latter case, the focus is on the geometric operations that the hippocampal formation is required to implement, such as representing locations and bearings and comparing them to obtain distances and angles or composing them via vector summation. Although explored computationally only in a number of instances (O'Keefe, 1990; McNaughton et al., 1996; O'Keefe & Burgess, 1996), this view has been enormously influential in inspiring experimental studies, particularly in rodents, and there are no sharp boundaries, in the complex body of ideas generated by the discovery of place cells (O'Keefe & Dostrovsky, 1971) and stimulated by the O'Keefe and Nadel (1978) book, between summaries of findings, conceptual models and rigorous theories, or between strictly spatial and strictly memory computations.

In the following, we will not attempt to force such boundaries where they do not belong, and try instead to blur those between research approaches, by relating the two dominant fields of enquiry, centered on humans and on rodents, to evidence available for the many other species, mammalian and non-mammalian, who can be said to use their hippocampus in somewhat similar ways.

2 Part I: The hippocampus from the outside

Before we look at the hippocampus inside, we should bring to mind the 'black box' picture of hippocampal anatomy and function.

In the first subsection, we review the comparative anatomy of the hippocampal for-

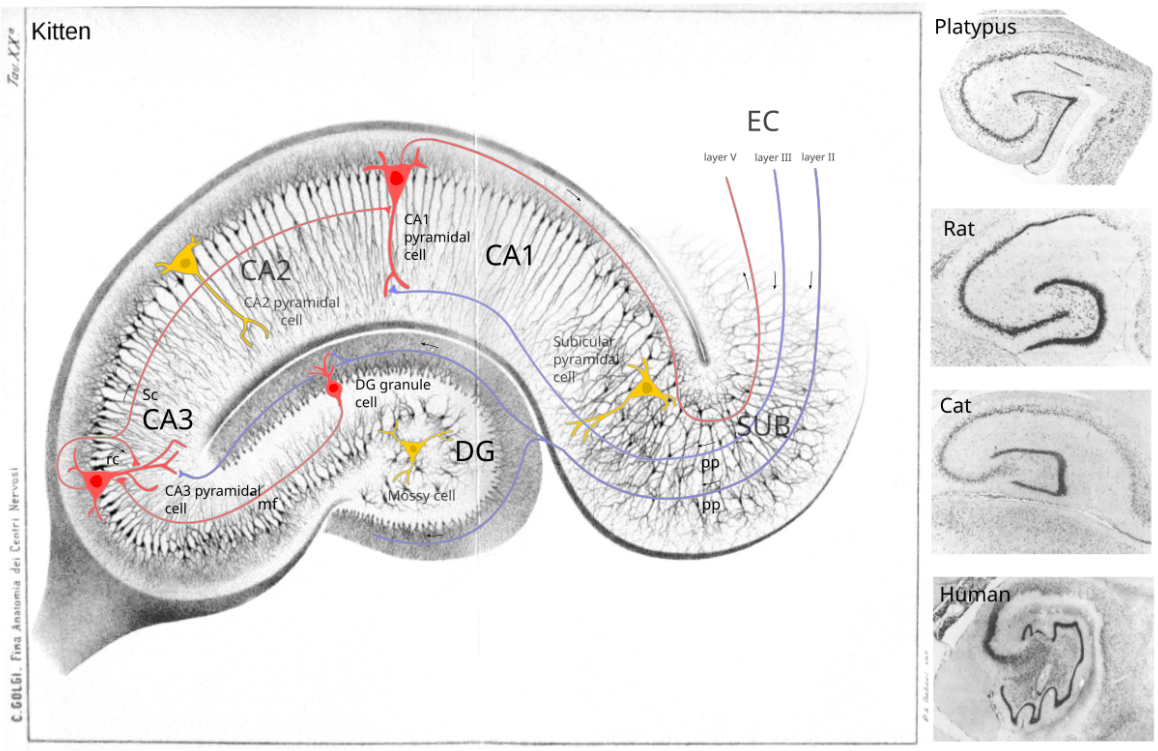

Figure 1: Left – In the background the hippocampus of a kitten as drawn by Camillo Golgi (1885) from slicing orthogonal to its elongated dimension. Information enters the hippocampus via two separate perforant pathways (pp) originating from layers III and II of EC, shown in purple. The three major regions of the trisynaptic pathway, DG, CA3 and CA1 are labeled and represented by enlarged sample cells in red, with their own axons. Granule cells project to CA3 with mossy fiber axons (mf); CA3 pyramidal cells project to each other via recurrent collaterals (rc) and to CA1 via Schaffer collaterals (Sc), both originating from the same axons; and CA1 pyramidal cells project back to EC in layer V. The regions CA2 and SUB (comprising subiculum, pre- and parasubiculum) as well as exemplary subicular and mossy cells are additional elements of the hippocampal circuitry, drawn in yellow, whose potential roles are not discussed in this review. Right – Cross sections through four mammalian hippocampi, rotated such that they are aligned. The conformity across mammalian species is remarkable. Images modified from Gloor (1997).

mation, across different species. In the second subsection, we take a look at what could be the general function of the hippocampus, as it emerges from black box approaches, that is, essentially from lesion studies.

2.1 The place of the hippocampus in the brain

During embryonic development, the precursor of the central nervous system of all vertebrate species forms a tiny pipe, the neural tube (shown in cross section in Figure 2, left). What will become the hippocampus sits right on top of the neural tube. When the embryo develops further, the two sides in the rostral or front-end portion of the tube pop out, through a process called evagination, beginning to shape the two cerebral hemispheres. The hippocampus, originally on top, comes to line the groove that separates the two bulging hemispheres, the so-called medial pallium, the internal portion of the mantle. This occurs in humans and in all other mammals, as well as in amniotes (including reptiles and birds), amphibians, and in the original Sarcopterygii, the lobe-finned fish we all are presumed to descend from. A different process, however, called eversion, occurs in the development of the Actinopterygii—comprising most of bony fish—whereby the neural tube splits up and opens at the top. That way, the piece of neural tissue homologous to the hippocampus (we can still call it hippocampus) comes to sit lateral, on the two sides of the telencephalon.

This rather dramatic difference, between an eversive and an evaginating forebrain, which may hark back to genetic mutations from over 400 million years ago, is puzzlingly NOT accompanied by equally dramatic or even apparent differences in the role of the hippocampus within the nervous system, as we shall see below. It would seem that, wherever the hippocampus ends up in the adult brain, it is still the hippocampus. A bit like exhaust pipes serving as exhaust pipes, whether displayed laterally as in Peter Fonda's chopper in Easy Rider or hidden under a Citroën 2CV. We can call it a sort of structural phase transition, around 400 million years ago, where the separation between two major lineages of bony fish marks a critical event in the evolution of vertebrates. Note, however, that this phase transition into an eversive developmental dynamics seems to have occurred to Actinopterygian fish only and not to us, since sharks, which as cartilaginous fish diverged earlier from our lineage than Actinopterygians, appear to evaginate as we do (Docampo-Seara et al., 2018).

Fast forward 400 million years to the present, we should however appreciate the variability in the final position of the adult hippocampus also among species whose forebrain evaginates. This is because in those mammals with larger body and brain, the hippocampus tends to slide from the top towards the back and then comes to lie at the bottom of the two hemispheres. In humans, it sits near the bottom of the cerebral cortex, at the internal hem of its temporal lobes (Figure 2, right). One should also remember that, as noted by Herrick (1933), 'below' reptiles the entire pallial field is dominated by the olfactory system; then in reptiles a simple layered cortex does appear, but it is still to a large extent receiving prevailingly olfactory fibers. A major discontinuity is between reptiles and mammals, since in the latter the olfactory cortex, i.e., the piriform cortex—but also, in his perspective and in that of early neuroscientists, the hippocampus—is clearly separated from the non-olfactory neocortex.

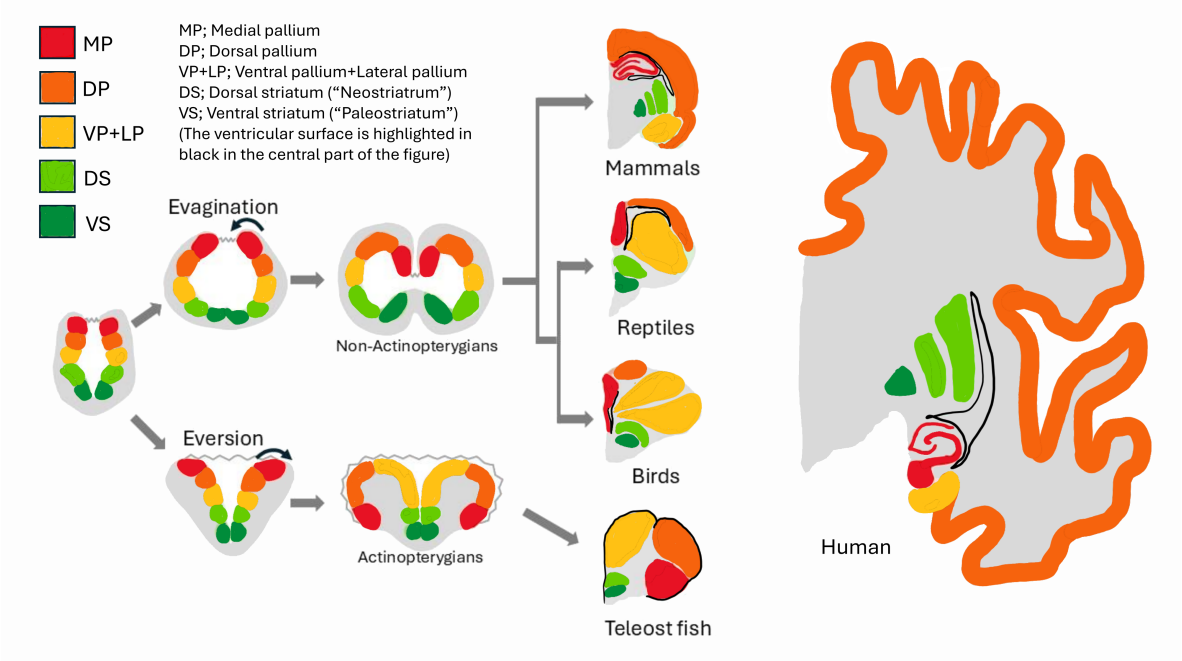

Figure 2: Left – Schematics of early brain development in actinopterygian and non-actinopterygian fish with respect to the eversion and evagination of the neural tube; the hippocampal homologue is shown in red; graphic adapted from Rodríguez, López, Vargas, Gómez, et al. (2002). Center – The position of the hippocampus in adult higher vertebrates reflects the evagination process they have inherited (one hemisphere only is shown); adapted from Rodríguez, López, Vargas, Broglio, et al. (2002). Right – In humans, the hippocampus comes to lay at the medial border of the temporal lobes; from Gloor (1997), redrawn—like the other diagrams—with substantial help from Cosme Salas.

2.2 Hippocampal function from lesion studies

Until the mid-50s, the functions of the hippocampus remained rather unclear. In his book 'Brain Mechanisms in Diachrome' (1955), Wendell Krieg tentatively referred to the hippocampus as the motor division of the olfactory system. This might make sense from an evolutionary perspective, since the entire pallium or upper portion of the forebrain is considered to have been originally olfactory territory, as noted above. However, a neuroanatomist like Alf Brodal had cast doubt on this view (Brodal, 1947): first, there was not much evidence of olfactory fibers actually reaching the hippocampal formation, however physically close it may be to piriform cortex (evincing support from Ramón y Cajal, 1911); second, behavioral experiments showed that lesioning the hippocampus in cats did not affect olfactory-conditioned reflex responses much. This paved the way for theories suggesting that the hippocampus is involved in other cognitive functions such as emotion, memory and space (see however Vanderwolf, 2001 for a contrarian position).

2.2.1 Hippocampal patients

At this point, compelling evidence for a role of the hippocampus in memory came from the observation by Brenda Milner of lobotomized patients, who had become amnesic. William B. Scoville and other neurosurgeons carried out fractional hippocampus lobotomies on patients who were severely ill—mostly with schizophrenia, sometimes with epilepsy—to alleviate their suffering after all other forms of therapy had failed. In most cases the surgeries were helpful, but sometimes they produced other serious and unexpected side effects: Several patients showed severe memory deficits, particularly when the hippocampus and hippocampal gyrus were removed on both sides (Scoville & Milner, 1957).

The most striking and consequently best studied case was Henry Gustav Molaison ('patient H.M.' as he was known to the scientific community before his death in 2008). Molaison underwent Scoville's bilateral temporal resection in 1953 to free him from his epilepsy when he was 29 years old. This resulted in a profound memory loss, preventing the formation of any new conscious or unconscious memories, a condition that would not improve during his following life. Meanwhile his working memory abilities, his perceptual abilities and even his language abilities in comprehension and production were largely left intact (Corkin, 2002; Gabrieli, Cohen, & Corkin, 1988). This most severe and pervasive anterograde amnesia was also likely accompanied by a temporally graded retrograde amnesia, i.e., a difficulty with recalling events that occurred before the injury, although the very distant past appeared to have been progressively less affected than more recent memories. Retrograde amnesia is harder to test in a controlled way, but Corkin and colleagues only found a (semantic) memory impairment for public and personal events from approximately 11 years before his operation (Corkin, 1984). Careful studies of the temporal gradient have been carried out with other hippocampal patients (Graham & Hodges, 1997; Manns, Hopkins, Reed, Kitchener, & Squire, 2003; Bayley, Hopkins, & Squire, 2006).

Note that while the evidence for anterograde amnesia of episodic memory is overwhelming (Spiers, Maguire, & Burgess, 2001), it seems possible to learn some new semantic memories even without the hippocampus, and even Molaison might have acquired a few during his long life (Vargha-Khadem et al., 1997; Sharon, Moscovitch, & Gilboa, 2011).

2.2.2 Lesion studies in other primates

Immediately after Scoville and Milner (1957), many attempts were made to reproduce the pervasive memory deficits of hippocampal patients in non-human primates, cats and rodents. This proved to be surprisingly difficult, however. As Mortimer Mishkin 1978 remarked, 'Hippocampal-system lesions in animals do markedly impair some forms of spatial memory, but the effects on other forms of memory have generally seemed minor'.

David Gaffan (1974) tested monkeys with fornix lesions in recognition memory and associative memory tasks. The monkeys were severely impaired in recognition memory and unimpaired in simple association memory, a dissociation that is similarly found in amnesic people. Mishkin (1978), on the other hand, performed an object-recognition task with monkeys. In a four-group design, the monkeys were either inflicted with a hippocampus lesion, an amygdala lesion, a combination of both types of lesions or nei-

ther. As it turned out, only the combined lesions produced a severe memory deficit: the monkey's performance dropped to chance levels and monkeys needed 10 times as many trials and made 10 times as many errors to relearn it. The extent of the hippocampal lesion is therefore a critical factor, and its relationship with the degree of the memory impairment has been assessed in a meta-analysis by Zola and Squire (2001).

In a series of extensive studies based principally on the fornix transection approach, Gaffan has further characterized the memory impairment in monkeys due to a disabled hippocampus. Using scenes from Raiders of the Lost Ark, he has observed a deficit in the acquisition of complex naturalistic scenes (but with normal forgetting of those acquired before the lesion; Gaffan, 1992, 1993). He has later argued against the simple-minded view that the hippocampus deals with spatial memory, while specific cortical regions take care of e.g. object memory (Gaffan, 2002), and in favor of a more holistic view, in which the contribution of the hippocampus, in part due to the neuromodulatory inputs it receives, is particularly in the formation of complex memories involving distributed storage of information across cortical areas—a view that resonates with the one at the basis of Marr's (1971) theory.

2.2.3 Lesion studies in rodents

In over half a century of extensive investigations, that has included in the early days also studies in cats (e.g., McDonough & Kesner, 1971), Ray Kesner and others have probed the effect of hippocampal lesions in rodents using a variety of tasks and generating a wealth of insights also on the contributions of specific hippocampal sub regions (see Kesner & Rolls, 2015). One non-spatial hippocampal function that has been analyzed in these studies, and tentatively ascribed to the output end of the rodent hippocampus, the CA1 subregion, is the linking of the representations of, e.g., distinct odors along a temporal continuum (Kesner, Gilbert, & Barua, 2002).

At a general hippocampal level, a large number of rodent studies have utilized a particularly brilliant experimental task that relies, transparently, on spatial memory. Morris, Garrud, Rawlins, and O'Keefe (1982) introduced the water maze, where rodents are placed in a large circular pool of opaque water and spontaneously try to escape the water by reaching a small platform hidden beneath the water surface. A quick escape requires learning the position of the platform in relation to visual cues outside the maze (e.g., above the walls of the pool), and experimenters can monitor learning behavior over trials and days, by measuring the time to reach the platform and where the rodent searches for the platform in trials in which it has been removed. Rats and mice with hippocampal lesions are impaired in the water maze, and the impairment is neither due to motor problems nor to a lack of motivation (Morris et al., 1982). Moreover, impairments have been observed by blocking synaptic plasticity, and they do not affect spatial information already acquired—they are clearly learning impairments (Morris, 1989).

The connection between memory performance and location and extent of hippocampal lesions could be clearly demonstrated in these rodent studies. It was found that especially dorsal lesions to the hippocampus, for example, reduce performance in the Morris water maze (Moser, Moser, Forrest, Andersen, & Morris, 1995).

In terms of the characterization of the types of memory that are affected, however,

rodent studies may have led to a somewhat biased impression, because of the ease with which spatial and context components can be introduced in the assessment of memory performance, and the relative difficulty of assessing behaviorally complex memories of a non-spatial nature.

2.2.4 Lesion and plasticity in birds

Studying passerine birds, selected both from species which store food and other which do not, Krebs, Sherry, Healy, Perry, and Vaccarino (1989) observed that food-caching birds have a significantly larger hippocampus (or hippocampal homologue, if one wants to refer to the avian structure in that way), relative to their body size. Further, they observed that hippocampal damage disrupts memory for the storage sites used by those species. These findings have been extended in an impressive number of studies with several species, as reviewed by Krebs et al. (1996), showing also that depriving food storing species of the opportunity to store food can lead to reduced hippocampal volumes and neuron numbers. Volume and neuron numbers might even increase seasonally in some species when memory demands become intense, (for instance at the end of summer; Smulders, Shiflett, Sperling, and DeVoogd, 2000). Remarkably, food caching behavior has been demonstrated in a number of studies, starting from the one by Clayton and Dickinson (1998), to involve memories that need to be characterized as episodic rather than merely spatial (see the discussion in Davies and Clayton, 2024).

The hippocampus has been shown to be critically involved in another avian behavior: homing in pigeons. A series of studies has demonstrated that hippocampal lesions impair the ability of pigeons to navigate long distances, often tens of kilometers, to return home (Gagliardo et al., 2020) and also that the hippocampus is larger in pigeons that have had navigational experience, relative not only to non-homing pigeons but also to those who have not had it yet (Cnotka, Möhle, & Rehkämper, 2008). Intriguingly, aging pigeons seem to be challenged by navigational tasks not due to their hippocampus having shrunk, but because they activate it less (Coppola & Bingman, 2020).

2.2.5 Lesions in reptiles and fish

Testing spatial memory in reptiles is not straightforward (Font, 2019), which might account for the unclear outcome of hippocampal (i.e., medial pallial) lesions in lizards (Day, Crews, & Wilczynski, 2001), but even with lizards there are indications that hippocampal size is correlated with foraging behavior (Day, Crews, & Wilczynski, 1999). In turtles, hippocampal lesions have been shown to lead to deficits in a place learning but not in a cue learning task (Rodríguez, López, Vargas, Broglio, et al., 2002; Rodríguez, López, Vargas, Gómez, et al., 2002). Most interestingly, the same study reports place learning deficits in goldfish following lateral pallial lesions—which makes sense, given that in teleost fish the embryonic hippocampus or hippocampal homologue ends up in the lateral pallium, as mentioned above. Research in spatial memory in goldfish has also produced tantalizing indications about differences in activation of distinct regions of the lateral pallium, during the acquisition of a spatial task (Ocaña, Uceda, Arias, Salas, & Rodríguez, 2017); but the evidence is still too preliminary—and probably too unknown to the mainstream modeling community—to have inspired the neural

computation models which are at the center of this review. Nonetheless, it appears a most promising direction for future research.

2.2.6 Functional summary

Given the many attempts to characterize exactly what kind of memories are affected by hippocampal lesions, a reasonable recapitulation of the system-level studies briefly reviewed above may be that the hippocampus (or the pallial regions homologue to the mammalian hippocampus) is involved in the formation of complex memories, across vertebrates.

Complex memory is a somewhat vague term, which can encompass the abstract notions of undirected (or 'free') and of 'simple' memory used by Marr (1971) in his paper, as discussed in the next Section. We use 'complex' not in contrast to his 'simple' and 'free' qualifications (which would appear to be opposite: what is complex is usually not simple, nor free), but in the sense of involving several distinct elements and their relations. The memory for a single element, such as a cue, the sight of a green triangle, or an association between just two elements, such as a sound and a weak electric shock, are not complex in our terminology. Complex memories can include also the more concrete cognitive characterizations of episodic memories (Tulving, 1972), autobiographical memories (Levine, Svoboda, Hay, Winocur, & Moscovitch, 2002), and other types of declarative memories including, with certain qualifications, semantic memories (Eichenbaum, 1997; Manns et al., 2003), 'conscious' memories (Moscovitch, 2008), and even incidental memories (Torromino et al., 2022). Complex memories typically require the representation of event-related relations among already consolidated elements (Teyler & DiScenna, 1986); relations which could be sequential (Eichenbaum, Otto, & Cohen, 1994; Lisman, 1999; Cheng & Werning, 2016), or spatial (Burgess, Maguire, & O'Keefe, 2002), sometimes building navigational and cognitive maps (Tolman, 1948; O'Keefe & Nadel, 1978), even if they are expressed as graphs (Muller, Stead, & Pach, 1996; Garvert, Dolan, & Behrens, 2017), and other constructs that express the relations among their elements.

The term 'complex memory' is, however, intended to exclude simpler types of memory that do not appear to require the hippocampus for their formation, such as, again, cue learning or other sorts that involve simple associations between two items, of the type considered in cue learning, classical and operant conditioning, as well as in many forms of statistical learning. Emphasizing the formation of complex memories is not meant to exclude, moreover, an important hippocampal role in their retrieval and retention, in certain conditions (Ryan et al., 2001). An abstract, streamlined representation of what storing these complex memories would entail is reproduced in Figure 3: taking a snapshot of distributed cortical activity, at the time a particular event is experienced (Murre, 1996).

This, then, could be the simplest working assumption for what the hippocampus is there to do, and one that we can use, following in the footsteps of Marr (1971), to try to make sense of the structure, the internal organization of the mammalian hippocampus.

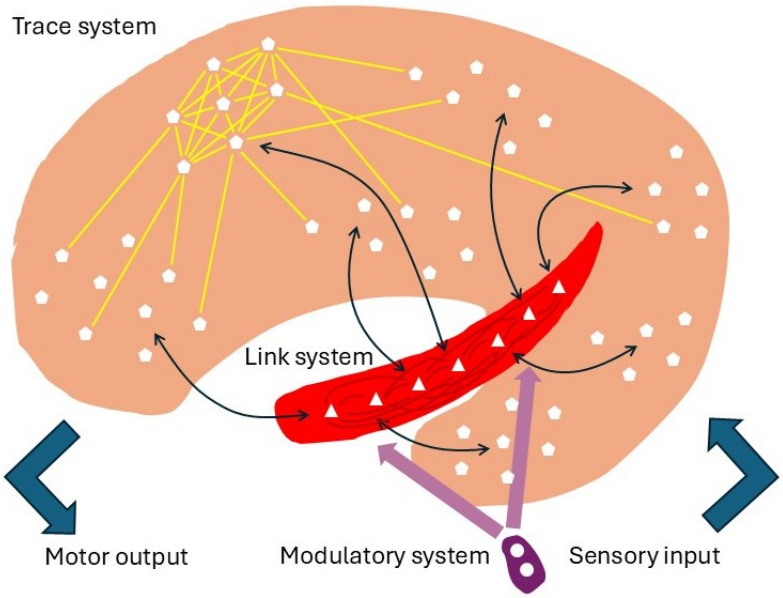

Figure 3: Marr's view of the role of the hippocampus in memory as conceptualized in the TraceLink model developed by Murre (1996). The hippocampus (the Link system) would capture and encode with the help of neuromodulators (chiefly Acetylcholine) a snapshot of activity in the cortex (the Trace system). Complex memory can be established in one shot in the Link and not in the Trace system because of insufficient long-range connectivity in the latter (Marr, 1971), or perhaps due to the interference resulting from their compositional nature (Ryom, Stendardi, Ciaramelli, & Treves, 2023), or because backpropagation requires slow learning (McClelland, McNaughton, & O'Reilly, 1995).

3 Part II: The hippocampus inside

The hippocampus is so clearly ordered, compared to the relative mess of brain tissue from other regions, that one feels an urge to understand the logic behind such clear design. This can be seen by slicing a section orthogonal to its elongated dimension, through, e.g., a rat, rabbit, or cat hippocampus, as in the famous drawings by Camillo Golgi (Figure 1; Golgi, 1885, see Bentivoglio et al., 2019) and by Santiago Ramón y Cajal (1893, 1911). Such a slice includes samples of neurons from the major 'subfields' of the hippocampus proper, traditionally called dentate gyrus (DG), Cornu Ammonis 3 (CA3) and Cornu Ammonis 1 (CA1), after the classification of Lorente de Nó, Cajal's last student (1934). The slice, even when relatively thin, preserves some of the synaptic connections between the regions, which are strikingly unidirectional, and can be used for neurophysiological in vitro assays that probe those connections. This led Per Andersen, who had been impressed by the anatomy of a slice shown to him in the early '50s by Theodor Blackstad (see Blackstad, 1958), to return to it a dozen years later, and formulate the notion of the 'trysynaptic circuit': from neocortex to DG to CA3 to CA1 (Andersen, Holmqvist, & Voorhoeve, 1966). The fascinating story, which has been told by Bliss and Lømo (2024), was completed by the demonstration that, unlike in Cajal's drawings, CA1 projects its axons not back to CA3 but rather to the subiculum, from where the wave of excitation can return to the cortex: the trisynaptic circuit is therefore to a first approximation an excitatory loop attached to the cortex, a bit like the autoencoder scheme of the artificial neural networks developed much later. But what is the first approximation missing out?

3.1 Principles and parameters in hippocampal anatomy

In a famous research program from the '80s, Noam Chomsky (1981) developed a description of cross-linguistic variation in syntax, the internal structure of natural languages, in terms of general principles, applicable to the syntax of all languages, and of a finite set of 'parameters', which take different values across languages, and are usually considered to be binary variables. Aside from this last feature, i.e. the binary nature of parameters, a similar description in terms of principles and graded-valued parameters can be useful to describe the internal structure of the hippocampus across mammalian species.

The trisynaptic circuit is a basic principle of mammalian hippocampal design, but not the only one. Another principle involves the shortcuts that supplement the trisynaptic circuit: the principal (i.e., the pyramidal) cells of the CA3 region receive the axons of the DG granule cells, the so-called mossy fibers (MF), but also directly the axons of the cortical cells that project to DG. Moreover, CA1 pyramidal cells receive the axons of CA3 pyramidal cells, but also their own distinct projections from the cortex (called, like those reaching to DG and CA3, the perforant path, PP). Therefore, synchronous cortical activation, in particular of the entorhinal cortex acting as the gateway to the hippocampus, can in principle reach CA1 with three temporally distinct excitatory volleys: first directly via PP, then by activating CA3 pyramidal cells via PP, and at last by activating CA3 pyramidal cells from DG granule cells via the trisynaptic circuit. Intriguingly, the last two volleys coming from CA3 convey largely the same information

to CA1, since DG granule cells are innervated largely by the very same fibers that then continue and innervate CA3. We return to this point later.

CA3 pyramidal cells are endowed with an extensive system of recurrent synaptic connections, through which they excite each other. In rats it has been estimated that 3/4 of the synapses onto its pyramidal cells are from their fellow cells in CA3 (Amaral, Ishizuka, & Claiborne, 1990). The efficacy of each individual synapse is associatively modifiable, a plasticity which is considered to be at the core of the memory functionality of the hippocampus (Debanne, Gähwiler, & Thompson, 1998). Recurrent connections are not an unusual feature of basic cortical circuitry, the presumed ancestor of mammalian hippocampal circuitry, although in CA3 they may be particularly well developed. What is striking, however, are their scarcity or near absence among the pyramidal cells of neighboring region CA1. It appears that in CA1 there is an architectural requirement to 'avoid' cross-talk between its principal cells. Likewise, there are no recurrent connections between the principal cells of the dentate gyrus, the granule cells, whose unidirectional dendritic trees are devoted to cortical inputs and to synapses from small populations of dentate interneurons (including the excitatory mossy cells).

While CA3 stands out for its numerically abundant recurrent collaterals, also the few but individually strong synapses made by the mossy fibers from the dentate granule cells are quite peculiar, apparently across species. They are made on large thorny excrescences, very different from normal spines, and through a complex counting process it has been reckoned that there are no more than some 40 per receiving CA3 cell, again in young adult rats (Gonzales, DeLeon Galvan, Rangel, & Claiborne, 2001). Notably, the presynaptic terminals of these synapses are rich in zinc, which they release when the dentate gyrus is stimulated with sufficient strength (Aniksztejn, Charton, & Ben-Ari, 1987), and impulse transmission can be blocked by a zinc chelator (Lassalle, Bataille, & Halley, 2000).

Alongside excitation, the organization of inhibition in the mammalian hippocampus is also thought to follow the same design principles across species, with separate classes of inhibitory interneurons carrying out conceptually distinct network operations (Gulyás, Megías, Emri, & Freund, 1999), as Marr (1971) had foreseen, with observable behavioral consequences (Fuchs et al., 2007). Interestingly, some of these inhibitory neurons are not interneurons after all, as it has been discovered that they can project long-distance, again with intriguing and yet-to-be-elucidated functional implications (Melzer et al., 2012).

Despite these similarities in the organization of the hippocampus in mammals, there are also important differences that are not expressed by simple binary parameters, as in the syntax of natural languages, but by scalar parameters that can sometimes vary by orders of magnitude. Most striking, of course, are the differences in size, both overall and in the various subregions (Seress, 1988). For example, the total number of principal cells in a rat hippocampus (on one side) has been estimated as 1.2×10^6 DG granule cells, 2.2×10^5 CA3/2 pyramidal cells and 4.0×10^5 CA1 pyramidal cells in rats, with limited variation with age (Rapp & Gallagher, 1996) vs. in humans 1.5×10^7 DG granule cells, 2.7×10^6 CA3/2 pyramidal cells and 1.6×10^7 CA1 pyramidal cells, with a major decrease with age in CA1, which is clearly also the region scaled up supra-linearly relative to the hippocampus overall (West & Gundersen, 1990). Independent scaling relations for the different regions, and for the hippocampus relative to the cortex, have

been observed also by comparing across more species. Interestingly, a recent study (Watson et al., 2024) estimates that within CA3, which itself scales up sublinearly with overall hippocampal size from mice to rats to humans, the number of recurrent inputs onto CA3 pyramidal cells sees only a modest increase (from 13,200 spines/cell in mice to 14,300 in rats to 17,500 in humans) with respect to the number of pyramidal cells (reported as 110,000 to 300,000 to 1,800,000 in their study). This may reflect a certain biophysical limit on the number of inputs that can be integrated in a pyramidal cell, while maintaining their individuality within an electrotonically compact cellular structure.

3.1.1 Do non-mammals have a Dentate?

The medial pallium of lizards presents a simpler anatomical structure in cross section as compared to birds and mammals (Figure 4) . Most cell bodies are contained in a single layer, with dendrites above and below. Still, there are differences within that layer, particularly in size, between a large-celled dorso-medial portion and a mostly small-celled medial portion (Day et al., 1999). The latter appears to be developmentally homologous to the dentate gyrus in mammals, whose principal cells, the granule cells, are also smaller than the pyramidal cells of CA3 and CA1. Further, if one looks at more rostral sections of the hippocampus of mammals like the opossum, one sees at the most anterior end an undifferentiated structure similar to that of a lizard; only in more posterior sections one recognizes the mammalian organization, with the dentate gyrus originating from the most medial portion, ventral to the hippocampal fissure (Hamel, 1966). Thus, in this sense one can say that also reptiles have a dentate gyrus. The situation is more complicated, however (Reiter, Liaw, Yamawaki, Naumann, & Laurent, 2017), and there are differences among reptilian lineages. In turtles and crocodiles, for instance, anatomists distinguish not 2 but 3 zones (zone 4, 3 and 2, going from the dorsal towards the medial end). The connectivity patterns are what is important from a network perspective, and they are only partially known. Still, there is no report of anything as exclusive as the MF system in mammals. Ultimately, as discussed by Striedter (2016), the question of homology may not have a yes or no answer: there may be features of parts of the reptilian hippocampus which look dentate-like (the developmental trajectory, the presence of zinc, the salient adult neurogenesis) without having a real parallelism at the circuit level.

An analogous situation prevents establishing a clear homology in the avian hippocampus, with the difference that the circuitry in birds appears to be more complex (Herold et al., 2019). One distinguishes a dorso-lateral, a dorso-medial and a ventral (which is also kind of V-shaped) division, but cell bodies in all three are largely scattered over the thickness of the cortex. There may be organizational principles not yet understood, but if so they are likely different from those of the mammalian hippocampus. Although studies of genetic mechanisms, such as those necessary for dentate gyrus development (Mercurio et al., 2021) may shed new light and help clarify e.g. whether the ventral region is more DG-like (Atoji & Wild, 2004) or more CA1-like (Kahn, Hough II, Ten Eyck, & Bingman, 2003), with reptiles similar transcriptomics does not seem to translate into the same circuits (Tosches et al., 2018). In the end, one has to accept that also genetic expression mechanisms may settle into distinct attractor states, which

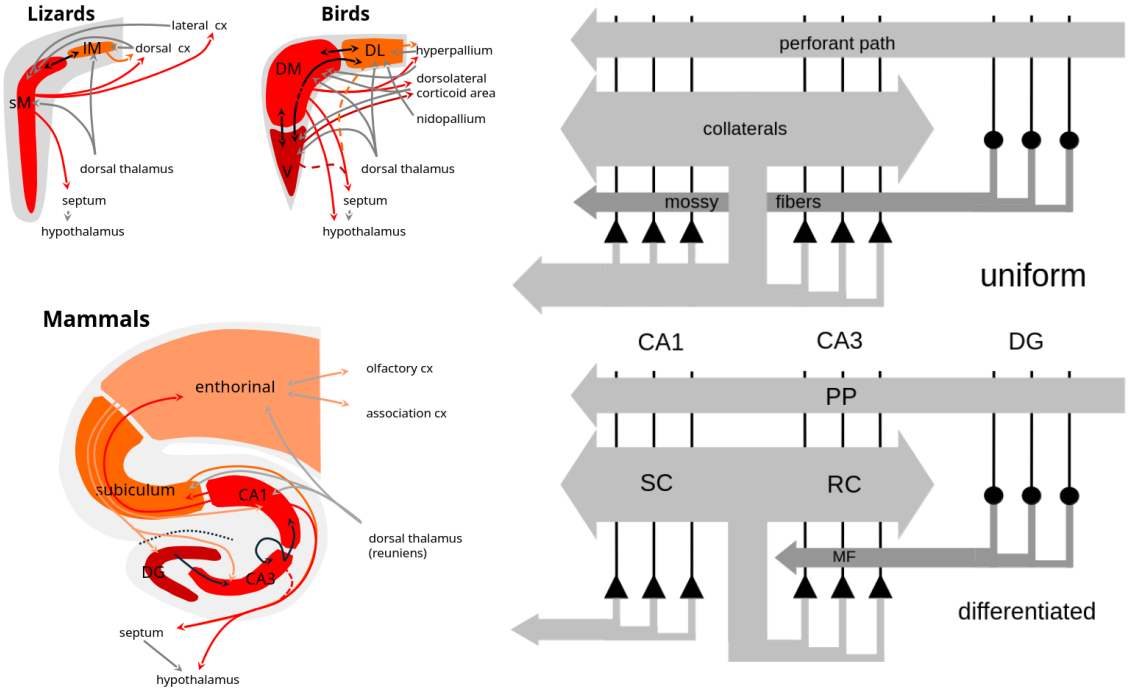

Figure 4: Left – The different internal design of the hippocampus among amniotes implies different connectivity patterns (redrawn from Striedter, 2016). Right – In computational models, often only some of the many connections are included, and even then only in schematic form. To understand the significance of a feature of mammalian design like the CA3-CA1 differentiation it is expedient to compare a simplified model of the mammalian circuit (bottom) with an imaginary uniform (i.e., undifferentiated) circuit (top) quantitatively equal in all other respects, even though it would not correspond to that seen in any existing species (compare to left side; redrawn from Treves, 2004).

cannot be forced to show exact homologies. For all practical purposes, therefore, the dentate gyrus is a mammalian 'invention' (Treves, Tashiro, Witter, & Moser, 2008).

3.1.2 Differentiation between CA3 and CA1

Like the duplicated cortical input to CA3, the original and the one 'translated' by the DG, so the difference in connectivity between CA3 and CA1 appears like a distinctly mammalian riddle. In crude terms, CA3 and CA1 pyramidal cells share the majority of their synaptic inputs, but it is CA3 that provides them all, the breadwinner in the family (see Figure 4). The axonal fibers that feed other CA3 cells are called recurrent collaterals (RC) and those that feed CA1 cells are called Schaffer collaterals (SC), but they are just branches of the same CA3 axons. Their direct cortical inputs, instead, appear to be largely separated, with those impinging onto the apical dendrites of CA3 cells, in stratum lacunosum moleculare (SM) coming mainly from entorhinal cortex layer II, and those to CA1 from layer III, at least in those species where they have been studied (Witter, 1993). Although some recurrent collaterals have been reported also in CA1 (Thomson & Radpour, 1991), they have been described as directionally

oriented towards the border of CA1 with the subiculum (the next subfield at the output of the hippocampus, see Figure 1), as if being directed there, and perhaps making a few synapses 'en passant' onto fellow CA1 cells (Orman, Von Gizycki, Lytton, & Stewart, 2008). A study which observed them in newborn rats but not in adult ones argued that they may have a developmental role in synaptogenesis (Aniksztejn, Demarque, Morozov, Ben-Ari, & Represa, 2001). The prevailing consensus is that CA3 recurrent collaterals are in any case much more abundant. Such a salient difference obviously has a genetic basis, and different patterns of gene expression have been reported, also in humans (Ginsberg & Che, 2005).

Besides the scarcity of recurrent collaterals, the other salient connectivity difference of CA1 from CA3 is that CA1 cells do not receive MF projections from DG. Thus neural computation models of the hippocampus, which mostly ignore the details of inhibitory circuits as well as other afferent and efferent connections (and usually also the subiculum and neighboring regions, at the exit from the hippocampal loop) typically focus on three networks, with distinct connectivity. DG receives PP inputs from EC layer II. CA3 receives the same PP inputs, plus an abundance of recurrent connections, plus the MF from DG. CA1 receives the SC from CA3, and its own PP inputs, from EC layer III (Figure 4, bottom right).

Except for CA2. De Nó (1934) had subdivided Cornu Ammonis into fields CA1, CA2, CA3 and CA4. While the latter is now generally regarded as the hilar part of the DG between the two prongs of the granule cells and the term CA4 has largely fallen out of use, the CA2 field, a relatively narrow strip of tissue between CA1 and CA3, has experienced a major revival in recent years. Although many scholars view it as a mere transition region between the two more distinct neighbors, or as a sort of incompletely expressed part of CA3 which does not receive MF inputs, others have reported a number of specific features (Jones & McHugh, 2011), a clearer definition of its boundaries using gene expression patterns (Lein, Callaway, Albright, & Gage, 2005) and a prominent role in social memory (Hitti & Siegelbaum, 2014). In humans, CA2 appears larger and more distinct than in rodents (Knowles, 1992), suggesting that it could be included not among the 'principles' of mammalian hippocampal organization, but rather among the 'parameters'. Moreover, genetic expression patterns tend to produce way too complex a description to be incorporated in any viable model (Thompson et al., 2008).

Whatever the case for CA2 in humans, evidence for a strong anatomical differentiation in other vertebrates is missing also between CA1 and CA3. In terms of neural activity, surprisingly, also in mammals, even in rodents, which account for the vast majority of recording experiments, CA3 and CA1 appeared to show rather similar phenomenology. Relatively minor quantitative differences had been reported in that CA1 pyramidal cells (morphologically a bit smaller on average) tend to be more active: e.g. in rats moving around in an environment of limited size, the fraction of CA1 cells with at least one place field could be 40% rather than 25% for CA3, their place specificity in an 8-arm maze lower and their mean firing rates marginally higher (Barnes, McNaughton, Mizumori, Leonard, & Lin, 1990). Nothing comparable to the major connectivity difference. This changed in 2004, with the discovery of a major qualitative difference in how CA3 and CA1 express spatial activity in similar environments (Leutgeb, Leutgeb, Treves, Moser, & Moser, 2004). In CA1, the similarity between two environments is reflected in a corresponding degree of similarity between the populations of cells active

in each; whereas in CA3 it appears that as soon as two environments are distinguishable, the two ensembles of active cells are totally unrelated to each other, with just a chance overlap. In complementary experiments by other labs, it was found that when the differences between environments are quite small, CA3 activity actually remains more coherent between the two (Lee, Yoganarasimha, Rao, & Knierim, 2004), so that one can describe, in CA3, more strongly non-linear dynamics than in CA1, which is indicative of attractor networks (Vazdarjanova & Guzowski, 2004; Guzowski, Knierim, & Moser, 2004). In summary, one could say that in 2004 experimental findings aligned with the expectation of Marr's 'collateral effect' in CA3, a point to be taken up below.

3.1.3 Interspecies variability, scaling and deviations

The wealth of results obtained by focusing on a few species, mostly of rodents, should not obscure the fact that even within the common mammalian design there is substantial variability in the details—some of which may not be minor after all. The different scaling relations between hippocampal subfields in relation to brain size alone lead to remarkable differences in design: CA3, for example, is a clear bottleneck between DG and CA1 in humans, in terms of number of principal cells; but not so clearly in species with a smaller brain, where the CA3 and CA1 populations may be similar in number, or even in an inverted relation. In highveld and naked mole-rats it has been estimated that there are more CA3 pyramidal cells than CA1 ones, and in cape mole-rats they are even more than DG granule cells (Van Dijk, Huang, Slomianka, & Amrein, 2016). Instead of a bottleneck, CA3 would be placed as a sort of expansion chamber! African mole-rats may be dismissed as local weirdos, but a large variability in numbers has been reported in other species as well, with e.g. the ratio of principal cells in CA1 to those in CA3 estimated to be 3.9 in white rabbits down to 1.6 in golden jackals, within hippocampi of comparable size (Malikovi ć et al., 2023). And this without considering the reflected blade of CA3. The reflected blade (also called CA3h) is a population of pyramidal cells inserted deep into the dentate gyrus, so that they have no access to the stratum lacunosum moleculare, where CA3 and CA1 pyramidal cells mostly receive PP inputs from EC. Included with the hilar polymorphic (non-pyramidal) cells in what Lorente de Nó had called CA4, it is a population that does not exist or reduces to negligible numbers in laboratory rats and mice, so that it has been largely ignored. In other species however it does exist; and in the wild boar and red and roe deer it has been estimated to even include roughly 2/3 of the pyramidal cells in CA3 proper (Malikovi ć et al., 2023). We are not aware of any computational model of the hippocampus that considers the reflected blade, although in terms of connectivity it appears to differ from both DG and CA3 proper, so that it would seem a gross oversight to group it with either of the two neighboring regions.

Hippocampal dynamics, in particular the prevailing frequencies in the power spectrum of the field potentials recorded during various behaviors, also comprise a domain where we are prone to confuse principles with parameters. A striking phenomenology has been described around hippocampal theta oscillations in rodents, part of a bewildering variety of rhythmic brain dynamics (Buzsáki, 2006). The discovery of a relationship between theta phase and spatial location in rats running along a track (O'Keefe & Recce, 1993) has stimulated the development of a rich set of ideas and

theories (Burgess & O'Keefe, 2011), which have often been considered applicable to species in which theta-band oscillations are much less prominent (Lisman & Jensen, 2013; Zheng et al., 2024), perhaps because they are not resonating with the sniffing cycle as they tend to do in rodents (Macrides, Eichenbaum, & Forbes, 1982; Niedermeyer, 2008). Clearly, the relevance of these oscillations has to be discussed case by case, for example by analyzing corpora of data where oscillating bouts have been removed, as done in a bat study (Yartsev, Witter, & Ulanovsky, 2011). Recent evidence suggests that even in mice the distribution across cells of the theta-related dynamical patterns may be different from what is envisaged by mainstream theories (Guardamagna, Stella, & Battaglia, 2023).

A final and rather obvious comment on the variability of hippocampal design across mammalian species is that it must, to some degree, reflect the distinct natural behavior of each species. Bats fly, rodents run around, and primates often prefer to explore the surrounding space visually rather than by walking to it—a point made early on by Rolls (1999) and supported by the discovery in his lab of spatial view cells. In humans, the reliance on visual scene analysis for memory encoding may have led to the development or strengthening of a ventromedial 'where' stream to the hippocampus (Rolls et al., 2024). An intriguing computational study shows how the different statistics of behavior may lead, within similar circuitry, to neural activity that is selective for different correlates (Franzius, Sprekeler, & Wiskott, 2007).

In summary, the principles of mammalian hippocampal design outlined in this section should be taken with a grain of salt, as a well-intentioned abstraction. They are heavily biased by the few species normally used in the lab, mainly rodents. Still, with this qualification, there is ground to explore the relation between structure and function that had fascinated the young David Marr (as well as many other scientists), although perhaps at a level of detail not as precise as he had hoped to reach.

3.2 CA3 as an auto-associative memory network

David Marr (1971) proposed that the hippocampus—which he called more comprehensively 'archicortex'—serves as a 'simple' memory device. He meant a content-addressable memory, that operates just by pattern completion: it encodes patterns of activity (which in his model are binary strings, where the 1's are active cells and the 0's are quiescent ones) and when presented with a partial cue that univocally identifies one of the encoded patterns (its active units are a subset of those in that pattern) it reactivates the entire pattern. The vanilla version is when the cue and the reactivated pattern are instantiated in the same population of cells—Marr calls it a 'free' simple memory. The reactivation then proceeds through recurrent connections that have been associatively modified, that is, strengthened between cells i and j that were both active in the encoded pattern. One can also consider, however, a 'directed' simple memory, in which the same event is represented twice, by pattern of activity $\mathcal{A}$ in population A and by pattern of activity $\mathcal{B}$ in population B. The cue can be a subset of $\mathcal{A}$ in population A and reactivate $\mathcal{B}$ through associatively modified connections from A to B.

In either case, the crucial components of the network are the synaptic connections between pyramidal cells, which Marr takes to be individually associatively modifiable according to the scheme proposed by his advisor Giles Brindley (1967). Note that the

earliest description of the discovery of long-term potentiation (LTP; Lømo, 1971) came out the same year as Marr's paper, where it is mentioned as a note added in proof; and the discovery was in the Dentate Gyrus, not in Cornu Ammonis. Marr instead focused his proposed auto-associative mechanism, whereby a partial or incomplete stimulus $\mathcal{E}'$ can reactivate the representation of event $\mathcal{E}$, on the collaterals of Cornu Ammonis, and called it the 'collateral effect'. He was less committed to pinpoint exactly which collaterals. In fact, combining his 'directed' and 'free' models, he referred both to the Schaffer collaterals (from CA3 to CA1) and to the recurrent collaterals 'of CA1 and CA2' which are now known to be scarce to non-existent. He cited a study (Raisman, Cowan, & Powell, 1965) which had made the point that the SC from CA3 to CA1 are NOT reciprocated from CA1 to CA3, and probably mis-interpreted it as stating that local recurrent collaterals (which their axonal degeneration technique was not suited to reveal) where more abundant in CA1/CA2 than in CA3. In any case, his network model of the collateral effect was visionary, and the fact that it was not followed up by others for a decade and a half, apart from Gardner-Medwin (1976), probably has more to do with the general demise of neural network studies in the 1970's, following the publication of the book 'Perceptrons' (Minsky & Papert, 1969), than with its rudimentary, binary character.

In the early 1980s, John Hopfield developed his recurrent auto-associative model based on binary model neurons and strictly symmetric but graded rather than binary synaptic weights (Hopfield, 1982), and two years later, a follow-up model which replaces binary with somewhat graded, sigmoid neurons, and has essentially the same properties (Hopfield, 1984). The former could be rigorously analyzed in a tour de force utilizing techniques from the statistical physics of spin glasses (Amit, Gutfreund, & Sompolinsky, 1987). Soon after also more neurally plausible versions of the Hopfield model were considered (Tsodyks & Feigelman, 1988; Treves & Amit, 1989; Treves, 1990), largely spurred by criticism from Moshe Abeles (1991) and other neurophysiologists in Jerusalem, who were interested in the neural plausibility of the model but not, however, particularly in the hippocampus. The hippocampus 'came back' with the proposal that it is particularly the CA3 region which operates, through its by-then-recognized abundant recurrent collaterals, as an auto-associative network (McNaughton & Morris, 1987; Rolls, 1987).

Like the Marr model, the Hopfield model network simplifies many of the biological details in order to better illustrate macroscopic network properties, but now also to enable the application of concepts and sophisticated mathematical tools from statistical physics, which Marr had no access to. Most important is the assumption of symmetrical connection weights between pyramidal cells ($w_{ij} = w_{ji}$, both of which can also take an implausible negative value), which makes the propagation of activity through the recurrent connections resemble the dissipative dynamics of a disordered system, such as a spin glass that converges towards one of its free-energy minima. Inhibition is not described at all, and the afferent inputs, including e.g. those conveying the cue that has to reactivate one of the stored memory patterns, are typically considered to subside after setting the initial state of the network for its memory retrieval dynamics. The network is taken to have stored such patterns of activity during a learning phase that is not explicitly described but that follows a simple learning rule according to Hebb's principle (Hebb, 1949). If, at retrieval, the network has relaxed to a free-energy

minimum that coincides or is close to the original representation of the true event, $\mathcal{E}$ in Marr's notation, the process can be called in connectionist terms 'pattern completion'.

While statistical physicists had been attracted to the Hopfield model as if it had been the ultimate mode of neural computation, at least insofar as the function was storing and retrieving memories, McNaughton and Morris (1987) and Rolls (1987) provided concrete evidence that memory circuits could be more articulated, by pointing out that within the hippocampus itself there are other subfields—in particular the Dentate Gyrus and CA1—with rather different organization from CA3. McNaughton and Morris (1987) illustrated Marr's theory with small descriptive toy models that could be used to obtain insight into the working of each subfield, and proposed that the auto-associative toy model of CA3 could also serve to store and retrieve sequences of memory items. Rolls (1987, 1989) related these theoretical notions to experimental findings in monkeys, and helped to de-emphasize the spatial character of hippocampal memories, that is so salient in rodents.

3.2.1 Auto-associative memory capacity

Marr (1971) hypothesized that events cannot be stored in memory faster than once per second. Together with his idea that the hippocampus transfers memories to the neocortex overnight, it follows that the upper limit of memories that the hippocampus should be able to store lies at around 10^5 memories, a number approximately equal to the number of seconds per day. Although this is just an order of magnitude calculation and it is based on probably mistaken assumptions, such as the clean overnight sweep of hippocampal memories to the cortex, Marr could use it to derive constraints on his rudimentary binary model. We shall see that it reasonably fits with the storage capacity calculated rigorously for neurally plausible versions of the Hopfield model, which is in principle quite different from that of Marr's model.

After Amit et al. (1987) had shown the way, calculating the storage capacity of different variants of the Hopfield model using the methods of statistical physics became a sort of coming-of-age exercise. In the original version, in which each of N pyramidal cells is synaptically connected to all others, and is active or quiescent with equal probability in each memory pattern, the maximum number of stored patterns which can be individually retrieved is $p_{\max} \simeq 0.14 \times N$. This confirms the estimate based on numerical simulations in Hopfield (1982). After considering various improvements to the original version of the model to bring it closer to a real cortical network, whether in CA3 or elsewhere, it was realized that the crucial changes are two. First, if the connectivity is not all-to-all, the average number of connections per cell C^{RC} (from other pyramidal cells in the same population; that is, the number of distinct recurrent collaterals a cell receives) has to replace N. Second, if much less than half the cells are active in the representation of each memory, as parametrized by the sparsity a, roughly the fraction of active cells, then $p_{\max}$ gets larger the smaller a is—more precisely it scales inversely to $a \ln(1/a)$. In a formula,

$$p_{\max} \simeq k \frac{C^{\mathrm{RC}}}{a \ln\left(\frac{1}{a}\right)}, \tag{1}$$

where the term k is factor that depends weakly on the detailed structure of the rate distribution and neural connectivity (typically between 0.2 and 0.3, Treves and Rolls,

1991), and where the sparsity parameter a, which for a binary activity pattern is simply the fraction of active cells, can in general be measured from the firing rate r_i of each neuron i in the population as

$$a = \left(\sum_{i=1}^{N} \frac{r_i}{N} \right)^2 \Big/ \sum_{i=1}^{N} \frac{r_i^2}{N}. \tag{2}$$

The sparsity a ranges from $1/N$, when only one of the neurons represents a certain memory, to 1.0 when all neurons are participating in the representation of a memory with equal rate. In CA3, recordings from rats running in an open arena have produced values of a in the range $0.02 - 0.06$ (Papp & Treves, 2008). A standard estimate for the rat C^{RC} is ca. $12,000$ associatively modifiable recurrent collateral synapses onto each neuron (Amaral et al., 1990) yielding with a sparseness of 0.1 or 0.02 a capacity of $12,000$ or $36,000$ memories, respectively. The number of collateral synapses for humans has been hitherto unclear, but a recent count gives 17,500 synapses per cell, only a little more than in rats (Watson et al., 2024). Assuming the sparsity parameter to be at the lower end of the rat range would yield a human CA3 capacity roughly half of the number David Marr had hypothesized to be required.

The above estimates are crude but convey two important take-home messages, plus an equally important qualification. First, what matters for storage capacity is not so much the size of the network (how many pyramidal cells it includes) but rather its connectivity (the number of recurrent collaterals per cell with independently modifiable synaptic weights). Second, the capacity is larger with sparser representations, which may indirectly favor larger networks. They can distribute the same information onto a smaller fraction of its cells, but with an attenuation imposed by the logarithmic co-factors in Equation 1. Third, the analytical derivation of equations like Equation 1 requires an assumption that discrete, point-like memories are distributed randomly in the $(\mathbb{R}^+)^N$ space of all possible activity patterns (technically, the joint distribution of the probability of activity levels $P(\{r_i\})$ across cells factorizes into $\Pi_i P(r_i)$). This assumption is violated by memories that are correlated, e.g. because they contain common elements and have not been recoded to remove the resulting correlation, or because they are not point-like, e.g. because they represent locations in a spatial continuum. Computer simulations can supplant analytical derivations when the latter are not feasible, but of course the variety of correlation types one may want to consider is infinite. Two cases of particular interest are discussed below in subsubsection 3.2.3 and 3.2.4, and in connection with the role of the Dentate Gyrus in subsection 3.3.

3.2.2 Auto-associative retrieval times

How long does it take before a pattern of activity triggered in CA3 by afferent inputs gets to be affected, or 'completed', by the activation of the recurrent collaterals, or even to settle into a stationary state? Masking experiments in monkeys and humans have shown that, in order to recognize a visual stimulus, neurons in higher cortical areas require time windows of as little as 20 ms (Rolls, Tovee, Purcell, Stewart, & Azzopardi, 1994; Rolls & Tovee, 1994), and to observe the signature of a categorization process in event-related potentials may take overall as little as 150 ms (Thorpe, Fize, & Marlot, 1996; Antal, Kéri, Kovács, Janka, & Benedek, 2000). Conceiving such a process as a

serial process in a computer, running sequentially across a series of cortical stations, it was argued that each station must complete its work in 20–30 ms at most, leaving no time for reverberations via recurrent collaterals in any of the stations: they must operate rapidly as feed-forward networks (Thorpe & Imbert, 1989). The hippocampus sits at the top of the visual processing hierarchy and, if its output serves memory rather than immediate behavior, it can perhaps act more relaxed and pattern-complete more slowly. But the argument was useful in raising a general issue about recurrent processing relevant also for the hippocampus and its CA3 recurrent network: What is a reverberation, really? Do temporal constraints limit the role of reverberating activity in the brain?

The early, physics-style computational implementations of autoassociative networks did not pay much attention to neural dynamics, and typically assumed that time is discretized into time steps. Model neurons would change their firing rate once per time step, either all together or in a fixed or variable sequence. Considering that it apparently took 4–5 time steps to substantially complete a pattern, the key question seemed to be: What actually is a time step, and how long is it in the real brain?

It took several years to realize that such questions can only be addressed with the help of models that contain a reasonably plausible description of neural dynamics, in particular models that break down the firing rate description of neuronal output in terms of its constituent elements, the individual action potentials or 'spikes'. A mean-field analysis of an autoassociative network of spiking model neurons showed that the dynamic convergence to a stationary state is very fast and takes about 20 ms, a time scale largely independent of the prevailing firing rates and membrane time constants of the neurons, and is essentially determined by the speed of synaptic dynamics Treves (1993).

The main limitation of the analysis was that it is based on a mean-field approximation, and two years later it was proposed that the convergence to an attractor state, i.e. the process of pattern completion or memory retrieval, can be driven by fluctuations and is even faster than predicted by mean-field theory (Tsodyks & Sejnowski, 1995b). Interestingly, computer simulations of a spiking recurrent network confirmed both the mean-field prediction and the suspicion that it may not apply because of the effect of fluctuations (Battaglia & Treves, 1998b). As shown in Figure 5, when an afferent input conveys a partial cue, the recurrent network dynamically selects the stored memory pattern best correlated with the cue, approaching it with a time constant proportional to the time constant of excitatory synapses. However, if the cue does not coincide with the full memory pattern, it keeps the network activity away from it as long as it is sustained; once it subsides, convergence to the full pattern is essentially immediate, driven by self-reinforcing fluctuations. In practice in the real brain, cue onset and removal are not sharply defined and it is highly unlikely that one may be able to dynamically distinguish the two experimentally as in the simulations shown in Figure 5. Still, the upshot is that local reverberations can exert their effect rapidly, even if they were subject to a strictly serial processing schedule—an improbable notion in neuroscience. They are definitely able to make their contribution within relevant cognitive time scales, particularly in memory.

These results apply to patterns of activity distributed over sufficiently large numbers of active cells, which spike at sufficiently asynchronous times, given sufficient recurrent

connectivity. This is important: because of asynchrony, any spike emitted by a cell that happened to be close to threshold can almost immediately (after a negligible 1–2 ms synaptic delay) influence the firing of those postsynaptic cells that it finds also close to threshold, advancing those with which it has potentiated synaptic weights. The consequence is a dynamical cascade, an avalanche, as a result of which much of the processing may be completed even before a large number of cells have had time to emit a single action potential—the distributed code can do without them. If all active cells were firing in strict synchrony, instead, within say 1 ms of each other, that influence would have to wait for the next activation wave. In general, neural network models that operate in discrete time steps or 'iteration cycles' are illustrative, but can lead to misleading arguments regarding the asynchronous dynamics of population activity (Treves, Rolls, & Tovee, 1996).

Applied to CA3, these results suggest that the activation of its recurrent collaterals, which have synapses with relatively short time constants of a few milliseconds, is able to represent and automatically complete a memory within a short time of a few tens of ms. In animals with a hippocampal theta rhythm, like rodents, when theta is present, e.g. during exploration behavior, these dynamics could occur within a theta cycle (Jezek, Henriksen, Treves, Moser, & Moser, 2011). Only a weak theta rhythm has been observed in the hippocampus of monkeys and humans (Ekstrom et al., 2005; Jutras, Fries, & Buffalo, 2013), however. During sharp wave activity, the level of synchronicity increases and the above results probably do not apply, but then successive waves of activity occur so fast that even the model based on iterative reverberations over discrete time steps would predict pattern completion within a few tens of ms.

3.2.3 Continuous attractors and remapping

The calculations above quantify the storage capacity of auto-associative networks in which discrete memories are represented by random patterns of network activity. Thus, all patterns are roughly orthogonal to each other, i.e. the Pearson correlation between their activity vectors is zero, plus or minus small chance fluctuations. As such, there is no meaningful notion of distance or similarity between these memories in the autoassociator. This scenario cannot certainly apply to all types of memories processed by the hippocampus.

Memories with a spatial component encode the position of an animal (e.g. through place cells), the orientation of its gaze (as found in monkeys through spatial view cells Rolls, 1999), or the position of a companion (as observed in bats, Omer, Maimon, Las, and Ulanovsky, 2018). Such spatial memories imply that the representations of nearby positions in space are highly correlated and that representations of distant ones are not. A way to model such correlations—taken to be the outcome of associative synaptic plasticity—is to proportion the strength of excitatory synaptic interactions between pyramidal cells to the similarity of patterns they represent, and to additionally apply a homogeneous amount of inhibition which keeps their overall activity constant (Wilson & Cowan, 1973; Amari, 1977). This local-excitatory and global-inhibitory weight structure produces quasi-continuous sets—i.e. continuous sets in the limit case—of attractors (Tsodyks & Sejnowski, 1995a; Blum & Abbott, 1996), the energy landscape of which resembles a flat riverbed towards which neural activity converges like rain-

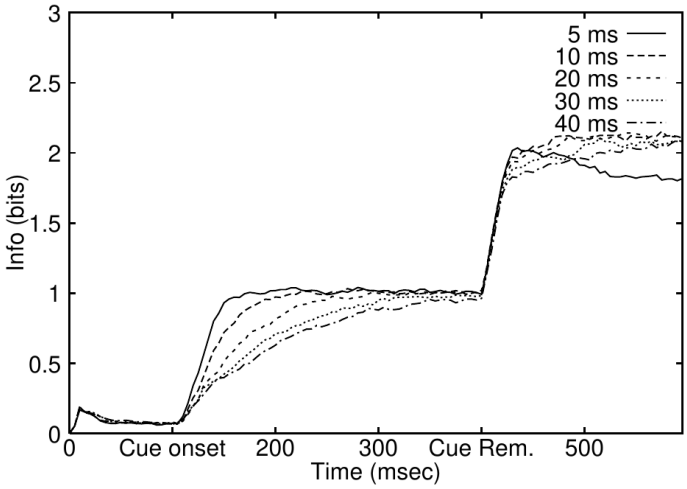

Figure 5: Information time course for different synaptic time constants, from numerical simulations of a spiking recurrent network. After cue onset, the amount of information in the pattern expressed by the network approaches the amount present in the (incomplete or corrupted) cue, with a time constant proportional to the time constant applied in the model for closing the conductance at excitatory synapses. When the cue is instantaneously removed the network is no longer held back by the cue, and a pattern completion process unfolds in a time so short that it cannot be resolved by the information measure used in the simulations. Figure adapted from Fig. 3 in Battaglia and Treves (1998b).

water from the river basin, and on whose ground a bump-like activity profile remains marginally stable, since the riverbed is flat, with no slope. Continuous attractor neural networks had been introduced in other contexts, in which they do not encode specific memories and do not need to be continuously re-established by synaptic plasticity—for example, as a model of working memory in the prefrontal cortex in oculomotor delayed response tasks with monkeys (Compte, Brunel, Goldman-Rakic, & Wang, 2000; Wimmer, Nykamp, Constantinidis, & Compte, 2014), or as a model of head-direction cells (Taube, 1995), which has been found to be relevant also for insects (Kim, Rouault, Druckmann, & Jayaraman, 2017). In those contexts it is generally assumed that there is a unique such attractor.

Since place cells are ubiquitous in CA3, at least in rodents, does this mean that CA3 at large comprises one single continuous attractor? That this is not the case, has been shown experimentally through the discovery of the phenomenon of *remapping* (Muller & Kubie, 1987; Bostock, Muller, & Kubie, 1991). First, when rodents explore different boxes, CA3 place cells are often observed to have a place field in more than one box, e.g. in 25–30% of boxes of typical laboratory size. Second, when any two cells are seen to both have place fields in two different boxes, the distances between them in the two boxes are unrelated to each other. More generally, changing the box or some aspect of it in a way that is sufficient to make the animal perceive it as a different box, causes a complete rearrangement of the place fields of an entire population of CA3 cells. CA1 place cells are more likely to exhibit multiple place fields in a given spatial context, but even there it is unpredictable whether a cell shows one, multiple, or no place fields and

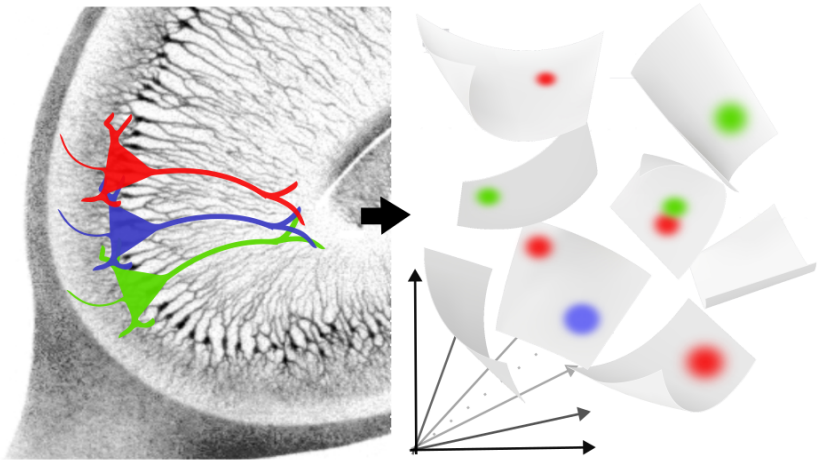

Figure 6: Three sample pyramidal cells in CA3 (left) may have place fields in several charts, the manifolds with which CA3 represents different 2D environments in the high-dimensional space of its neural activity patterns (right). Neighboring cells need not have neighboring place fields, nor place fields in the same charts; while cells far away in the neural tissue may happen to have overlapping place fields. Schematics loosely based on Alme et al. (2014). Flying charts design by macrovector/Freepik.

at which location(s), knowing its place field(s) in other spatial contexts or the activity of other cells with nearby place fields in those other contexts.

It was later understood that one has to distinguish between the global remapping described above, and what has been called 'rate' or 'local' remapping (Leutgeb et al., 2005). While place cell activity is usually globally remapped when the animal encounters an entirely new environment, it is locally remapped when place fields remain where they are, but change their peak firing rate—sometimes to such an extent that a field appears where it was not seen before, or another one disappears. Rate remapping usually occurs when only some parts of the environment have changed (e.g. some landmarks; Anderson and Jeffery, 2003; Leutgeb et al., 2005; see also Baraduc, Duhamel, and Wirth (2019) for evidence in macaques) and it can be taken to reflect the role of other, non-spatial correlates, in determining the firing rate of the cells in each of their fields. Sometimes the distinction is not so clear cut, reminding us that 'pure' place cells may be just a Platonic idea.

In any case, it is clear that CA3 can keep in memory multiple continuous attractors simultaneously, and reactivates each, at different times, by selecting from the same overall population the particular mix of cells that represents the same or a similar place in the same context, as illustrated in Figure 6. Samsonovich and McNaughton (1997) introduced a formal network model in which these representations are called 'charts'. Since neighboring place cells in the brain do not represent neighboring locations outside the brain, a chart can be thought of as an imaginary arrangement of place cells on a plane such that each cell represents the location of its highest activity—at least if no cell has more than one field per chart. Within a chart, the actual spatial position of an animal is represented as a bump of neural population activity moving on the plane. The idea is that multiple charts are stored in CA3 but only one of them is active at a

given time, thereby solving the problem of cells with multiple place fields but in distinct environments (Muller & Kubie, 1987; Kubie & Muller, 1991).

What is the memory capacity for multiple charts? The memory capacity should be highest if one assumes that different charts are uncorrelated with each other, a hypothesis that has been confirmed in CA3 (but not in CA1) in a study mentioned above (Leutgeb et al., 2004), and further validated in CA3 by experimental recordings of rat place cell activity in as many as 11 different environments (Alme et al., 2014). This latter study also empirically demonstrates a storage capacity of at least 11 charts in rats. Still, 11 is 3 orders of magnitude below 12,000, the lowest estimate of the capacity for discrete memories. In fact, the capacity for multiple charts in a continuous attractor network is necessarily lower than that for discrete point-like memories in a classical attractor network, because neurons have to represent a continuum of different positions on each chart, with some stability. How much lower was calculated by Battaglia and Treves (1998a) for a model in which each neuron participates in each chart with a place field of constant size that covers roughly a fraction a of the stored environment. Then, the memory load of each chart is roughly equivalent to that of $1/a$ discrete point-like memories, and the storage capacity, expressed as the maximum number of charts that can be retrieved, is expressed by a formula similar to Equation 1:

$$p_{\max}^{\mathrm{CHARTS}} \simeq k' \frac{C^{\mathrm{RC}}}{\ln\left(\frac{1}{a}\right)}, \tag{3}$$

without the $1/a$ factor that was giving the increase for sparse point-like patterns, and where the numerical factor k' is reduced for 1D charts and further reduced for 2D ones. The result is visualized in Figure 7. The prediction is that the rat CA3 should be able to retrieve 100–300 distinct charts. Note, however, that a key assumption behind this calculation—that each neuron contributes with a place field of constant size and peak firing rate—is incompatible with recent experimental observations (Eliav et al., 2021), an issue we further discuss below.

3.2.4 Sequence learning

The attractor networks discussed so far simplify the dynamical nature of episodic memories by treating them as discrete events, represented either by a point-like configuration of neural activity unrelated to the representation of other events, or as events situated in a spatial continuum where events are only correlated with other events situated nearby. Human episodic memories, however, are naturally conceived to unfold along a temporal continuum, of variable duration, with a succession of few or many 'phases' of the episode, which could be a fleeting moment or an entire story. One somewhat related phenomenon in rodents, which is experimentally accessible at the neural level, is that of replay (Louie & Wilson, 2001) and preplay (Dragoi & Tonegawa, 2011; Pfeiffer & Foster, 2013). Replay is when a continuous succession of spike patterns that occurred during recent experience is reactivated (usually in the same order, but not always) during subsequent periods of rest or sleep. Preplay, on the other hand, is said to occur when rest or sleep activity appears to predict the future flow of spike patterns in the awake animal. While the studies above were based on recordings in CA1, one can ask how such dynamical flows might be stored in CA3. Early formulations, starting with McNaughton and

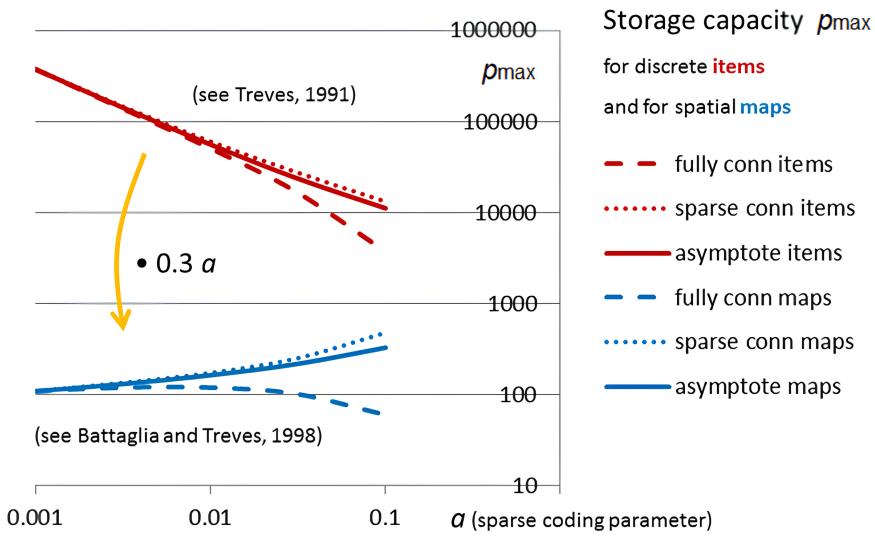

Figure 7: The storage capacity, $p_{\max}^{\text{CHARTS}}$, for multiple continuous attractors (charts) in blue is contrasted with the storage capacity for discrete attractors (patterns) in red as a function of sparseness on a log-log scale. The former can be roughly derived from the latter by multiplication with $0.3a$. Three distinct connectivity models—fully connected ($C = N - 1$), sparse ($C \ll N$), and a plausible intermediate estimate (solid curves)—show major differences only for non-sparse memories where $a \gtrsim 0.03a$. The plot uses a slightly more conservative value of $C^{RC} = 10,000$ in a rodent model, rather than $C^{RC} = 12,000$ as mentioned in subsubsection 3.2.1.

Morris (1987), considered rather than a continuous flow the stitching together of 'discontiguous' subevents in a discrete sequence (Granger, Whitson, Larson, & Lynch, 1994; Wallenstein, Hasselmo, & Eichenbaum, 1998), and sometimes even relied on putative reverberations from CA3 back to DG (Lisman, 1999). One way to model such patched sequences is to use an additional set of temporally asymmetric synaptic weights. A simple model by Sompolinsky and Kanter (1986) combines symmetric-autoassociative weights as in the Hopfield model with asymmetric-heteroassociative weights operating at a slower time scale. The former allow the model to relax to a quasi-stable fixed point, that represents one stage in the temporal sequence, while the latter bring it out of that fixed point and onto a different one, that is the next stage in the sequence (see also Kleinfeld, 1986).

Zhang (1996) was the first to propose a continuous variant of this idea by introducing a ring attractor network with symmetric and asymmetric weight components, which can both be learned with standard 'Hebbian' plasticity. It encodes a laterally moving, stable bump of activity when the asymmetric weights are proportional to the derivative of the symmetric components. The magnitude of the asymmetric weights is then simply the velocity of the bump movement. The network had been proposed as a model of head-direction cells which have been found in many places outside of the hippocampus (Taube, Muller, & Ranck, 1990; Taube, 1995; Robertson, Rolls, Georges-François, & Panzeri, 1999), but the uni-dimensional ring topology can apply to memory networks within the hippocampus as well, for example to rodents running in a linear track (Rich, Liaw, & Lee, 2014) or bats flying through a tunnel (Eliav et al., 2021).

Others models make use of spike-frequency adaptation rather than asymmetric weights to move the bump of activity in 2D (Treves, 2004), or in 1D (Hopfield, 2010; Itskov, Curto, Pastalkova, & Buzsáki, 2011; Azizi, Wiskott, & Cheng, 2013). If one assumes that the time constant of activity decay via adaptation is much longer than the time constant of information integration (i.e. the interspike interval), the speed of the bump is proportional to the timescale of adaptation.

Adapting Zhang et al.'s (1996) concept, Spalla, Cornacchia, and Treves (2021) have found analytically, that the retrieval capacity limit of a network with stored dynamical attractors through associative plasticity (i.e., with asymmetric weights) can be larger than the capacity limit of discrete attractor nets: the factor $0.3a$ reducing the capacity of marginally stable continuous attractor networks (Figure 7 and Battaglia and Treves, 1998a) does not apply, since network dynamics converge towards dynamical attractors as entire entities, not towards specific positions along their 1D attractive manifolds. Thus, there is one asymptotically stable fixed point per manifold, similar to the discrete model. The surprising result is that including time does not reduce capacity, but quite the opposite.

All of these models, however, except the one in 2D, follow a predefined flow along a (short) learned episode, but omit what could happen at decision points, for instance when two episodic memories overlap in part, before branching off to a different continuation. Hence, they cannot model the spontaneous dynamics across an unlearned territory. We just mention here an alternative framework, the so-called 'latching dynamics' which models spontaneous network trajectories (Treves, 2005).

3.3 The Dentate Gyrus as a random number generator

The dentate gyrus sits at the front end of the the trisynaptic circuit and contains in most mammalian species, though not in all (Van Dijk et al., 2016), many more granule cells than there are pyramidal cells in CA3, particularly in wild-living species it seems (Amrein, Slomianka, & Lipp, 2004). Although their ratio is nowhere near the huge ratio of granule cells to Purkinje cells in the cerebellum, it has inspired the notion that also in the hippocampus the distribution and sparsification of activity over a large number of input cells serves to separate out patterns of activity that are overlapping in the afferent inputs (Borzello et al., 2023). This is called 'expansion recoding', and can lead to pattern separation as a consequence of sparseness. Sparse activity in the expanded representation facilitates, for example, having a few cells that remain above a high threshold code for A+B, distinct from the few cells that code for A and the few that code for B. Expansion recoding does not require synaptic plasticity, and such a recoding stage has been inserted in connectionist feed-forward models arguing that it helps with pattern separation. Again, similarly to cerebellar granule cells, though less radical, dentate granule cells are not just numerous but also small (Amaral, Scharfman, & Lavenex, 2008; Rogers Flattery et al., 2020) and have no recurrent collaterals, which fits in a connectionist framework as well. Finally, it has been observed that adult neurogenesis, which is relatively rare in the mammalian brain, does persist in the dentate gyrus: new granule cells continuously develop in adulthood and some of them survive in rodents (Altman & Das, 1965; Kuhn, Dickinson-Anson, & Gage, 1996) and also in humans (Eriksson et al., 1998; Spalding et al., 2013). After all, if more neurons are

better, why not ask for more? All this appears to point towards a simple connectionist feed-forward architecture, in the form of a linear autoencoder network for example (Wiskott, Rasch, & Kempermann, 2006), in which the addition of neurons further contributes to the expansion.

Except that DG circuitry is very peculiar. First, granule cells receive the bulk of their afferent inputs from the very same perforant path fibers from layer II of EC that project to CA3. Why does CA3 need to receive the same information twice, once directly and once relayed by the dentate gyrus? Second, the synapses made by the PP fibers are associatively modifiable (LTP was discovered there; Lømo, 1971; Bliss and Lømo, 1973) and are quite numerous—both on DG cells, and on CA3 cells. The synapses of the mossy fibers projecting from the granule cells onto CA3 instead are very few, strong because of their close location to CA3 cell somata, large vesicular pools, and multiple release sites (Rollenhagen et al., 2007), and moreover very different from the synapses made by the cerebellar granule cells onto the Purkinje cells. So the parallel falls flat before taking off, and a different perspective is needed to solve the riddle of the presence of the dentate gyrus.

3.3.1 The conflict between learning and retrieval

In an auto-associative network like CA3, the same recurrent connections that are used to reactivate a memory must have their synaptic weights modified when acquiring or learning the representation of a new memory. If those synaptic weights already encode other memories, however, reverberations through the recurrent collaterals tend to distort the new pattern of activity and dilute the information it contains. Hasselmo and Bower (1993) suggested that the neurotransmitter acetylcholine (ACh) could resolve this conflict.

In vitro experiments in the late 1970s and later in rat CA1 and cat piriform cortex had revealed that acetylcholine suppresses synaptic transmission and increases synaptic plasticity at collateral fibers, while it has negligible effects on synaptic transmission and plasticity of afferent fibers (Hounsgaard, 1978; Valentino & Dingledine, 1981; Hasselmo & Bower, 1992). Acetylcholine could then be released during learning and render recurrent collateral synapses plastic, so that their strength can be modified, even in one shot, according to some learning rule (e.g., a Hebbian rule as in the Hopfield network); but at the same time acetylcholine would suppress synaptic transmission at those synapse, so that the network activity is determined primarily by the afferent inputs, and stays close to the new pattern to be learned. Later, at retrieval, no or little acetylcholine would be released, allowing synaptic transmission and collateral reverberations to run freely, so that the cued pattern of activity can be completed and relayed onward by the network.

Although Hasselmo and Bower (1992) originally proposed this mechanism for piriform cortex, the release of acetylcholine is widespread throughout the cortex and hippocampus, and it can be assumed that the mechanism operates in a similar fashion across many locally recurrent networks, including CA3. Yet at the same time, acetylcholine is a common neuromodulator, present across vertebrate brains and beyond; the proposed mechanism does not, therefore, help to explain the specifically mammalian riddle of what the dentate gyrus is there to do.

3.3.2 Detonator synapses and orthogonalization

Marr (1971, p. 69) had no inspired idea to propose for dentate granule cells, and he brushed them off as 'extensions of the dendrite trees of the CA pyramidal cells'. A characterization rather at odds with their sparse but large mossy fiber synapses (Andersen & Loyning, 1962; Blackstad & Kjaerheim, 1961) which would suggest a powerful non-linearity, rather than the presumably quasi-linear summation occurring in an extended dendritic tree.

McNaughton and Morris (1987), proposed to consider the opposite limit case. They assumed each mossy synapse to be powerful enough to drive a CA3 cell on its own, calling them 'detonator synapses'—a reference to synapses at the neuromuscular junction with this capacity (Eccles, 1937). Conceivably, such detonator synapses could dominate CA3 dynamics when learning a new memory if individual granule cells were to be particularly active, thereby effectively imprinting a pattern of pyramidal cell activity that, one-to-one, reflects the granule cell activity in DG. It would therefore be independent of previous memories encoded in the CA3 recurrent collaterals. During retrieval, however, granule cells with detonator synapses would have to be largely silent, thereby allowing CA3 to auto-complete a memory that has entered via the perforant path. Both PP and RC synapses would have to be modified during learning, first for the newly established pattern to be associated with the cue conveyed by the PP afferents and second, for RC reverberations to be able to complete it.

The detonator concept is extreme, and in practice it remains dubious whether single mossy fiber post-synaptic potentials are ever sufficient to drive CA3 cells (Mori, Abegg, Gähwiler, & Gerber, 2004). Brief spike trains have been shown nevertheless to have this capacity, qualifying MF synapses, perhaps, as 'conditional detonators', that need to cumulate a few of their synaptic potentials over a short time period (Henze, Wittner, & Buzsáki, 2002). Large and prolonged post-tetanic potentiation (PTP) can, however, actually push them into 'full detonator mode' (Vyleta, Borges-Merjane, & Jonas, 2016). In the end, it might not be so important, since from a neural network perspective, that looks at patterns of activity distributed over large numbers of CA3 cells, there is no substantial difference between the full and the conditional detonator mode. In the words of Edmund Rolls (1989) 'the probability that any two CA3 pyramidal cells receive synapses from a similar subset of the dentate granule cells is very low (because of the low probability of contact of any one dentate granule cell with a pyramidal cell), so that each CA3 pyramidal cell is influenced by a very different subset of the active dentate granule cells, [...] and it is therefore likely that each CA3 pyramidal cell will respond differently to the others, so that in this way pattern separation is achieved.' Conversely, if the cue that triggers retrieval were to be relayed via the mossy fibers, any weak correlation with the cue pattern would be washed away at the MF synapses due to their low number, as shown by a simple signal-to-noise analysis (Treves & Rolls, 1992). Thus, to initiate pattern completion, the cue must be transmitted to CA3 by the numerous associatively-modifiable PP synapses instead of the MF synapses. A simple quantitative analysis further indicates that, given plausible values for the relative strength of MF, PP and RC synapses, during learning only the mossy fibers, and not the perforant path would have the capacity to overcome the interference effect of the other memory patterns stored in the recurrent connections during learning (Treves &

Rolls, 1992).

Therefore, it can be said that the dentate gyrus in this model orthogonalizes the patterns to be stored in memory rather than separating them on-line. On-line, that is, during memory retrieval (to the extent that storage and retrieval phases can be rigidly separated) the dentate gyrus may as well shut off. This is in fact a critical prediction of this model: shutting off the dentate gyrus should not affect the retrieval of previously acquired memories. Its role, essentially, reduces to that of generating a pattern of activity and impressing it onto the CA3 network during learning. The pattern can be effectively random, since Hebbian plasticity associates it anyway with the afferent PP pattern and with itself, to later effect retrieval. The role of DG can then be described, for all practical purposes, as that of a random pattern generator.

3.3.3 Orthogonalization of spatial patterns

The first validation of this prediction came with an experiment in which a zinc chelator is used to block synaptic transmission with sufficient selectivity specifically at MF synapses (for about 45', before the effect of the drug is washed away). Administering such a drug daily to mice who are then released in a Morris water maze shows that they are unable, over a week, to learn the position of the hidden platform, in this standard test of hippocampal function (Lassalle et al., 2000). Control mice are also given a zinc chelator, but one that does not cross the blood-brain barrier, and so is considered to produce the same systemic effects without blocking the mossy synapses to CA3. After a week, the daily administration of the effective and the ineffective drug was switched between the two groups of rats, showing that both groups now find the platform: the first group presumably because without the MF block they can learn during the second week; and the original control group presumably because they can remember in the second week what they had stored in their CA3 memory network during the first week (Lassalle et al., 2000).

A methodologically rather different experiment, with rats running in a dry maze, led to compatible results (Lee & Kesner, 2004). In this experiment, a permanent DG lesion is shown to selectively impair acquisition of the maze lay-out; whereas a lesion of the PP fibers in the stratum lacunosum-moleculare of CA3, that is, after they have traversed through DG dendrites, is seen to disproportionately affect consolidation, or memory retention during sleep, which is presumably a retrieval dependent process.

These experiments, in rodents, are consistent with the random pattern generation hypothesis but applied to spatial patterns of activity. Furthermore, they raise the issue of what DG granule cell activity looks like, in space, and how it specifically operates to promote orthogonalization.

Experimental evidence indicates that dentate granule cells express sparse spatial representations, even compared to those in CA3 (Jung & McNaughton, 1993; Chawla et al., 2005; Leutgeb, Leutgeb, Moser, & Moser, 2007). They have place fields, but most of them are markedly silent in a given environment, and when they do have place fields, these are small and multiple (Leutgeb et al., 2007). This suggests that the configuration of active granule cells can vary substantially even between nearby spatial locations, helping to distinguish them. Correspondingly, selective DG lesion experiments in rats were shown to impair their ability to encode and retrieve spatially

close locations (Gilbert, Kesner, & Lee, 2001; Morris, Churchwell, Kesner, & Gilbert, 2012).

At the level of network models, mathematical analysis and simulations show that the argument derived for discrete patterns of activity by Treves and Rolls (1992) holds also in the case of spatial representations of the type observed by Leutgeb et al. (2007), (Cerasti & Treves, 2010).

3.4 The self-effacing role of CA1

In the human brain, CA1 comes to have over five times more pyramidal cells than CA3 (Rogers Flattery et al., 2020), which suggests that its contribution to memory processes should be at least as prominent. The salient feature of its extrinsic connectivity is that it receives both cortical afferents, from EC layer III, and the Schaffer collaterals from CA3. Already at the time of David Marr, this has generated the idea that CA1 is involved in contrasting or somehow comparing the input information it receives via the two streams (Vinogradova & Dudaeva, 1972). Note, however, that convergence of direct cortical inputs with those coming from another hippocampal subfield occurs also in CA3; to which the comparator function has sometimes also been attributed (Mizumori, Ragozzino, Cooper, & Leutgeb, 1999; Vinogradova, 2001). The function has even been attributed to the subiculum (Naber, Witter, & Lopes Da Silva, 2000), or to EC itself (Lörincz & Buzsáki, 2000).

In one version of the idea, CA1 receives predictions through the Schaffer collaterals and compares them to the actual sensory data received via the perforant path, until a mismatch occurs. This occurrence is then relayed to the ventral tegmental area (VTA), where novelty is signaled. Apart from making predictions, the upstream CA3 may act as a buffer that keeps newly presented information active for LTP induction. A short dopamine impulse would temporarily inhibit perforant path input, forcing CA1 to divert its attention away from EC and towards CA3, while dopamine receptors there may facilitate LTP and inhibit the depotentiation of Schaffer collaterals simultaneously (Lisman & Otmakhova, 2001).

A fundamental weakness of the comparator idea is that it should involve a minus sign somewhere along the way. Although the net effect of the afferent inputs can be inhibitory—upon recruiting local feedforward inhibition—the afferents themselves are excitatory, except for those recently discovered (Melzer et al., 2012); a discovery which has yet to be 'digested' from a neural computation point of view. One possibility to rescue the comparator hypothesis is to posit that it involves comparing activations elicited at different times, as suggested by the apparent dominance of SC and PP inputs to CA1 at moments of prevailing slow or fast gamma rhythmic activity, respectively (Colgin et al., 2009).

From an evolutionary perspective, as we have seen in subsubsection 3.1.2, what is most salient about CA1 is actually what is not there: its own recurrent collaterals and the mossy fiber terminals. Based on this observation, a somewhat unimaginative alternative idea is that CA1 offers a further 'clean-up' of the pattern of activity retrieved from CA3. This would actually be a variant of Marr's notion of a combined directed and free memory network in CA1, utilizing first feed-forward and then recurrent collaterals onto CA1 cells. In this variant, the combination would be inverted, as the free memory

would be implemented in CA3 by its recurrent collaterals, and the directed memory, serving as a clean-up, would be based on the (Schaffer) collaterals of the same axons, which arrive to CA1 cells. A quantitative analysis, however, indicates that the gain in the amount of information recovered at the hippocampal output would likely be minimal with such a further stage of pattern completion by the Schaffer collaterals (Treves, 1995).

3.4.1 CA1 as a public relations executive

On the basis of lesion studies, CA1 had already been suggested to link representations across time (Gilbert et al., 2001; Kesner et al., 2002). Lesions, however, are prone to confound genuine CA1 contributions with the effect of providing hippocampal feedback to the cortex, which has to go through CA1 as well. A network analysis found no advantage of the distinct CA1 circuitry in linking memory elements specifically across time (Treves, 2004) but maybe the focus on the temporal domain was too narrow.

A compelling experimental piece of evidence was the discovery that in rats the place field representations of different environments are NOT orthogonalized in CA1 as they are in CA3 (Leutgeb et al., 2004). This fits of course with the lack of mossy fibers in CA1, given the putative role of the dentate gyrus in orthogonalization. More importantly, however, it suggests that a suitably abstract version of such 'non-orthogonalization', which amounts after all to preserving in CA1 the neural representation correlations that are present in the external world, is 'linking across experiences'. For example, CA1 might use two partially correlated cell assemblies to encode two rare experiences of drinking kombucha, even though they occurred at different times, in different places and with different company. This might thus prove to be a general descriptor of the CA1 contribution that goes beyond the specifics of rodents and place cells.

Still in rodents, CA1 in fact continues to express place cells even when the inputs from CA3 have been partially or even completely inactivated (Mizumori, McNaughton, Barnes, & Fox, 1989; Brun et al., 2002), showing that the spatial information they encode is not entirely inherited from CA3. Moreover, the early observation that CA1 place fields are more frequent and less specific than in CA3 (Barnes et al., 1990) suggests perhaps that spatial information is multiplexed there with other types of information, which typical lab experiments in an open empty arena cannot readily resolve, nor certainly quantify.

The importance of mixed selectivity for cognitive function has been emphasized (Fusi, Miller, & Rigotti, 2016; Eichenbaum, 2018), but it remains difficult to study it experimentally, using tasks that cannot possibly probe more than a limited number of correlates. In the hippocampus, a variety of factors have been demonstrated that can make neurons active, both in macaques (Rolls, 1989) and in rats (Eichenbaum, Mathews, & Cohen, 1989), but it is difficult to prove that they operate on the same cells, and the degree of mixed selectivity has been properly quantified only for spatial correlates (Stefanini et al., 2020; Spalla, Treves, & Boccara, 2022). To contrast it between populations is even more arduous.

A preliminary analysis of a limited number of hippocampal neurons, in macaques, found that CA1 cells, together with cells in the subiculum and parahippocampal gyrus, express substantially higher metric content than CA3 ones (Treves, Georges-Francois,

Panzeri, Robertson, & Rolls, 1998). The metric content index can quantify the degree to which a neural representation is related or shares elements with other representations expressed by the same population. In general it is expected to be much higher for semantic memory in the cortex than for episodic memory, particularly as encoded by the putatively orthogonalized representations in CA3 (Ciaramelli, Lauro-Grotto, & Treves, 2006).

Computational theories can resort to vague formulations, such as stating that CA1 converts CA3 output into a suitable format to be used by the neocortex (Treves & Rolls, 1994; O'Reilly & McClelland, 1994; McClelland & Goddard, 1996). Still, the most effective strategy to obtain quantitative data is to focus on place cells in rodents, even though at this point one would like to extrapolate conclusions beyond rodents and beyond space.

A recent quantitative study with calcium imaging, in freely behaving mice, contrasts CA3 with CA1 cell assemblies, and finds that in CA3 they are more stable in the long-term (over weeks) and more precise (Sheintuch, Geva, Deitch, Rubin, & Ziv, 2023). The flip side of the coin is that a more flexible and imprecise CA1 could more easily incorporate into its assemblies the information from successive experiences that have elements in common. Extrapolating from the mouse experience of running in the experimental apparatus to (human) episodic memory, the flexibility may be expressed as follows: CA1 integrates the information-compressed description of the specific memory retrieved from CA3 with a multitude of additional elements originating in other episodes and in semantic memory, conveyed from the neocortex by the EC inputs. This integration may facilitate both the consolidation of long-term memory storage in the neocortex and the immediate use of the retrieved memory in relation to the current context (Treves & Rolls, 1994). According to this idea, CA1 would be a sort of mediator between the straightforward but rigid memory operation in CA3, and the complexity lurking outside.

4 Part III: The view beyond

This discussion of how the mammalian hippocampus may serve memory is perforce incomplete without considering what the beneficiary of the service can do with it—where the primary beneficiary is not the behavior of the individual, but rather the cortex of that same individual.

Discussing memory in the cortex, however, is not feasible here and outside the scope of this chapter. Still, a few critical questions are worth raising:

- Can the hippocampus act as a teacher to the cortex? Learning and memory in the cortex have often been discussed in a hierarchically arranged connectionist framework (McClelland, McNaughton, & O'Reilly, 1995), in which synaptic plasticity is supervised, guided by a teacher which provides a 'desired' output. The difference between desired and actual output is then backpropagated through lower levels of the hierarchy, to tell them how they should change their ways. One problem of such a scenario, apart from the biological infeasibility of the backpropagation concept, is the need for a minus sign, just like in the notion of CA1 as a comparator. Is it possible, however, to conceive of the hippocampus

as teaching by example, not by correcting errors—requiring the minus sign—but just by reinforcing spontaneously occurring contingencies? Then it could make use of the associative synaptic plasticity in the superficial layers of the neocortex, where backprojections originating in the hippocampus and in other associational cortices terminate (Rolls, 2000).

- Can the hippocampus help the cortex with the compositionality of episodic memories? Whether episodic memories are ever stored in their original format in the neocortex remains controversial; one problem with it is their compositional nature. If most episodic memories are plausibly conceived as including a relatively small number of salient elements, each having a-priori stable representations in the cortex (see Figure 8), the patterns of cortical activity for these memories cannot all be approximately uncorrelated, because the law of large numbers does not apply: a pair of such memories will be correlated in 0 or in 1, or 2, or 3, ... of such elements rather than in $\approx \sqrt{N}$ out of N cortical cells. Auto-associative memory capacity is then dramatically reduced (Ryom, Stendardi, Ciaramelli, & Treves, 2023) and to still retrieve composite patterns the hippocampus might be required to provide a reactivating signal, as in the 'index' theory (Teyler & DiScenna, 1986), making good use of the compositional encoding that has been observed, e.g. in CA2 with social memories (Boyle, Posani, Irfan, Siegelbaum, & Fusi, 2024). This could help explain why even remote episodic memories are lost with hippocampal damage (Sanders & Warrington, 1971).

- Can the hippocampus provide spontaneous cortical dynamics with sensory content? When the mind is disconnected from the external world it can still wander through sensory experiences, e.g. through scenes (Maguire & Mullally, 2013), which have to come from memory. Hippocampal amnesia, on the other hand, impairs the imaginative faculty (Hassabis, Kumaran, Vann, & Maguire, 2007). Still, models of cortical dynamics during mind-wandering are only just beginning to be formulated (Ciaramelli & Treves, 2019; Ryom, Basu, Stendardi, Ciaramelli, & Treves, 2024).

4.1 Open issues with the hippocampus

In addition to these and many more questions about hippocampo-cortical interactions, there are however several issues about mechanisms at work in the hippocampus itself, which might well be resolved or at least better understood in the near future. We list just three, which are to some extent related.

4.1.1 Representational turnover and cell reuse

Calcium imaging enabled the discovery that hippocampal representations of spatial environments are not stable across days and weeks: substantial turnover was reported of cells that 'enter' and 'exit' the relevant cell assembly(Ziv et al., 2013), like actors transiently playing their roles in a long-running Broadway musical. Another even more striking facet of this looser correspondence between activity patterns and what they represent is the finding that multiple representations of the same environment can

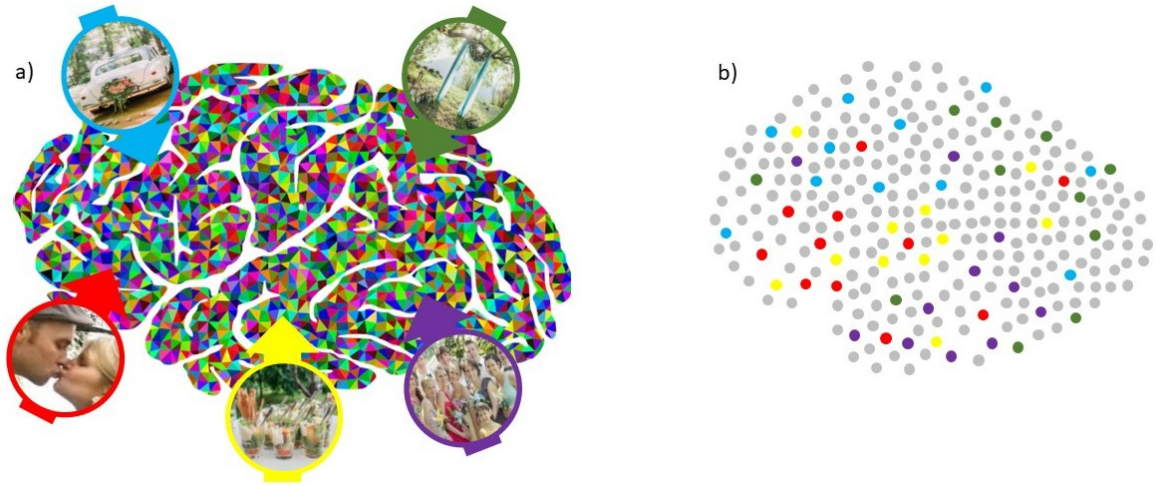

Figure 8: Left – The episodic memory of a friend's wedding is a new composition of events that have already been experienced in one way or another, such as the meal, the kiss, the gathering of friends, the canopy, and the flower-decorated wedding car. Right – In a model cortical network, each of these five elements is a sparsely distributed representation across several patches of cortex; each circle represents a small patch of neocortex. Image adapted from Ryom, Stendardi, Ciaramelli, and Treves (2023).

coexist, and mice can 'choose' which one to reactivate (Sheintuch et al., 2020), like several shows of the same musical being put up at the same time by different companies.

The quantitative relevance of this amazing phenomenon in ecological conditions is still being investigated, but it seems that a new generation of computational hippocampal models will have to take it into account, one way or another.

4.1.2 The function of neurogenesis

As mentioned above, adult neurogenesis, which is generally suppressed in the mammalian brain, has been reported in the dentate gyrus of several species. In laboratory rats, it has been estimated that about 6% new granule cells are generated each month, based on K-67 gene expression (Cameron & Mckay, 2001). In humans, adult neurogenesis has been estimated using atmospheric 14-C levels caused by nuclear bomb tests (Spalding et al., 2013), and found to be about 700 new neurons per day, which would amount to a rate of order 1% per year. Other studies express skepticism about these counts; or they report no neurogenesis in humans, or in species such as bats, dolphins or whales; or they point at a rate decrease with age, and at the many new cells which die after a few weeks (Amrein & Lipp, 2009; Augusto-Oliveira, Arrifano, Malva, & Crespo-Lopez, 2019).

Proposals about the functional significance of adult neurogenesis have varied, ranging from a role in forgetting old memories (Frankland, Köhler, & Josselyn, 2013) to temporally tagging new ones (Aimone, Wiles, & Gage, 2006). What seems to be clear, at least in rodents, is that new neurons are not smoothly and quietly integrated into the existing circuit: they exhibit enhanced activity during their 'puberty' period (Alme et al., 2010; Kropff, Yang, & Schinder, 2015), maybe to show their usefulness; and then

go into early retirement—or perhaps those who have not proven themselves are even eliminated.

Moreover, proliferation rates seem to correlate with the novelty rather than the complexity of the surrounding environment (Kempermann & Gage, 1999), to increase with physical exercise also in humans (Pereira et al., 2007; Van Praag, 2008); and are even 3 times higher in wild-living species, such as voles and wood mice (Amrein, Slomianka, Poletaeva, Bologova, & Lipp, 2004).

However the adult neurogenesis issue will be clarified in the future, these observation indicate that focusing on animals raised in the lab may not be sufficient. The need for controlled laboratory measures has to be balanced with avoiding inferences and conclusions based solely on individuals, strains and species which, in a sense, are pathological: they suffer from having been locked up in the lab. Neuroscience has to move into the open. Again, to address ecological conditions.

4.1.3 Disorder

A significant step towards ecological conditions has been taken in the lab of Nachum Ulanovsky, by recording from hippocampal cells in bats that fly in a tunnel, initially 200 m long (Eliav et al., 2021) and more recently even longer (but still well below their normal navigational range). One key observation, in such quasi-ecological setting, has been the confirmation that CA1 cells exhibit multiple fields, as already seen in rats (Park, Dvorak, & Fenton, 2011). Even more striking, the multiple fields are seen to vary considerably in width and peak rate, with the wide variability surprisingly well-described by simple probability distributions—exponential for the number of fields per cell and log-normal for their widths and peak rates. Such distributions, or something very close to them, have now been argued to arise from a purely random process (Mainali, Azeredo Da Silveira, & Burak, 2025). A simple recurrent model shows that, if similar statistics apply to CA3, it results in a novel capacity limit, due to the fact that the combined peak rate and width variability in sufficiently large environments effectively prevents the establishment of a continuous attractor to represent the tunnel (Schönsberg, Monasson, & Treves, 2024). Disorder—if the observed variability stems from randomness it can be labeled as disorder—may be a more serious constraint on autoassociative networks than had been imagined so far. Whether this entirely unexpected effect is relevant for the real CA3 network, in bats and in other species, remains to be seen, but the general implication is that theories and quantitative results obtained with models informed by experimental data but also based on idealized notions (such as, in this case, one field per cell, of standard width and peak rate) may have to be revised. Sometimes leading to completely new insights and conclusions.

4.2 Incidental and interim conclusions

Once we are sensitized to the potential relevance of ecological conditions, of variability and disorder, we may find a host of intriguing observations that had been overlooked in previously acquired data, as well as devise ways to better understand their effects with new experiments.

One small example has to do with incidental learning, which occurs spontaneously

when animals or humans acquire information that is of no use to them—at least of no foreseeable use. In a standard task probing incidental learning and memory, mice are exposed to several different objects in a box and are then taken away, to be reintroduced to it after a delay, during which one of the objects has been changed. If they explore the new object more than the others—which they tend to do spontaneously—it is a sign that they remember which objects they had been exposed to. Typically they do remember, at least up to 6 different objects. Two interesting incidental observations with this incidental learning task concern c-Fos gene expression (an index of neuronal activation) in the dorsal hippocampus: first, male mice show expression levels 5 times higher than females; and second, exposing mice to 6 objects leads to dramatically higher levels of c-Fos expression than are seen with 3 objects, in males even tens of times higher (Torromino et al., 2022). Although the exact quantitative relation of c-Fos expression to neuronal activity remains to be clarified, such a dramatic non-linearity is hardly compatible with the tacit assumption made in many mathematical models, that the average activity in hippocampal networks (and in associative memories in the brain generally) is tightly regulated, while it is the distribution of the activity across neurons that varies from one pattern to another.

Findings like these have the potential to pull the rug from under the feet of network models that, over the last few decades, have largely supported and offered further resolution and detail to the theory of simple memory, proposed by the young David Marr (1971). Overall, models of the hippocampus, particularly when integrated with experimental data from rodents, have been at the forefront of efforts to understand computationally how the brain works. Marr had written that while his theory of such a complex structure as the cerebral cortex had to remain at an abstract level, that of the archicortex could be spelled out in full, in terms of neurons and synapses; the construction of such a theory was, he wrote, 'little more than a technical exercise'. He did not convincingly carry out the exercise, but he inspired many later models. Some of these models, or maybe their average, have come to be called by some researchers the 'standard framework'. Evoking a standard framework betrays a self-righteous sensation that the field has committed to it: the work has been done, and everything is now clear.

That might be an illusion. We may be at a stage in our understanding of the hippocampus and its role in the brain somewhat analogous to the stage the physics of space and time was in, nearly a century and a half ago. The Michelson and Morley experiment demonstrated that some idealized assumptions were incorrect, and opened the door to the entirely new physics of the 20th century. It could well be that in order to make progress, with the hippocampus and with the brain in general, we must, to some extent, get out of the lab and into the space and time of real life.

Acknowledgments

We are indebted to many colleagues for extensive discussions and suggestions, and particularly to Cosme Salas, Edmund Rolls, Menno Witter and Lutz Slomianka. We are aware of the personal perspective adopted in this review and that space constraints have not allowed us to even mention a large number of relevant studies.

5 References

Abeles, M. (1991). *Corticonics: Neural circuits of the cerebral cortex*. Cambridge: Cambridge University Press.

Aimone, J. B., Wiles, J., & Gage, F. H. (2006). Potential role for adult neurogenesis in the encoding of time in new memories. *Nature Neuroscience*, *9*(6), 723–727. doi:10.1038/nn1707

Alme, C., Buzzetti, R., Marrone, D., Leutgeb, J., Chawla, M., Schaner, M., ... Barnes, C. (2010). Hippocampal granule cells opt for early retirement. *Hippocampus*, *20*(10), 1109–1123. doi:10.1002/hipo.20810

Alme, C. B., Miao, C., Jezek, K., Treves, A., Moser, E. I., & Moser, M.-B. (2014). Place cells in the hippocampus: Eleven maps for eleven rooms. *Proceedings of the National Academy of Sciences*, *111*(52), 18428–18435. doi:10.1073/pnas.1421056111

Altman, J., & Das, G. D. (1965). Autoradiographic and histological evidence of postnatal hippocampal neurogenesis in rats. *Journal of Comparative Neurology*, *124*(3), 319–335. doi:10.1002/cne.901240303

Amaral, D. G., Ishizuka, N., & Claiborne, B. (1990). Neurons, numbers and the hippocampal network. *Progress in Brain Research*, *83*, 1–11.

Amaral, D. G., Scharfman, H. E., & Lavenex, P. (2008). The dentate gyrus: Fundamental neuroanatomical organization (dentate gyrus for dummies). *Progress in brain research*, *163*, 3–22. doi:10.1016/S0079-6123(07)63001-5

Amari, S.-I. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, *27*(2), 77–87. doi:10.1007/BF00337259

Amit, D. J., Gutfreund, H., & Sompolinsky, H. (1987). Statistical mechanics of neural networks near saturation. *Annals of Physics*, *173*(1), 30–67. doi:10.1016/0003-4916(87)90092-3

Amrein, I., & Lipp, H.-P. (2009). Adult hippocampal neurogenesis of mammals: Evolution and life history. *Biology Letters*, *5*(1), 141–144. doi:10.1098/rsbl.2008.0511

Amrein, I., Slomianka, L., & Lipp, H.-P. (2004). Granule cell number, cell death and cell proliferation in the dentate gyrus of wild-living rodents. *European Journal of Neuroscience*, *20*(12), 3342–3350. doi:10.1111/j.1460-9568.2004.03795.x

Amrein, I., Slomianka, L., Poletaeva, I. I., Bologova, N. V., & Lipp, H.-P. (2004). Marked species and age-dependent differences in cell proliferation and neurogenesis in the hippocampus of wild-living rodents. *Hippocampus*, *14*(8), 1000–1010. doi:10.1002/hipo.20018

Andersen, P., & Loyning, Y. (1962). Interaction of various afferents on CA1 neurons and dentate granule cells. *Colloq Int CNRS*, *107*, 23–45.

Andersen, P., Holmqvist, B., & Voorhoeve, P. E. (1966). Excitatory Synapses on Hippocampal Apical Dendrites Activated by Entorhinal Stimulation. *Acta Physiologica Scandinavica*, *66*(4), 461–472. doi:10.1111/j.1748-1716.1966.tb03224.x

Anderson, M. I., & Jeffery, K. J. (2003). Heterogeneous Modulation of Place Cell Firing by Changes in Context. *The Journal of Neuroscience*, *23*(26), 8827–8835. doi:10.1523/JNEUROSCI.23-26-08827.2003

Aniksztejn, L., Charton, G., & Ben-Ari, Y. (1987). Selective release of endogenous zinc from the hippocampal mossy fibers in situ. *Brain Research*, *404*(1-2), 58–64. doi:10.1016/0006-8993(87)91355-2

Aniksztejn, L., Demarque, M., Morozov, Y., Ben-Ari, Y., & Represa, A. (2001). Recurrent CA1 collateral axons in developing rat hippocampus. *Brain Research*, *913*(2), 195–200. doi:10.1016/S0006-8993(01)02817-7

Antal, A., Kéri, S., Kovács, G., Janka, Z., & Benedek, G. (2000). Early and late components of visual categorization: An event-related potential study. *Cognitive Brain Research*, *9*(1), 117–119. doi:10.1016/S0926-6410(99)00053-1

Atoji, Y., & Wild, J. M. (2004). Fiber connections of the hippocampal formation and septum and subdivisions of the hippocampal formation in the pigeon as revealed by tract tracing and kainic acid lesions. *Journal of Comparative Neurology*, *475*(3), 426–461. doi:10.1002/cne.20186

Augusto-Oliveira, M., Arrifano, G., Malva, J., & Crespo-Lopez, M. (2019). Adult Hippocampal Neurogenesis in Different Taxonomic Groups: Possible Functional Similarities and Striking Controversies. *Cells*, *8*(2), 125. doi:10.3390/cells8020125

Azizi, A. H., Wiskott, L., & Cheng, S. (2013). A computational model for preplay in the hippocampus. *Frontiers in Computational Neuroscience*, *7*. doi:10.3389/fncom.2013.00161

Baraduc, P., Duhamel, J.-R., & Wirth, S. (2019). Schema cells in the macaque hippocampus. *Science*, *363*(6427), 635–639. doi:10.1126/science.aav5404

Barnes, C. A., McNaughton, B. L., Mizumori, S. J. Y., Leonard, B. W., & Lin, L.-H. (1990). Comparison of spatial and temporal characteristics of neuronal activity in sequential stages of hippocampal processing. *Progress in brain research*, *83*, 287–300. doi:10.1016/S0079-6123(08)61257-1

Battaglia, F. P., & Treves, A. (1998a). Attractor neural networks storing multiple space representations: A model for hippocampal place fields. *Physical Review E*, *58*(6), 7738–7753. doi:10.1103/PhysRevE.58.7738

Battaglia, F. P., & Treves, A. (1998b). Stable and Rapid Recurrent Processing in Realistic Autoassociative Memories. *Neural Computation*, *10*(2), 431–450. doi:10.1162/089976698300017827

Bayley, P. J., Hopkins, R. O., & Squire, L. R. (2006). The Fate of Old Memories after Medial Temporal Lobe Damage. *Journal of Neuroscience*, *26*(51), 13311–13317. doi:10.1523/JNEUROSCI.4262-06.2006

Bentivoglio, M., Cotrufo, T., Ferrari, S., Tesoriero, C., Mariotto, S., Bertini, G., … Mazzarello, P. (2019). The Original Histological Slides of Camillo Golgi and His Discoveries on Neuronal Structure. *Frontiers in Neuroanatomy*, *13*, 3. doi:10.3389/fnana.2019.00003

Blackstad, T. W. (1958). On the termination of some afferents to the hippocampus and fascia dentata. *Cells Tissues Organs*, *35*(3), 202–214. doi:10.1159/000141409

Blackstad, T. W., & Kjaerheim, Å. (1961). Special axo-dendritic synapses in the hippocampal cortex: Electron and light microscopic studies on the layer of mossy fibers. *Journal of Comparative Neurology*, *117*(2), 133–159. doi:10.1002/cne.901170202

Bliss, T. V. P., & Lømo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *The Journal of Physiology*, *232*(2), 331–356. doi:10.1113/jphysiol.1973.sp010273

Bliss, T., & Lømo, T. (2024). Per Andersen. 12 January 1930—17 February 2020. *Biographical Memoirs of Fellows of the Royal Society*, *76*, 27–47. doi:10.1098/rsbm.2023.0041

Blum, K. I., & Abbott, L. F. (1996). A Model of Spatial Map Formation in the Hippocampus of the Rat. *Neural Computation*, *8*(1), 85–93. doi:10.1162/neco.1996.8.1.85

Borzello, M., Ramirez, S., Treves, A., Lee, I., Scharfman, H., Stark, C., … Rangel, L. M. (2023). Assessments of dentate gyrus function: Discoveries and debates. *Nature Reviews Neuroscience*, *24*(8), 502–517. doi:10.1038/s41583-023-00710-z

Bostock, E., Muller, R. U., & Kubie, J. L. (1991). Experience-dependent modifications of hippocampal place cell firing. *Hippocampus*, *1*(2), 193–205. doi:10.1002/hipo.450010207

Boyle, L. M., Posani, L., Irfan, S., Siegelbaum, S. A., & Fusi, S. (2024). Tuned geometries of hippocampal representations meet the computational demands of social memory. *Neuron*, *112*(8), 1358–1371.e9. doi:10.1016/j.neuron.2024.01.021

Brindley, G. S. (1967). The Classification of Modifiable Synapses and Their Use In Models For Conditioning. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, *168*(1013), 361–376. doi:10.1098/rspb.1967.0070

Brodal, A. (1947). The Hippocampus and the Sense of Smell: A Review. *Brain*, *70*(2), 179–222. doi:10.1093/brain/70.2.179

Brun, V. H., Otnæss, M. K., Molden, S., Steffenach, H.-A., Witter, M. P., Moser, M.-B., & Moser, E. I. (2002). Place Cells and Place Recognition Maintained by Direct Entorhinal-Hippocampal Circuitry. *Science*, *296*(5576), 2243–2246. doi:10.1126/science.1071089

Burgess, N., Maguire, E. A., & O'Keefe, J. (2002). The Human Hippocampus and Spatial and Episodic Memory. *Neuron*, *35*(4), 625–641. doi:10.1016/s0896-6273(02)00830-9

Burgess, N., & O'Keefe, J. (2011). Models of place and grid cell firing and theta rhythmicity. *Current Opinion in Neurobiology*, *21*(5), 734–744. doi:10.1016/j.conb.2011.07.002

Buzsáki, G. (2006). *Rhythms of the Brain*. New York: Oxford University Press.

Cameron, H. A., & Mckay, R. D. (2001). Adult neurogenesis produces a large pool of new granule cells in the dentate gyrus. *Journal of Comparative Neurology*, *435*(4), 406–417. doi:10.1002/cne.1040

Cerasti, E., & Treves, A. (2010). How Informative Are Spatial CA3 Representations Established by the Dentate Gyrus? *PLoS Computational Biology*, *6*(4), e1000759. doi:10.1371/journal.pcbi.1000759

Chawla, M., Guzowski, J., Ramirez-Amaya, V., Lipa, P., Hoffman, K., Marriott, L., … Barnes, C. (2005). Sparse, environmentally selective expression ofArc RNA in the upper blade of the rodent fascia dentata by brief spatial experience. *Hippocampus*, *15*(5), 579–586. doi:10.1002/hipo.20091

Cheng, S., & Werning, M. (2016). What is episodic memory if it is a natural kind? *Synthese*, *193*(5), 1345–1385. doi:10.1007/s11229-014-0628-6

Chomsky, N. (1981). *Lectures on government and binding: The Pisa lectures*. Dordrecht: Reidel.

Ciaramelli, E., Lauro-Grotto, R., & Treves, A. (2006). Dissociating episodic from semantic access mode by mutual information measures: Evidence from aging and Alzheimer's disease. *Journal of Physiology-Paris*, *100*(1-3), 142–153. doi:10.1016/j.jphysparis.2006.09.008

Ciaramelli, E., & Treves, A. (2019). A Mind Free to Wander: Neural and Computational Constraints on Spontaneous Thought. *Frontiers in Psychology*, *10*, 39. doi:10.3389/fpsyg.2019.00039

Clayton, N. S., & Dickinson, A. (1998). Episodic-like memory during cache recovery by scrub jays. *Nature*, *395*(6699), 272–274. doi:10.1038/26216

Cnotka, J., Möhle, M., & Rehkämper, G. (2008). Navigational Experience Affects Hippocampus Size in Homing Pigeons. *Brain, Behavior and Evolution*, *72*(3), 233–238. doi:10.1159/000165102

Colgin, L. L., Denninger, T., Fyhn, M., Hafting, T., Bonnevie, T., Jensen, O., … Moser, E. I. (2009). Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature*, *462*(7271), 353–357. doi:10.1038/nature08573

Compte, A., Brunel, N., Goldman-Rakic, P. S., & Wang, X.-J. (2000). Synaptic Mechanisms and Network Dynamics Underlying Spatial Working Memory in a Cortical Network Model. *Cerebral Cortex*, *10*(9), 910–923. doi:10.1093/cercor/10.9.910

Coppola, V. J., & Bingman, V. P. (2020). C-Fos revealed lower hippocampal participation in older homing pigeons when challenged with a spatial memory task. *Neurobiology of Aging*, *87*, 98–107. doi:10.1016/j.neurobiolaging.2019.11.019

Corkin, S. (1984). Lasting Consequences of Bilateral Medial Temporal Lobectomy: Clinical Course and Experimental Findings in H.M. *Seminars in Neurology*, *4*(02), 249–259. doi:10.1055/s-2008-1041556

Corkin, S. (2002). What's new with the amnesic patient H.M.? *Nature Reviews Neuroscience*, *3*(2), 153–160. doi:10.1038/nrn726

Davies, J. R., & Clayton, N. S. (2024). Is episodic-like memory *like* episodic memory? *Philosophical Transactions of the Royal Society B: Biological Sciences*, *379*(1913), 20230397. doi:10.1098/rstb.2023.0397

Day, L. B., Crews, D., & Wilczynski, W. (1999). Spatial and reversal learning in congeneric lizards with different foraging strategies. *Animal Behaviour*, *57*(2), 393–407. doi:10.1006/anbe.1998.1007

Day, L. B., Crews, D., & Wilczynski, W. (2001). Effects of medial and dorsal cortex lesions on spatial memory in lizards. *Behavioural Brain Research*, *118*(1), 27–42. doi:10.1016/S0166-4328(00)00308-9

De Nó, L. (1934). Studies on the structure of the cerebral cortex. II. Continuation of the study of the ammonic system. *Journal für Psychologie und Neurologie*, *46*, 113–177.

Debanne, D., Gähwiler, B. H., & Thompson, S. M. (1998). Long-term synaptic plasticity between pairs of individual CA3 pyramidal cells in rat hippocampal slice cultures. *The Journal of Physiology*, *507*(1), 237–247. doi:10.1111/j.1469-7793.1998.237bu.x

Docampo-Seara, A., Lagadec, R., Mazan, S., Rodríguez, M. A., Quintana-Urzainqui, I., & Candal, E. (2018). Study of pallial neurogenesis in shark embryos and the evolutionary origin of the subventricular zone. *Brain Structure and Function*, *223*(8), 3593–3612. doi:10.1007/s00429-018-1705-2

Dragoi, G., & Tonegawa, S. (2011). Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature*, *469*(7330), 397–401. doi:10.1038/nature09633

Eccles, J. C. (1937). Synaptic and Neuro-Muscular Transmission. *Physiological Reviews*, *17*(4), 538–555. doi:10.1007/BF02320730

Eichenbaum, H. (1997). Declarative Memory: Insights from Cognitive Neurobiology. *Annual Review of Psychology*, *48*(1), 547–572. doi:10.1146/annurev.psych.48.1.547

Eichenbaum, H. (2018). Barlow versus Hebb: When is it time to abandon the notion of feature detectors and adopt the cell assembly as the unit of cognition? *Neuroscience Letters*, *680*, 88–93. doi:10.1016/j.neulet.2017.04.006

Eichenbaum, H., Mathews, P., & Cohen, N. J. (1989). Further Studies of Hippocampal Representation During Odor Discrimination Learning. *Behavioral neuroscience*, *103*(6), 1207–1216. doi:10.1037/0735-7044.103.6.1207

Eichenbaum, H., Otto, T., & Cohen, N. J. (1994). Two functional components of the hippocampal memory system. *Behavioral and Brain Sciences*, *17*(3), 449–472. doi:10.1017/S0140525X00035391

Ekstrom, A. D., Caplan, J. B., Ho, E., Shattuck, K., Fried, I., & Kahana, M. J. (2005). Human hippocampal theta activity during virtual navigation. *Hippocampus*, *15*(7), 881–889. doi:10.1002/hipo.20109

Eliav, T., Maimon, S. R., Aljadeff, J., Tsodyks, M., Ginosar, G., Las, L., & Ulanovsky, N. (2021). Multiscale representation of very large environments in the hippocampus of flying bats. *Science*, *372*(6545), eabg4020. doi:10.1126/science.abg4020

Eriksson, P. S., Perfilieva, E., Björk-Eriksson, T., Alborn, A.-M., Nordborg, C., Peterson, D. A., & Gage, F. H. (1998). Neurogenesis in the adult human hippocampus. *Nature Medicine*, *4*(11), 1313–1317. doi:10.1038/3305

Font, E. (2019). Rapid learning of a spatial memory task in a lacertid lizard (Podarcis liolepis). *Behavioural Processes*, *169*, 103963. doi:10.1016/j.beproc.2019.103963

Frankland, P. W., Köhler, S., & Josselyn, S. A. (2013). Hippocampal neurogenesis and forgetting. *Trends in Neurosciences*, *36*(9), 497–503. doi:10.1016/j.tins.2013.05.002

Franzius, M., Sprekeler, H., & Wiskott, L. (2007). Slowness and Sparseness Lead to Place, Head-Direction, and Spatial-View Cells. *PLoS Computational Biology*, *3*(8), e166. doi:10.1371/journal.pcbi.0030166

Fuchs, E. C., Zivkovic, A. R., Cunningham, M. O., Middleton, S., LeBeau, F. E., Bannerman, D. M., … Monyer, H. (2007). Recruitment of Parvalbumin-Positive Interneurons Determines Hippocampal Function and Associated Behavior. *Neuron*, *53*(4), 591–604. doi:10.1016/j.neuron.2007.01.031

Fusi, S., Miller, E. K., & Rigotti, M. (2016). Why neurons mix: High dimensionality for higher cognition. *Current Opinion in Neurobiology*, *37*, 66–74. doi:10.1016/j.conb.2016.01.010

Gabrieli, J. D., Cohen, N. J., & Corkin, S. (1988). The impaired learning of semantic knowledge following bilateral medial temporal-lobe resection. *Brain and Cognition*, *7*(2), 157–177. doi:10.1016/0278-2626(88)90027-9

Gaffan, D. (1974). Recognition impaired and association intact in the memory of monkeys after transection of the fornix. *Journal of Comparative and Physiological Psychology*, *86*(6), 1100–1109. doi:10.1037/h0037649

Gaffan, D. (1992). Amnesia for Complex Naturalistic Scenes and for Objects Following Fornix Transection in the Rhesus Monkey. European Journal of Neuroscience. *European Journal of Neuroscience*, *4*, 381–388. doi:10.1111/j.1460-9568.1992.tb00886.x

Gaffan, D. (1993). Normal forgetting, impaired acquisition in memory for complex naturalistic scenes by fornix-transected monkeys. *Neuropsychologia*, *31*(4), 403–406. doi:10.1016/0028-3932(93)90163-T

Gaffan, D. (2002). Against memory systems. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, *357*(1424), 1111–1121. doi:10.1098/rstb.2002.1110

Gagliardo, A., Pollonara, E., Casini, G., Rossino, M. G., Wikelski, M., & Bingman, V. P. (2020). Importance of the hippocampus for the learning of route fidelity in homing pigeons. *Biology Letters*, *16*(7), 20200095. doi:10.1098/rsbl.2020.0095

Gardner-Medwin, A. R. (1976). The recall of events through the learning of associations between their parts. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, *194*(1116), 375–402. doi:10.1098/rspb.1976.0084

Garvert, M. M., Dolan, R. J., & Behrens, T. E. (2017). A map of abstract relational knowledge in the human hippocampal–entorhinal cortex. *eLife*, *6*, e17086. doi:10.7554/eLife.17086

Gilbert, P. E., Kesner, R. P., & Lee, I. (2001). Dissociating hippocampal subregions: A double dissociation between dentate gyrus and CA1. *Hippocampus*, *11*(6), 626–636. doi:10.1002/hipo.1077

Ginsberg, S. D., & Che, S. (2005). Expression profile analysis within the human hippocampus: Comparison of CA1 and CA3 pyramidal neurons. *Journal of Comparative Neurology*, *487*(1), 107–118. doi:10.1002/cne.20535

Gloor, P. (1997). *The temporal lobe and limbic system*. New York: Oxford university press.

Golgi, C. (1885). *Sulla Fina Anatomia Degli Organi Centrali del Sistema Nervoso*. Reggio Emilia: S. Calderini e Figlio.

Gonzales, R. B., DeLeon Galvan, C. J., Rangel, Y. M., & Claiborne, B. J. (2001). Distribution of thorny excrescences on CA3 pyramidal neurons in the rat hippocampus. *The Journal of Comparative Neurology*, *430*(3), 357–368. doi:10.1002/1096-9861(20010212)430:3<357::AID-CNE1036>3.0.CO;2-K

Graham, K. S., & Hodges, J. R. (1997). Differentiating the roles of the hippocampus complex and the neocortex in long-term memory storage: Evidence from the study of semantic dementia and Alzheimer's disease. *Neuropsychology*, *11*(1), 77–89. doi:10.1037//0894-4105.11.1.77

Granger, R., Whitson, J., Larson, J., & Lynch, G. (1994). Non-Hebbian properties of long-term potentiation enable high-capacity encoding of temporal sequences. *Proceedings of the National Academy of Sciences*, *91*(21), 10104–10108. doi:10.1073/pnas.91.21.10104

Guardamagna, M., Stella, F., & Battaglia, F. P. (2023). Heterogeneity of network and coding states in mouse CA1 place cells. *Cell Reports*, *42*(2), 112022. doi:10.1016/j.celrep.2023.112022

Gulyás, A. I., Megías, M., Emri, Z., & Freund, T. F. (1999). Total Number and Ratio of Excitatory and Inhibitory Synapses Converging onto Single Interneurons

of Different Types in the CA1 Area of the Rat Hippocampus. *The Journal of Neuroscience*, *19*(22), 10082–10097. doi:10.1523/JNEUROSCI.19-22-10082.1999

Guzowski, J., Knierim, J. J., & Moser, E. (2004). Ensemble dynamics of hippocampal regions CA3 and CA1. *Neuron*, *44*, 581–584. doi:10.1016/j.neuron.2004.11.003

Hamel, E. G. (1966). A study of the hippocampal formation in the opossum, Didelphis virginiana. In R. Hassler & H. Stephan (Eds.), *Evolution of the Forebrain: Phylogenesis and Ontogenesis of the Forebrain* (pp. 81–91). doi:10.1007/978-1-4899-6527-1_9

Hassabis, D., Kumaran, D., Vann, S. D., & Maguire, E. A. (2007). Patients with hippocampal amnesia cannot imagine new experiences. *Proceedings of the National Academy of Sciences*, *104*(5), 1726–1731. doi:10.1073/pnas.0610561104

Hasselmo, M. E., & Bower, J. M. (1992). Cholinergic suppression specific to intrinsic not afferent fiber synapses in rat piriform (olfactory) cortex. *Journal of Neurophysiology*, *67*(5), 1222–1229. doi:10.1152/jn.1992.67.5.1222

Hasselmo, M. E., & Bower, J. M. (1993). Acetylcholine and memory. *Trends in Neurosciences*, *16*(6), 218–222. doi:10.1016/0166-2236(93)90159-J

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

Henze, D. A., Wittner, L., & Buzsáki, G. (2002). Single granule cells reliably discharge targets in the hippocampal CA3 network in vivo. *Nature Neuroscience*, *5*(8), 790–795. doi:10.1038/nn887

Herold, C., Schlömer, P., Mafoppa-Fomat, I., Mehlhorn, J., Amunts, K., & Axer, M. (2019). The hippocampus of birds in a view of evolutionary connectomics. *Cortex*, *118*, 165–187. doi:10.1016/j.cortex.2018.09.025

Herrick, C. J. (1933). The Functions of the Olfactory Parts of the Cerebral Cortex. *Proceedings of the National Academy of Sciences*, *19*(1), 7–14. doi:10.1073/pnas.19.1.7

Hitti, F. L., & Siegelbaum, S. A. (2014). The hippocampal CA2 region is essential for social memory. *Nature*, *508*(7494), 88–92. doi:10.1038/nature13028

Hopfield, J. J. (1984). Neurons with graded response have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences*, *81*(10), 3088–3092. doi:10.1073/pnas.81.10.3088

Hopfield, J. J. (1982). Neural networks and physical systems with emergent computational abilities. *Proceedings of the national academy of sciences*, *79*(8), 2554–2558. doi:10.1073/pnas.79.8.2554

Hopfield, J. J. (2010). Neurodynamics of mental exploration. *Proceedings of the National Academy of Sciences*, *107*(4), 1648–1653. doi:10.1073/pnas.0913991107

Hounsgaard, J. (1978). Presynaptic inhibitory action of acetylcholine in area CA1 of the hippocampus. *Experimental Neurology*, *62*(3), 787–797. doi:10.1016/0014-4886(78)90284-4

Itskov, V., Curto, C., Pastalkova, E., & Buzsáki, G. (2011). Cell Assembly Sequences Arising from Spike Threshold Adaptation Keep Track of Time in the Hippocampus. *The Journal of Neuroscience*, *31*(8), 2828–2834. doi:10.1523/JNEUROSCI.3773-10.2011

Jezek, K., Henriksen, E. J., Treves, A., Moser, E. I., & Moser, M.-B. (2011). Theta-paced flickering between place-cell maps in the hippocampus. *Nature*, *478*(7368), 246–249. doi:10.1038/nature10439

Jones, M. W., & McHugh, T. J. (2011). Updating hippocampal representations: CA2 joins the circuit. *Trends in Neurosciences*, *34*(10), 526–535. doi:10.1016/j.tins.2011.07.007

Jung, M. W., & McNaughton, B. L. (1993). Spatial selectivity of unit activity in the hippocampal granular layer. *Hippocampus*, *3*(2), 165–182. doi:10.1002/hipo.450030209

Jutras, M. J., Fries, P., & Buffalo, E. A. (2013). Oscillatory activity in the monkey hippocampus during visual exploration and memory formation. *Proceedings of the National Academy of Sciences*, *110*(32), 13144–13149. doi:10.1073/pnas.1302351110

Kahn, M. C., Hough II, G. E., Ten Eyck, G. R., & Bingman, V. P. (2003). Internal connectivity of the homing pigeon (Columba livia) hippocampal formation: An anterograde and retrograde tracer study. *Journal of Comparative Neurology*, *459*(2), 127–141. doi:10.1002/cne.10601

Kempermann, G., & Gage, F. H. (1999). Experience-dependent regulation of adult hippocampal neurogenesis: Effects of long-term stimulation and stimulus withdrawal. *Hippocampus*, *9*(3), 321–332. doi:10.1002/(SICI)1098-1063(1999)9:3<321::AID-HIPO11>3.0.CO;2-C

Kesner, R. P., Gilbert, P. E., & Barua, L. A. (2002). The role of the hippocampus in memory for the temporal order of a sequence of odors. *Behavioral Neuroscience*, *116*(2), 286–290. doi:10.1037/0735-7044.116.2.286

Kesner, R. P., & Rolls, E. T. (2015). A computational theory of hippocampal function, and tests of the theory: New developments. *Neuroscience & Biobehavioral Reviews*, *48*, 92–147. doi:10.1016/j.neubiorev.2014.11.009

Kim, S. S., Rouault, H., Druckmann, S., & Jayaraman, V. (2017). Ring attractor dynamics in the *Drosophila* central brain. *Science*, *356*(6340), 849–853. doi:10.1126/science.aal4835

Kleinfeld, D. (1986). Sequential state generation by model neural networks. *Proceedings of the National Academy of Sciences*, *83*(24), 9469–9473. doi:10.1073/pnas.83.24.9469

Knowles, W. D. (1992). Normal Anatomy and Neurophysiology of the Hippocampal Formation. *Journal of Clinical Neurophysiology*, *9*(2), 253–263. doi:10.1097/00004691-199204010-00006

Krebs, J. R., Sherry, D. F., Healy, S. D., Perry, V. H., & Vaccarino, A. L. (1989). Hippocampal specialization of food-storing birds. *Proceedings of the National Academy of Sciences*, *86*(4), 1388–1392. doi:10.1073/pnas.86.4.1388

Krebs, J. R., Clayton, N. S., Healy, S. D., Cristol, D. A., Patel, S. N., & Jolliffe, A. R. (1996). The ecology of the avian brain: Food-storing memory and the hippocampus. *Ibis*, *138*(1), 34–46. doi:10.1111/j.1474-919X.1996.tb04311.x

Krieg, W. J. S. (1955). *Brain mechanisms in diachrome*. Evanston, IL: Brain Books.

Kropff, E., Yang, S. M., & Schinder, A. F. (2015). Dynamic role of adult-born dentate granule cells in memory processing. *Current Opinion in Neurobiology*, *35*, 21–26. doi:10.1016/j.conb.2015.06.002

Kubie, J. L., & Muller, R. U. (1991). Multiple representations in the hippocampus. *Hippocampus*, *1*(3), 240–242. doi:10.1002/hipo.450010305

Kuhn, H., Dickinson-Anson, H., & Gage, F. (1996). Neurogenesis in the dentate gyrus of the adult rat: Age-related decrease of neuronal progenitor proliferation. *The Journal of Neuroscience*, *16*(6), 2027–2033. doi:10.1523/JNEUROSCI.16-06-02027.1996

Lassalle, J.-M., Bataille, T., & Halley, H. (2000). Reversible Inactivation of the Hippocampal Mossy Fiber Synapses in Mice Impairs Spatial Learning, but neither Consolidation nor Memory Retrieval, in the Morris Navigation Task. *Neurobiology of Learning and Memory*, *73*(3), 243–257. doi:10.1006/nlme.1999.3931

Lee, I., & Kesner, R. P. (2004). Encoding versus retrieval of spatial memory: Double dissociation between the dentate gyrus and the perforant path inputs into CA3 in the dorsal hippocampus. *Hippocampus*, *14*(1), 66–76. doi:10.1002/hipo.10167

Lee, I., Yoganarasimha, D., Rao, G., & Knierim, J. J. (2004). Comparison of population coherence of place cells in hippocampal subfields CA1 and CA3. *Nature*, *430*(6998), 456–459. doi:10.1038/nature02739

Lein, E. S., Callaway, E. M., Albright, T. D., & Gage, F. H. (2005). Redefining the boundaries of the hippocampal CA2 subfield in the mouse using gene expression and 3-dimensional reconstruction. *Journal of Comparative Neurology*, *485*(1), 1–10. doi:10.1002/cne.20426

Leutgeb, J. K., Leutgeb, S., Moser, M.-B., & Moser, E. I. (2007). Pattern Separation in the Dentate Gyrus and CA3 of the Hippocampus. *Science*, *315*(5814), 961–966. doi:10.1126/science.1135801

Leutgeb, S., Leutgeb, J. K., Barnes, C. A., Moser, E. I., McNaughton, B. L., & Moser, M.-B. (2005). Independent Codes for Spatial and Episodic Memory in Hippocampal Neuronal Ensembles. *Science*, *309*(5734), 619–623. doi:10.1126/science.1114037

Leutgeb, S., Leutgeb, J. K., Treves, A., Moser, M.-B., & Moser, E. I. (2004). Distinct Ensemble Codes in Hippocampal Areas CA3 and CA1. *Science*, *305*(5688), 1295–1298. doi:10.1126/science.1100265

Levine, B., Svoboda, E., Hay, J. F., Winocur, G., & Moscovitch, M. (2002). Aging and autobiographical memory: Dissociating episodic from semantic retrieval. *Psychology and Aging*, *17*(4), 677–689. doi:10.1037/0882-7974.17.4.677

Lisman, J. E. (1999). Relating Hippocampal Circuitry to Function. *Neuron*, *22*(2), 233–242. doi:10.1016/S0896-6273(00)81085-5

Lisman, J. E., & Jensen, O. (2013). The Theta-Gamma Neural Code. *Neuron*, *77*(6), 1002–1016. doi:10.1016/j.neuron.2013.03.007

Lisman, J. E., & Otmakhova, N. A. (2001). Storage, recall, and novelty detection of sequences by the hippocampus: Elaborating on the SOCRATIC model to account for normal and aberrant effects of dopamine. *Hippocampus*, *11*(5), 551–568. doi:10.1002/hipo.1071

Lømo, T. (1971). Patterns of activation in a monosynaptic cortical pathway: The perforant path input to the dentate area of the hippocampal formation. *Experimental Brain Research*, *12*(1). doi:10.1007/BF00234414

Lörincz, A., & Buzsáki, G. (2000). Two-Phase Computational Model Training Long-Term Memories in the Entorhinal-Hippocampal Region. *Annals of the New York Academy of Sciences*, *911*(1), 83–111. doi:10.1111/j.1749-6632.2000.tb06721.x

Louie, K., & Wilson, M. A. (2001). Temporally Structured Replay of Awake Hippocampal Ensemble Activity during Rapid Eye Movement Sleep. *Neuron*, *29*(1), 145–156. doi:10.1016/S0896-6273(01)00186-6

Macrides, F., Eichenbaum, H. B., & Forbes, W. B. (1982). Temporal Relationship between Sniffing and the Limbic Rhythm during Odor Discrimination Reversal Learning. *The Journal for Neuroscience*, *2*(12), 1705–1717. doi:10.1523/JNEUROSCI.02-12-01705.1982

Maguire, E. A., & Mullally, S. L. (2013). The hippocampus: A manifesto for change. *Journal of Experimental Psychology: General*, *142*(4), 1180–1189. doi:10.1037/a0033650

Mainali, N., Azeredo Da Silveira, R., & Burak, Y. (2025). Universal statistics of hippocampal place fields across species and dimensionalities. *Neuron*, *113*, 1–11. doi:10.1016/j.neuron.2025.01.017

Maliković, J., Amrein, I., Vinciguerra, L., Lalošević, D., Wolfer, D. P., & Slomianka, L. (2023). Cell numbers in the reflected blade of CA3 and their relation to other hippocampal principal cell populations across seven species. *Frontiers in Neuroanatomy*, *16*, 1070035. doi:10.3389/fnana.2022.1070035

Manns, J. R., Hopkins, R. O., Reed, J. M., Kitchener, E. G., & Squire, L. R. (2003). Recognition Memory and the Human Hippocampus. *Neuron*, *37*(1), 171–180. doi:10.1016/S0896-6273(02)01147-9

Marr, D. (1971). Simple Memory: A Theory for Archicortex. *Philos Trans R Soc Lond B Biol Sci*, *262*(841), 23–81. doi:10.1098/rstb.1971.0078

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why There Are Complementary Learning Systems in the Hippocampus and Neocortex: Insights From the Successesand Failuresof Connectionist Models of Learning and Memory. *Psychological review*, *102*(3), 419–457. doi:10.1037/0033-295X.102.3.419

McClelland, J. L., & Goddard, N. H. (1996). Considerations arising from a complementary learning systems perspective on hippocampus and neocortex. *Hippocampus*, *6*(6), 654–665. doi:10.1002/(SICI)1098-1063(1996)6:6<654::AID-HIPO8>3.0.CO;2-G

McDonough, J. H., & Kesner, R. P. (1971). Amnesia produced by brief electrical stimulation of amygdala or dorsal hippocampus in cats. *Journal of Comparative and Physiological Psychology*, *77*(1), 171–178. doi:10.1037/h0031575

McNaughton, B. L., Barnes, C. A., Gerrard, J. L., Gothard, K., Jung, M. W., Knierim, J. J., ... Weaver, K. L. (1996). Deciphering The Hippocampal Polyglot: The Hippocampus as a Path Integration System. *Journal of Experimental Biology*, *199*(1), 173–185. doi:10.1242/jeb.199.1.173

McNaughton, B. L., & Morris, R. G. M. (1987). Hippocampal synaptic enhancement and information storage within a distributed memory system. *10*(10), 408–415. doi:10.1016/0166-2236(87)90011-7

Melzer, S., Michael, M., Caputi, A., Eliava, M., Fuchs, E. C., Whittington, M. A., & Monyer, H. (2012). Long-Range–Projecting GABAergic Neurons Modulate Inhibition in Hippocampus and Entorhinal Cortex. *Science*, *335*(6075), 1506–1510. doi:10.1126/science.1217139

Mercurio, S., Alberti, C., Serra, L., Meneghini, S., Berico, P., Bertolini, J., ... Nicolis, S. K. (2021). An early Sox2-dependent gene expression programme required for

hippocampal dentate gyrus development. *Open Biology*, *11*(2), 200339. doi:10.1098/rsob.200339

Minsky, M., & Papert, S. A. (1969). *Perceptrons: An introduction to computational geometry*. Cambridge, MA: MIT Press.

Mishkin, M. (1978). Memory in monkeys severely impaired by combined but not by separate removal of amygdala and hippocampus. *Nature*, *273*(5660), 297–298. doi:10.1038/273297a0

Mizumori, S., Ragozzino, K., Cooper, B., & Leutgeb, S. (1999). Hippocampal Representational Organization and Spatial Context. *Hippocampus*, *9*(4), 444–451. doi:10.1002/(SICI)1098-1063(1999)9:4<444::AID-HIPO10>3.0.CO;2-Z

Mizumori, S., McNaughton, B., Barnes, C., & Fox, K. (1989). Preserved spatial coding in hippocampal CA1 pyramidal cells during reversible suppression of CA3c output: Evidence for pattern completion in hippocampus. *The Journal of Neuroscience*, *9*(11), 3915–3928. doi:10.1523/JNEUROSCI.09-11-03915.1989

Mori, M., Abegg, M. H., Gähwiler, B. H., & Gerber, U. (2004). A frequency-dependent switch from inhibition to excitation in a hippocampal unitary circuit. *Nature*, *431*(7007), 453–456. doi:10.1038/nature02854

Morris, A. M., Churchwell, J. C., Kesner, R. P., & Gilbert, P. E. (2012). Selective lesions of the dentate gyrus produce disruptions in place learning for adjacent spatial locations. *Neurobiology of Learning and Memory*, *97*(3), 326–331. doi:10.1016/j.nlm.2012.02.005

Morris, R. G. M. (1989). Synaptic plasticity and learning: Selective impairment of learning rats and blockade of long-term potentiation in vivo by the N-methyl-D- aspartate receptor antagonist AP5. *The Journal of Neuroscience*, *9*(9), 3040–3057. doi:10.1523/JNEUROSCI.09-09-03040.1989

Morris, R. G. M., Garrud, P., Rawlins, J. A., & O'Keefe, J. (1982). Place navigation impaired in rats with hippocampal lesions. *Nature*, *297*(5868), 681–683. doi:10.1038/297681a0

Moscovitch, M. (2008). The hippocampus as a "stupid," domain-specific module: Implications for theories of recent and remote memory, and of imagination. *Canadian Journal of Experimental Psychology / Revue canadienne de psychologie expérimentale*, *62*(1), 62–79. doi:10.1037/1196-1961.62.1.62

Moser, M. B., Moser, E. I., Forrest, E., Andersen, P., & Morris, R. G. (1995). Spatial learning with a minislab in the dorsal hippocampus. *Proceedings of the National Academy of Sciences*, *92*(21), 9697–9701. doi:10.1073/pnas.92.21.9697

Muller, R. U., Stead, M., & Pach, J. (1996). The hippocampus as a cognitive graph. *The Journal of general physiology*, *107*(6), 663–694. doi:10.1085/jgp.107.6.663

Muller, R., & Kubie, J. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *The Journal of Neuroscience*, *7*(7), 1951–1968. doi:10.1523/JNEUROSCI.07-07-01951.1987

Murre, J. M. (1996). TraceLink: A model of amnesia and consolidation of memory. *Hippocampus*, *6*(6), 675–684. doi:10.1002/(SICI)1098-1063(1996)6:6<675::AID-HIPO10>3.0.CO;2-Y

Naber, P. A., Witter, M. P., & Lopes Da Silva, F. H. (2000). Networks of the Hippocampal Memory System of the Rat: The Pivotal Role of the Subiculum. *An-*

nals of the New York Academy of Sciences, *911*(1), 392–403. doi:10.1111/j.1749-6632.2000.tb06739.x

Niedermeyer, E. (2008). Hippocampic Theta Rhythm. *Clinical EEG and Neuroscience*, *39*(4). doi:10.1177/155005940803900408

O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, *34*(1), 171–175. doi:10.1016/0006-8993(71)90358-1

O'Keefe, J. (1990). Chapter 22 A computational theory of the hippocampal cognitive map. In *Progress in Brain Research* (Vol. 83, pp. 301–312). doi:10.1016/S0079-6123(08)61258-3

O'Keefe, J., & Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. *Nature*, *381*(6581), 425–428. doi:10.1038/381425a0

O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford : New York: Clarendon Press ; Oxford University Press.

O'Keefe, J., & Recce, M. L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*, *3*(3), 317–330. doi:10.1002/hipo.450030307

O'Reilly, R. C., & McClelland, J. L. (1994). Hippocampal conjunctive encoding, storage, and recall: Avoiding a trade-off. *Hippocampus*, *4*(6), 661–682. doi:10.1002/hipo.450040605

Ocaña, F. M., Uceda, S., Arias, J. L., Salas, C., & Rodríguez, F. (2017). Dynamics of Goldfish Subregional Hippocampal Pallium Activity throughout Spatial Memory Formation. *Brain, Behavior and Evolution*, *90*(2), 154–170. doi:10.1159/000478843

Omer, D. B., Maimon, S. R., Las, L., & Ulanovsky, N. (2018). Social place-cells in the bat hippocampus. *Science*, *359*(6372), 218–224. doi:10.1126/science.aao3474

Orman, R., Von Gizycki, H., Lytton, W., & Stewart, M. (2008). Local axon collaterals of area CA1 support spread of epileptiform discharges within CA1, but propagation is unidirectional. *Hippocampus*, *18*(10), 1021–1033. doi:10.1002/hipo.20460

Papp, G., & Treves, A. (2008). Network Analysis of the Significance of Hippocampal Subfields. In S. J. Y. Mizumori (Ed.), *Hippocampal Place Fields: Relevance to Learning and Memory* (pp. 328–342). New York: Oxford University Press.

Park, E., Dvorak, D., & Fenton, A. A. (2011). Ensemble Place Codes in Hippocampus: CA1, CA3, and Dentate Gyrus Place Cells Have Multiple Place Fields in Large Environments. *PLoS ONE*, *6*(7), e22349. doi:10.1371/journal.pone.0022349

Pereira, A. C., Huddleston, D. E., Brickman, A. M., Sosunov, A. A., Hen, R., McKhann, G. M., … Small, S. A. (2007). An *in vivo* correlate of exercise-induced neurogenesis in the adult dentate gyrus. *Proceedings of the National Academy of Sciences*, *104*(13), 5638–5643. doi:10.1073/pnas.0611721104

Pfeiffer, B. E., & Foster, D. J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*, *497*(7447), 74–79. doi:10.1038/nature12112

Raisman, G., Cowan, W. M., & Powell, T. P. S. (1965). The extrinsic afferent, commissural and association fibres of the hippocampus. *Brain*, *88*(5), 963–996. doi:10.1093/brain/88.5.963

Ramón y Cajal, S. (1893). Estructura del asta de Ammon y fascia dentata. *Ann. Soc. Esp. Hist. Nat. (Madrid)*, *22*, 53–114.

Ramón y Cajal, S. (1911). *Histologie du Systeme Nerveux de l'Homme et des Vertebres Vol. II*. Madrid: CSIC reprint 1954.

Rapp, P. R., & Gallagher, M. (1996). Preserved neuron number in the hippocampus of aged rats with spatial learning deficits. *Proceedings of the National Academy of Sciences*, *93*(18), 9926–9930. doi:10.1073/pnas.93.18.9926

Reiter, S., Liaw, H.-P., Yamawaki, T. M., Naumann, R. K., & Laurent, G. (2017). On the Value of Reptilian Brains to Map the Evolution of the Hippocampal Formation. *Brain, Behavior and Evolution*, *90*(1), 41–52. doi:10.1159/000478693

Rich, P. D., Liaw, H.-P., & Lee, A. K. (2014). Large environments reveal the statistical structure governing hippocampal representations. *Science*, *345*(6198), 814–817. doi:10.1126/science.1255635

Robertson, R. G., Rolls, E. T., Georges-François, P., & Panzeri, S. (1999). Head direction cells in the primate pre-subiculum. *Hippocampus*, *9*(3), 206–219. doi:10.1002/(SICI)1098-1063(1999)9:3<206::AID-HIPO2>3.0.CO;2-H

Rodríguez, F., López, J., Vargas, J., Broglio, C., Gómez, Y., & Salas, C. (2002). Spatial memory and hippocampal pallium through vertebrate evolution: Insights from reptiles and teleost fish. *Brain Research Bulletin*, *57*(3-4), 499–503. doi:10.1016/S0361-9230(01)00682-7

Rodríguez, F., López, J. C., Vargas, J. P., Gómez, Y., Broglio, C., & Salas, C. (2002). Conservation of Spatial Memory Function in the Pallial Forebrain of Reptiles and Ray-Finned Fishes. *The Journal of Neuroscience*, *22*(7), 2894–2903. doi:10.1523/JNEUROSCI.22-07-02894.2002

Rogers Flattery, C. N., Rosen, R. F., Farberg, A. S., Dooyema, J. M., Hof, P. R., Sherwood, C. C., … Preuss, T. M. (2020). Quantification of neurons in the hippocampal formation of chimpanzees: Comparison to rhesus monkeys and humans. *Brain Structure and Function*, *225*(8), 2521–2531. doi:10.1007/s00429-020-02139-x

Rollenhagen, A., Sätzler, K., Rodríguez, E. P., Jonas, P., Frotscher, M., & Lübke, J. H. R. (2007). Structural Determinants of Transmission at Large Hippocampal Mossy Fiber Synapses. *The Journal of Neuroscience*, *27*(39), 10434–10444. doi:10.1523/JNEUROSCI.1946-07.2007

Rolls, E. T. (1987). Information representation, processing and storage in the brain: Analysis at the single neuron level. In J.-P. Changeux & M. Konishi (Eds.), *The Neural and Molecular Bases of Learning* (pp. 503–540). New York: Wiley.

Rolls, E. T., & Tovee, M. J. (1994). Processing Speed in the Cerebral Cortex and the Neurophysiology of Visual Masking. *Proc. R. Soc. Lond. B*, *257*, 9–15. doi:10.1098/rspb.1994.0087

Rolls, E. T. (1989). Functions of Neuronal Networks in the Hippocampus and Neocortex in Memory. In *Neural Models of Plasticity* (pp. 240–265). doi:10.1016/B978-0-12-148955-7.50017-5

Rolls, E. T. (1999). Spatial view cells and the representation of place in the primate hippocampus. *Hippocampus*, *9*(4), 467–480. doi:10.1002/(SICI)1098-1063(1999)9:4<467::AID-HIPO13>3.0.CO;2-F

Rolls, E. T. (2000). Hippocampo-cortical and cortico-cortical backprojections. *Hippocampus*, *10*(4), 380–388. doi:10.1002/1098-1063(2000)10:4<380::AID-HIPO4>3.0.CO;2-0

Rolls, E. T., Yan, X., Deco, G., Zhang, Y., Jousmaki, V., & Feng, J. (2024). A ventromedial visual cortical 'Where' stream to the human hippocampus for spatial scenes revealed with magnetoencephalography. *Communications Biology*, *7*(1), 1047. doi:10.1038/s42003-024-06719-z

Rolls, E., Tovee, M., Purcell, D., Stewart, A., & Azzopardi, P. (1994). The responses of neurons in the temporal cortex of primates, and face identification and detection. *Experimental Brain Research*, *101*(3). doi:10.1007/BF00227340

Ryan, L., Nadel, L., Keil, K., Putnam, K., Schnyer, D., Trouard, T., & Moscovitch, M. (2001). Hippocampal complex and retrieval of recent and very remote autobiographical memories: Evidence from functional magnetic resonance imaging in neurologically intact people. *Hippocampus*, *11*(6), 707–714. doi:10.1002/hipo.1086

Ryom, K. I., Basu, A., Stendardi, D., Ciaramelli, E., & Treves, A. (2024). Taking time to compose thoughts with prefrontal schemata. *Experimental Brain Research*, *242*(5), 1101–1114. doi:10.1007/s00221-024-06785-z

Ryom, K. I., Stendardi, D., Ciaramelli, E., & Treves, A. (2023). Computational constraints on the associative recall of spatial scenes. *Hippocampus*, *33*(5), 635–645. doi:10.1002/hipo.23511

Samsonovich, A., & McNaughton, B. L. (1997). Path Integration and Cognitive Mapping in a Continuous Attractor Neural Network Model. *The Journal of Neuroscience*, *17*(15), 5900–5920. doi:10.1523/JNEUROSCI.17-15-05900.1997

Sanders, H. I., & Warrington, E. K. (1971). Memory for remote events in amnesic patients. *Brain*, *94*(4), 661–668. doi:10.1093/brain/94.4.661

Schönsberg, F., Monasson, R., & Treves, A. (2024). Continuous Quasi-Attractors dissolve with too much – or too little – variability. *PNAS Nexus*, *3*(12), pgae525. doi:10.1093/pnasnexus/pgae525

Scoville, W. B., & Milner, B. (1957). Loss of Recent Memory after bilateral Hippocampal Lesions. *Journal of neurology, neurosurgery, and psychiatry*, *20*(1), 11–21. doi:10.1136/jnnp.20.1.11

Seress, L. (1988). Interspecies comparison of the hippocampal formation shows increased emphasis on the regio superior in the Ammon's horn of the human brain. *Journal fur Hirnforschung*, *29*(3), 335–440.

Sharon, T., Moscovitch, M., & Gilboa, A. (2011). Rapid neocortical acquisition of long-term arbitrary associations independent of the hippocampus. *Proceedings of the National Academy of Sciences*, *108*(3), 1146–1151. doi:10.1073/pnas.1005238108

Sheintuch, L., Geva, N., Baumer, H., Rechavi, Y., Rubin, A., & Ziv, Y. (2020). Multiple Maps of the Same Spatial Context Can Stably Coexist in the Mouse Hippocampus. *Current Biology*, *30*(8), 1467–1476.e6. doi:10.1016/j.cub.2020.02.018

Sheintuch, L., Geva, N., Deitch, D., Rubin, A., & Ziv, Y. (2023). Organization of hippocampal CA3 into correlated cell assemblies supports a stable spatial code. *Cell Reports*, *42*(2), 112119. doi:10.1016/j.celrep.2023.112119

Smulders, T. V., Shiflett, M. W., Sperling, A. J., & DeVoogd, T. J. (2000). Seasonal changes in neuron numbers in the hippocampal formation of a food-hoarding bird: The black-capped chickadee. *Journal of Neurobiology*, *44*(4), 414–422. doi:10.1002/1097-4695(20000915)44:4<414::AID-NEU4>3.0.CO;2-I

Sompolinsky, H., & Kanter, I. (1986). Temporal Association in Asymmetric Neural Networks. *Physical Review Letters*, *57*(22), 2861–2864. doi:10.1103/PhysRevLett.57.2861

Spalding, K. L., Bergmann, O., Alkass, K., Bernard, S., Salehpour, M., Huttner, H. B., … Frisén, J. (2013). Dynamics of Hippocampal Neurogenesis in Adult Humans. *Cell*, *153*(6), 1219–1227. doi:10.1016/j.cell.2013.05.002

Spalla, D., Cornacchia, I. M., & Treves, A. (2021). Continuous attractors for dynamic memories. *eLife*, *10*, e69499. doi:10.7554/eLife.69499

Spalla, D., Treves, A., & Boccara, C. N. (2022). Angular and linear speed cells in the parahippocampal circuits. *Nature Communications*, *13*(1), 1907. doi:10.1038/s41467-022-29583-z

Spiers, H. J., Maguire, E. A., & Burgess, N. (2001). Hippocampal Amnesia. *Neurocase*, *7*(5), 357–382. doi:10.1076/neur.7.5.357.16245

Stefanini, F., Kushnir, L., Jimenez, J. C., Jennings, J. H., Woods, N. I., Stuber, G. D., … Fusi, S. (2020). A Distributed Neural Code in the Dentate Gyrus and in CA1. *Neuron*, *107*(4), 703–716.e4. doi:10.1016/j.neuron.2020.05.022

Striedter, G. F. (2016). Evolution of the hippocampus in reptiles and birds. *Journal of Comparative Neurology*, *524*(3), 496–517. doi:10.1002/cne.23803

Taube, J., Muller, R., & Ranck, J. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. II. Effects of environmental manipulations. *The Journal of Neuroscience*, *10*(2), 436–447. doi:10.1523/JNEUROSCI.10-02-00436.1990

Taube, S. (1995). Head Direction Cells Recorded in the Anterior Thalamic Nuclei of Freely Moving Rats. *Journal of Neuroscience*, *15*(1), 70–86. doi:10.1523/JNEUROSCI.15-01-00070.1995

Teyler, T. J., & DiScenna, P. (1986). The Hippocampal Memory Indexing Theory. *Behavioral neuroscience*, *100*(2), 147–152. doi:10.1037//0735-7044.100.2.147

Thompson, C. L., Pathak, S. D., Jeromin, A., Ng, L. L., MacPherson, C. R., Mortrud, M. T., … Lein, E. S. (2008). Genomic Anatomy of the Hippocampus. *Neuron*, *60*(6), 1010–1021. doi:10.1016/j.neuron.2008.12.008

Thomson, A. M., & Radpour, S. (1991). Excitatory Connections Between CA1 Pyramidal Cells Revealed by Spike Triggered Averaging in Slices of Rat Hippocampus are Partially NMDA Receptor Mediated. *European Journal of Neuroscience*, *3*(6), 587–601. doi:10.1111/j.1460-9568.1991.tb00845.x

Thorpe, S., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *Nature*, *381*(6582), 520–522. doi:10.1038/381520a0

Thorpe, S. J., & Imbert, M. (1989). Biological Constraints on Connectionist Modelling. *Connectionism in perspective*, 63–92.

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, *55*(4), 189–208. doi:10.1037/h0061626

Torromino, G., Loffredo, V., Cavezza, D., Sonsini, G., Esposito, F., Crevenna, A. H., … De Leonibus, E. (2022). Thalamo-hippocampal pathway regulates incidental memory capacity in mice. *Nature Communications*, *13*(1), 4194. doi:10.1038/s41467-022-31781-8

Tosches, M. A., Yamawaki, T. M., Naumann, R. K., Jacobi, A. A., Tushev, G., & Laurent, G. (2018). Evolution of pallium, hippocampus, and cortical cell types

revealed by single-cell transcriptomics in reptiles. *Science*, *360*(6391), 881–888. doi:10.1126/science.aar4237

Treves, A. (1990). Threshold-linear formal neurons in auto-associative nets. *Journal of Physics A: Mathematical and General*, *23*(12), 2631–2650. doi:10.1088/0305-4470/23/12/037

Treves, A., & Amit, D. J. (1989). Low firing rates: An effective Hamiltonian for excitatory neurons. *Journal of Physics A: Mathematical and General*, *22*(12), 2205–2226. doi:10.1088/0305-4470/22/12/020

Treves, A., Tashiro, A., Witter, M., & Moser, E. (2008). What is the mammalian dentate gyrus good for? *Neuroscience*, *154*(4), 1155–1172. doi:10.1016/j.neuroscience.2008.04.073

Treves, A. (1993). Mean-field analysis of neuronal spike dynamics. *Network: Computation in Neural Systems*, *4*(3), 259–284. doi:10.1088/0954-898X_4_3_002

Treves, A. (1995). Quantitative estimate of the information relayed by the Schaffer collaterals. *Journal of Computational Neuroscience*, *2*(3), 259–272. doi:10.1007/BF00961437

Treves, A. (2004). Computational constraints between retrieving the past and predicting the future, and the CA3-CA1 differentiation. *Hippocampus*, *14*(5), 539–556. doi:10.1002/hipo.10187

Treves, A. (2005). Frontal latching networks: A possible neural basis for infinite recursion. *Cognitive Neuropsychology*, *22*(3-4), 276–291. doi:10.1080/02643290442000329

Treves, A., Georges-Francois, P., Panzeri, S., Robertson, R. G., & Rolls, E. T. (1998). The Metric Content of Spatial Views as Represented in the Primate Hippocampus. In V. Torre & J. Nicholls (Eds.), *Neural Circuits and Networks: Proceedings of the NATO advanced Study Institute on Neuronal Circuits and Networks, held at the Ettore Majorana Center, Erice, Italy, June 15–27 1997* (pp. 239–247). doi:10.1007/978-3-642-58955-3_17

Treves, A., & Rolls, E. T. (1991). What determines the capacity of autoassociative memories in the brain? *Network: Computation in Neural Systems*, *2*(4), 371–397. doi:10.1088/0954-898X_2_4_004

Treves, A., & Rolls, E. T. (1992). Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. *Hippocampus*, *2*(2), 189–199. doi:10.1002/hipo.450020209

Treves, A., & Rolls, E. T. (1994). Computational analysis of the role of the hippocampus in memory. *Hippocampus*, *4*(3), 374–391. doi:10.1002/hipo.450040319

Treves, A., Rolls, E. T., & Tovee, M. J. (1996). On the Time Required for Recurrent Processing in the Brain. In V. Torre & F. Conti (Eds.), *Neurobiology* (pp. 371–382). doi:10.1007/978-1-4615-5899-6_29

Tsodyks, M. V., & Sejnowski, T. J. (1995a). Associative memory and hippocampal place cells. *International Journal of Neural Systems*, *6*(Supp. 1995), 81–86.

Tsodyks, M. V., & Sejnowski, T. J. (1995b). Rapid State Switching in Balanced Cortical Network Models. *Network: Computation in Neural Systems*, *6*, 111–124. doi:10.1088/0954-898X/6/2/001

Tsodyks, M. V., & Feigelman, M. V. (1988). The Enhanced Storage Capacity in Neural Networks with Low Activity Level. *Europhysics Letters (EPL)*, *6*(2), 101–105. doi:10.1209/0295-5075/6/2/002

Tulving, E. (1972). Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.), *Organization of memory* (pp. 381–403). Oxford: Academic Press.

Valentino, R., & Dingledine, R. (1981). Presynaptic inhibitory effect of acetylcholine in the hippocampus. *The Journal of Neuroscience*, *1*(7), 784–792. doi:10.1523/JNEUROSCI.01-07-00784.1981

Van Dijk, R. M., Huang, S.-H., Slomianka, L., & Amrein, I. (2016). Taxonomic Separation of Hippocampal Networks: Principal Cell Populations and Adult Neurogenesis. *Frontiers in Neuroanatomy*, *10*. doi:10.3389/fnana.2016.00022

Van Praag, H. (2008). Neurogenesis and Exercise: Past and Future Directions. *Neuromolecular Medicine*, *10*(2), 128–140. doi:10.1007/s12017-008-8028-z

Vanderwolf, C. (2001). The hippocampus as an olfacto-motor mechanism: Were the classical anatomists right after all? *Behavioural Brain Research*, *127*(1-2), 25–47. doi:10.1016/S0166-4328(01)00354-0

Vargha-Khadem, F., Gadian, D. G., Watkins, K. E., Connelly, A., Van Paesschen, W., & Mishkin, M. (1997). Differential Effects of Early Hippocampal Pathology on Episodic and Semantic Memory. *Science*, *277*(5324), 376–380. doi:10.1126/science.277.5324.376

Vazdarjanova, A., & Guzowski, J. F. (2004). Differences in Hippocampal Neuronal Population Responses to Modifications of an Environmental Context: Evidence for Distinct, Yet Complementary, Functions of CA3 and CA1 Ensembles. *The Journal of Neuroscience*, *24*(29), 6489–6496. doi:10.1523/JNEUROSCI.0350-04.2004

Vinogradova, O. S., & Dudaeva, K. I. (1972). On the comparator function of the hippocampus. *Doklady Akademii Nauk SSSR*, *202*(2), 486–489.

Vinogradova, O. (2001). Hippocampus as comparator: Role of the two input and two output systems of the hippocampus in selection and registration of information. *Hippocampus*, *11*(5), 578–598. doi:10.1002/hipo.1073

Vyleta, N. P., Borges-Merjane, C., & Jonas, P. (2016). Plasticity-dependent, full detonation at hippocampal mossy fiber–CA3 pyramidal neuron synapses. *eLife*, *5*, e17977. doi:10.7554/eLife.17977

Wallenstein, G. V., Hasselmo, M. E., & Eichenbaum, H. (1998). The hippocampus as an associator of discontiguous events. *Trends in Neurosciences*, *21*(8), 317–323. doi:10.1016/S0166-2236(97)01220-4

Watson, J. F., Vargas-Barroso, V., Morse-Mora, R. J., Navas-Olive, A., Tavakoli, M. R., Danzl, J. G., ... Jonas, P. (2024, May). Human hippocampal CA3 uses specific functional connectivity rules for efficient associative memory. doi:10.1101/2024.05.02.592169

West, M. J., & Gundersen, H. J. G. (1990). Unbiased stereological estimation of the number of neurons in the human hippocampus. *Journal of Comparative Neurology*, *296*(1), 1–22. doi:10.1002/cne.902960102

Wilson, H. R., & Cowan, J. D. (1973). A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Kybernetik*, *13*(2), 55–80. doi:10.1007/BF00288786

Wimmer, K., Nykamp, D. Q., Constantinidis, C., & Compte, A. (2014). Bump attractor dynamics in prefrontal cortex explains behavioral precision in spatial working memory. *Nature Neuroscience*, *17*(3), 431–439. doi:10.1038/nn.3645

Wiskott, L., Rasch, M. J., & Kempermann, G. (2006). A functional hypothesis for adult hippocampal neurogenesis: Avoidance of catastrophic interference in the dentate gyrus. *Hippocampus*, *16*(3), 329–343. doi:10.1002/hipo.20167

Witter, M. P. (1993). Organization of the entorhinal—hippocampal system: A review of current anatomical data. *Hippocampus*, *3*(S1), 33–44. doi:10.1002/hipo.1993.4500030707

Yartsev, M. M., Witter, M. P., & Ulanovsky, N. (2011). Grid cells without theta oscillations in the entorhinal cortex of bats. *Nature*, *479*(7371), 103–107. doi:10.1038/nature10583

Zhang, K. (1996). Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory. *The Journal of Neuroscience*, *16*(6), 2112–2126. doi:10.1523/JNEUROSCI.16-06-02112.1996

Zheng, J., Yebra, M., Schjetnan, A. G. P., Patel, K., Katz, C. N., Kyzar, M., … Rutishauser, U. (2024). Theta phase precession supports memory formation and retrieval of naturalistic experience in humans. *Nature Human Behaviour*. doi:10.1038/s41562-024-01983-9

Ziv, Y., Burns, L. D., Cocker, E. D., Hamel, E. O., Ghosh, K. K., Kitch, L. J., … Schnitzer, M. J. (2013). Long-term dynamics of CA1 hippocampal place codes. *Nature Neuroscience*, *16*(3), 264–266. doi:10.1038/nn.3329

Zola, S. M., & Squire, L. R. (2001). Relationship between magnitude of damage to the hippocampus and impaired recognition memory in monkeys. *Hippocampus*, *11*(2), 92–98. doi:10.1002/hipo.1027