

UC Berkeley

UC Berkeley Previously Published Works

Title

Flexible weighting of diverse inputs makes hippocampal function malleable

Permalink

<https://escholarship.org/uc/item/9zc7w28m>

Authors

Aly, Mariam

Turk-Browne, Nicholas B

Publication Date

2018-07-01

DOI

10.1016/j.neulet.2017.05.063

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial-NoDerivatives License, available at

<https://creativecommons.org/licenses/by-nc-nd/4.0/>

Peer reviewed



Published in final edited form as:

Neurosci Lett. 2018 July 27; 680: 13–22. doi:10.1016/j.neulet.2017.05.063.

Flexible weighting of diverse inputs makes hippocampal function malleable

Mariam Aly^{a,b,*} and Nicholas B. Turk-Browne^{a,c,d}

^aPrinceton Neuroscience Institute, Princeton University, Princeton, NJ, 08544

^bDepartment of Psychology, Columbia University, New York, NY, 10027

^cDepartment of Psychology, Princeton University, Princeton, NJ, 08544

^dDepartment of Psychology, Yale University, New Haven, CT, 06520

Abstract

Classic theories of hippocampal function have emphasized its role as a dedicated memory system, but recent research has shown that it contributes broadly to many aspects of cognition, including attention and perception. We propose that the reason the hippocampus plays such a broad role in cognition is that its function is particularly malleable. We argue that this malleability arises because the hippocampus receives diverse anatomical inputs and these inputs are flexibly weighted based on behavioral goals. We discuss examples of how hippocampal representations can be flexibly weighted, focusing on hippocampal modulation by attention. Finally, we suggest some general neural mechanisms and core hippocampal computations that may enable the hippocampus to support diverse cognitive functions, including attention, perception, and memory. Together, this work suggests that great progress can and has been made in understanding the hippocampus by considering how the domain-general computations it performs allow it to dynamically contribute to many different behaviors.

Keywords

episodic memory; relational representations; pattern separation; pattern completion; match/mismatch detection

1. Introduction

The brain is hierarchically organized along multiple dimensions, including space, time, and features. For example, spatial receptive fields are smallest in primary visual cortex and increase in size throughout the ventral visual stream [1]. Likewise, primary visual cortex responds to simple features such as oriented edges, whereas downstream areas are driven by

*Please address correspondence to: Mariam Aly, Department of Psychology, 406 Schermerhorn Hall, Columbia University, New York, NY, 10027, 212-854-5318, ma3631@columbia.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

more complex feature combinations, such as shapes, objects, faces, and scenes, as well as information from other modalities [1–3]. Finally, information from the current environment is preferentially represented in early sensory areas, whereas information integrated over the past several seconds or minutes is represented in higher-order areas in posterior medial cortex [4].

The brain also contains a more abstract hierarchy of “malleability” — the extent to which neural function is fixed and immutable vs. influenced by other processes and variable across tasks. The retina is perhaps the least malleable: light signals are passed ballistically along a fixed route from photoreceptors to bipolar cells to ganglion cells; this pathway supports a very specific function (transduction of light and transmission to visual areas in the brain); and the direction that information is routed and what function it performs are unaffected by abstract factors such as memories, motivations, or goals.

Early visual cortex is relatively more malleable. On one hand, neurons respond to specific features and locations in a standard and reliable way that can be described by a unidimensional tuning curve. On the other hand, this tuning can be modulated by top-down or feedback processes such as expectation, reward, and attention [5–9].

The hippocampus is perhaps at the apex of this hierarchy of malleability, along with other regions like prefrontal cortex (which is beyond the scope of this paper, but will be discussed briefly in Section 4). For example, unlike cells in early visual cortex, individual hippocampal cells can have mixed selectivity, responding to multiple dimensions such as object identities, object positions, spatial contexts, and rewards [10–11]. Such mixed selectivity means that hippocampal cells can change how they respond as a function of task demands or goal states. By definition, this is a key aspect of malleability as we discussed above. The hippocampus is also malleable in that it contributes broadly to many cognitive domains (including attention, perception, working memory, long-term memory, and decision making), and does so in a manner influenced by our goals and motivational states [11–16].

We propose that the malleability of hippocampal processing is tied to the *diversity* and *flexibility* of its inputs: the myriad types of information it receives from other brain areas, and how flexibly the weights on those inputs can change as a function of behavioral goals. We argue that the hippocampus, far from being a dedicated memory system, can be configured to contribute to many functions, and that future progress will come by considering how these functions load on different computations implemented in the hippocampus (Figure 1).

2. Diversity of input: What features can the hippocampus represent?

The hippocampus receives indirect input (via medial temporal lobe cortex) from multiple sensory modalities, including vision, audition, somatosensation, and olfaction [1, 3]. It also receives input from prefrontal cortex (e.g., information about goals, task rules, and contexts), both directly [17], and indirectly via the nucleus reuniens of the thalamus [18–19] and medial temporal lobe cortex [3]. Furthermore, it is modulated by the dopaminergic, cholinergic and noradrenergic systems [20–24]. Given this diversity of anatomical inputs

and neuromodulation, it is no surprise that the hippocampus has been identified as a hub in brain networks [25–27].

The multiplicity of information relayed to the hippocampus makes it difficult to identify what “features” it represents. One dominant perspective — especially from the rodent literature — is that the hippocampus forms an allocentric (i.e., world-centered) representation of space, and its contribution to different domains of cognition can be understood via its role in representing spatial context [2, 28–30]. An alternative, but complementary, perspective is that the hippocampus is fundamentally relational, and thus the “features” it represents are the associations between objects, locations, spatial and temporal contexts, rewards, and actions [11, 13, 16, 31]. Thus, the hippocampus integrates multiple types of information, perhaps forming a conjunctive representation that subsumes all input features [25, 32–33], enabling the retrieval of associatively related information and reinstatement of episode-specific patterns of activity that were present during encoding [34–37].

These two views — purely spatial vs. conjunctive / relational — have been compared in detail elsewhere [14, 31]. In the current paper, we consider both spatial and (non-spatial) relational representations in the hippocampus.

3. Flexible weighting of input: How are hippocampal representations modulated?

Given some input to a brain area, flexibility refers to modulation of how that input is weighted and thus what information gets represented in that area. One variety of modulation results from the deployment of attention, which enables the selection and enhancement of stimulus features. Here, we focus on goal-directed attention — namely, how hippocampal representations are reconfigured by top-down attentional states.

There are at least two reasons to expect attention to modulate the hippocampus. First, the visual system is robustly modulated by attention [38], and the hippocampus is at the apex of the ventral visual processing stream [1]. There is no *a priori* reason to expect an abrupt transition in the brain between receptivity to attention and imperviousness to it. Second, behavioral expressions of memory are robustly modulated by attention [12], and the hippocampus is critical for memory encoding and retrieval. Thus, one possibility is that attentional modulation of memory occurs via its modulation of the hippocampus.

In visual cortex, attention increases the overall level of activity in brain areas coding for attended features, items, or locations [38]. Relying on these findings, some studies have investigated gain modulation in the hippocampus as a function of attention to items and/or locations [39–40]. Importantly, these studies investigated attentional modulation without any overt long-term memory demands, thus seeking to identify a signature of attention *per se* in the hippocampus, rather than attentional modulation of memory (c.f. [41–43]). Surprisingly, these studies failed to find attentional modulation of the overall level of activity in the hippocampus.

One possible interpretation of these findings is that attention simply does not modulate the hippocampus. Alternatively, those studies might have missed evidence of hippocampal modulation by attention because they examined the wrong neural measure (Figure 2). Specifically, in animal models, the task relevance of environmental cues does not seem to modulate the overall level of activity in the hippocampus, which is the classic attentional signature in visual cortex. In contrast, task relevance modulates the *stability* (or reliability) of hippocampal activity patterns [44–49]. For example, when spatial locations are task-relevant (vs. task-irrelevant), place fields in mouse hippocampus are more stable [45], and this stability increases as a function of how important spatial features are [46]. However, the task relevance of spatial information does not change the overall level of neural activity, i.e., firing rates [45–46]. Such modulation of stability is not limited to the spatial domain: when olfactory information is task-relevant (vs. task-irrelevant) hippocampal cells that respond to odor fire more reliably [45]. Finally, this stability of hippocampal representations has also been observed at the level of networks: making different spatial reference frames task relevant switches which cell assemblies are involved [47–49].

It is worth noting that manipulations of task relevance in animals are not equivalent to manipulations of goal-directed attention in humans, in the sense that animals are not given explicit goal instructions and instead must learn across many trials which features help achieve a more implicit goal of obtaining reward or avoiding punishment. Nevertheless, whether learned or instructed, we will work under the assumption that task relevance may manifest itself in similar ways in the hippocampus.

Inspired by this work in animal models, we set out to test whether attention modulates the stability of hippocampal activity patterns in humans, at the macro scale of voxels in fMRI as opposed to individual neurons and neural ensembles. Additionally, in contrast to prior studies of hippocampal attentional modulation in humans (e.g., [39–40]), which examined attention to individual items or locations, we manipulated attention to relatively high-level spatial and relational dimensions that are a key component of hippocampal processing [31].

While undergoing high-resolution fMRI, participants viewed several 3D-rendered rooms, each with a unique layout, multiple pieces of furniture, and a unique painting on the wall (Figure 3A–B; [50]). On different trials, participants were instructed to attend either to the artistic style of the painting or to the spatial layout of the room. Specifically, on “art” trials they were to identify paintings in different rooms that could have been painted by the same artist (i.e., that were similar in terms of style, use of color, and brushstrokes). On “room” trials they were to identify rooms that had the same spatial layout and furniture configuration, but viewed from different perspectives; all other visual features (e.g., wall color, visual details of the furniture) were different. Critically, identical stimuli were used in the “art” and “room” tasks across different trials, allowing us to isolate effects of top-down attention vs. bottom-up stimulation, and there were no overt demands on long-term memory.

These tasks were designed to tap into both the spatial and relational representations of the hippocampus. The art task required assessing stylistic relations between different paintings (e.g., how the color scheme of one painting relates to those of others). The room task required assessing spatial relations between different rooms (e.g., how the placement of

furniture and the wall angles in one room were similar to or different from those of other rooms). Thus, these tasks should strongly tax hippocampal representations: the art task requires relational processing, while the room task requires both relational and spatial processing.

If attention modulates the stability of hippocampal representations, there should be distinct activity patterns for the art and room attentional states: that is, activity patterns should be similar within a task and different between tasks. This hypothesis was supported by the data (Figure 3C; [50]): patterns of activity in each hippocampal subfield were more similar between trials of the same task (i.e., art trials compared to other art trials, and room trials compared to other room trials) vs. between trials of different tasks (i.e., art trials compared with room trials).

Moreover, this modulation of hippocampal pattern stability was dissociable from differences in the overall level of activity in the hippocampal subfields, and pattern stability but not overall activity was correlated with attentional behavior. Indeed, one region of interest in the hippocampus — comprising the CA2/3 and dentate gyrus (DG) subfields — was the only area in the brain where attentional states correlated with behavioral performance (Figure 3D). These data offer strong evidence that even with identical sensory input, hippocampal activity patterns can be reconfigured as a function of top-down attention to represent goal-relevant information, and that the hippocampal representation of this information can in turn relate to behavioral decisions.

This conclusion was subsequently bolstered by the finding that hippocampal representations of objects are dynamically modulated as attention is shifted to different goal-relevant stimulus features [51]. That is, as participants gradually acquired new conceptual knowledge about objects, hippocampal representations became configured to represent task-relevant features. When task rules changed, making once-relevant features now irrelevant, hippocampal representations dynamically reconfigured to represent the new goal-relevant features.

These studies therefore provide evidence that attention modulates the hippocampus by stabilizing representations of goal-relevant features. Attentional modulation of the hippocampus may also have relevance beyond in-the-moment attentional behavior, because attention robustly influences long-term memory: we are more likely to successfully remember information that was consistent with our attentional goals during encoding [12]. Could the effects of attention on what is represented in the hippocampus provide an explanation? Specifically, does the attentional state of the hippocampus — i.e., the extent to which it is configured for the task at hand — serve to prioritize goal-relevant information for encoding into long-term memory [41–43]?

This possibility finds suggestive support in animal models. For example, when space is task-relevant, place field stability in mice correlates with spatial memory [46]. Additionally, when rats engage in “attentive scanning” of a particular spatial location, place fields tend to form at that location on the next pass through it [52], an effect that bears resemblance to

rapid episodic encoding in humans. Thus, in animals, spatial attention modulates the stability of hippocampal spatial representations and predicts memory for spatial information.

Using fMRI, we explored whether attentional modulation of the stability of hippocampal representations predicts memory for goal-relevant information. We predicted that such effects should be observed in the CA2/3 and DG subfields of the hippocampus, given our previous finding that attentional modulation of these subfields predicted online attentional behavior [50], as well as other work relating activity and pattern similarity in these subfields to long-term memory encoding [53–55]. To test this hypothesis, we first identified attentional state representations in the hippocampus, and then examined whether more evidence for the task-relevant state during encoding predicted better behavioral memory for task-relevant information (Figure 4A).

During encoding of novel 3D-rendered rooms with paintings (as used in [50]), individuals were more likely to remember goal-relevant features of the image (i.e., paintings if the task required attention to art, and spatial layouts if the task required attention to rooms) when the attentional state of CA2/CA3/DG more closely resembled the goal-relevant vs. -irrelevant activity pattern [56]. That is, paintings were more likely to be remembered when the pattern of activity in CA2/CA3/DG during encoding more closely resembled the art attentional state than the room attentional state (defined independently earlier in the experiment). Conversely, room layouts were more likely to be remembered when CA2/CA3/DG activity patterns more closely resembled the room vs. the art state (Figure 4B–C). This effect was selective to the hippocampus, and not found in medial temporal lobe cortex or object- and scene-selective regions in temporal and parietal cortices. Thus, attention configures the hippocampus to prioritize information consistent with behavioral goals; this configuration is manifest as distinct patterns of activity for different attentional states; and the fidelity of the goal-relevant hippocampal state during encoding predicts memory for information consistent with that goal.

These studies show that the hippocampus is modulated by attention, and that this modulation is relevant for ongoing attentional behavior as well as the formation of long-term episodic memories. The relationship between the hippocampus and attention, however, is bidirectional: not only is the hippocampus influenced by attention, but hippocampal representations also influence attentional behaviors. For example, hippocampal memories (both explicit and implicit) can serve as a cue for visual search and affect how we move our eyes, even outside of awareness [12]. The mechanisms by which the hippocampus influences attentional orienting is not yet clear, but one possibility is that this depends on the interplay between the hippocampus, frontoparietal regions involved in attentional control, and feedback from these regions to sensory cortices. Indeed, there are several promising studies along these lines [51, 57–59]. This area of research holds exciting promise for future work.

Sensitivity to attention is just one instance of modulation of hippocampal function by behavioral goals. Another example is the ability of the hippocampus to represent future states, or predict upcoming events that are relevant for goal-directed behavior [60]. For example, hippocampal activity can represent the identity of a goal location [61] or proximity to the goal [62]. Sequential activity in hippocampal place cells can also represent future

navigational paths to a goal [63–64], even when the goal location was viewed but never explored [65]. Place cell activity at any given moment can also be influenced by actions taken or about to be taken [66]. Predictive coding in the hippocampus is not unique to the spatial domain: hippocampal anticipatory signals have also been found in statistical learning paradigms, where one object predicts another [67–70].

Another example of the modulation of hippocampal representations by goals is their sensitivity to reward, motivational states, and reward-based decision-making. For example, place fields tend to cluster around rewarded locations [71] and place cells fire when an animal waits in a goal location for delivery of reward [72]. The firing of hippocampal cells is also influenced by the motivational state of an animal (e.g., hunger vs. thirst), but only when the internal state can help guide behavior to the reward relevant for that state (i.e., food vs. water); thus, hippocampal coding of relationships between internal states and external rewards offers a means by which motivation can affect goal-directed behavior [73]. Hippocampal activity also reflects anticipation of an upcoming stimulus during reward-based decision-making [67] and can enable the construction of novel experiences and a judgment of their value [74].

Thus, although the hippocampus receives diverse input, the weighting of these inputs — i.e., its bias to represent or prioritize certain types of information — can be dynamically modulated based on behavioral goals. Thus, the hippocampus represents not just what is in the world, but rather the part of the world that we are attending to, motivated by, predicting, or acting on to reach a reward.

4. What computations enable hippocampal contribution to diverse functions?

So far, we have discussed some factors that influence hippocampal representations, but not the mechanisms by which those representations arise or how they are used to support diverse cognitive functions. We turn to these issues next, by examining case studies in attention and perception. We consider how domain-general computations can give rise to attentional and perceptual representations and how those representations are used for goal-directed behavior.

Attention

How does attention modulate the hippocampus, as in the examples above, and how does the hippocampus in turn guide attention? There are several potential mechanisms, all of which likely act in concert: neuromodulation, dynamic changes in functional connectivity, and hippocampal pattern completion.

First, neurotransmitter systems offer a potential route by which input to the hippocampus can be flexibly weighted to form state-dependent activity patterns in the hippocampus itself [21–23, 75]. For example, manipulations of dopaminergic [46] and cholinergic [76] systems alter hippocampal place field stability. Given that place field stability is also influenced by spatial attention [45–46], this raises the possibility that attention influences place field stability via these neuromodulatory systems. There are several reasons to expect such a

relationship, which we describe below, along with brief descriptions of the anatomy of the hippocampus and its neuromodulatory inputs.

Sensory information from the environment reaches the hippocampus via entorhinal cortex, which itself receives input from the perirhinal and parahippocampal cortices, among other regions [3]. The trisynaptic loop in the hippocampus consists of projections from entorhinal cortex to the dentate gyrus, from dentate gyrus to CA3, from CA3 to CA1, and from CA1 back to entorhinal cortex [3]. Additionally, entorhinal cortex projects directly to CA1. Neurons in CA3 are interconnected with one another, forming a recurrent network that is hypothesized to allow *pattern completion* — that is, the generation of a more complete representation of an event given only partial or degraded input as a cue [77–78].

The hippocampus receives cholinergic projections from the medial septum and vertical limb of the diagonal band of Broca [75]. Acetylcholine amplifies afferent signals into the hippocampus and suppresses excitatory recurrent connections in CA3 [75]. Such a mechanism might be the means by which information from the environment is up-weighted, and memory retrieval (via pattern completion in CA3) is down-weighted. In other words, this would shift the balance of external / internal signals in the hippocampus to prioritize representations of relevant information in the outside world. How such modulation can enable flexible attentional switching between different aspects of the external environment remains to be explored.

Another way in which neurotransmitter systems can affect hippocampal attention is via novelty- or reward-related dopaminergic modulation. Information that is novel or associated with reward should attract attention, either for further exploration or so that rewards can be obtained. CA1 is ideally situated to signal novelty because of its connectivity [21]. Information about the environment can reach CA1 directly from entorhinal cortex. Concurrently, CA1 receives input from CA3, which uses signals from entorhinal cortex and its internal pattern completion mechanisms to generate predictions about the world, given the input's similarity to stored representations. CA1 can therefore compare inputs (from entorhinal cortex) to expectations (from CA3), and generate a novelty response when those signals diverge. Detection of novelty in CA1 indirectly activates dopamine neurons in the ventral tegmental area, which in turn send dopaminergic projections to CA1, enhancing long-term potentiation [21]. Such plasticity may lead to the creation or maintenance of hippocampal representations of the novel, attended information. Moreover, these representations are then subject to modulation by motivational, contextual, or goal signals from prefrontal cortex. Thus, mutual interactions between the ventral tegmental area, hippocampus, and prefrontal cortex might be a means by which novelty and motivation influence attention and goal-directed behavior (though dopamine from the locus coeruleus may also play an important role [20, 24]).

Dynamic changes in functional connectivity — the influence of one brain region on another — may be another signature of attentional modulation [79]. For example, visual regions in ventral temporal cortex are more tightly coupled to perirhinal cortex when faces are attended, and more tightly coupled to parahippocampal cortex when scenes are attended [80]. Such switching of connectivity might be a general mechanism by which a given brain

region can represent different types of information depending on its relevance for goal-directed behavior. One possibility, for example, is that the hippocampus is more functionally connected to parahippocampal cortex when spatial or contextual information is task-relevant, and more functionally connected to perirhinal cortex when item or object information is task-relevant [81]. Attention may also dynamically modulate hippocampal connectivity with prefrontal cortex — either directly or indirectly via the thalamus — to up- or down-regulate the influence of particular goals or rules. Indeed, such a circuit — between prefrontal cortex, thalamus, and hippocampus — is necessary for goal-directed behavior and planning [18].

Finally, computations within the hippocampus itself can contribute to attentional effects. One such example, briefly discussed above, is pattern completion: the ability of the hippocampus to retrieve stored representations from degraded input [77–78]. This computation has been studied extensively in the long-term memory literature as a mechanism for retrieval: the hippocampus can reinstate the full pattern of activity for a memory when cued with only part of the memory. However, such a mechanism can also allow attention to configure the hippocampus. For example, one may learn that a particular behavior is rewarded given a particular cue — e.g., turning left at the intersection by your workplace starts you on a route that eventually ends in coffee. Subsequent encounters with the cue may then retrieve an attentional set that facilitates performing the behavior and receiving the reward — e.g., arriving at the intersection may cue attention to the spatial environment so that you can navigate to the café. Given that the attentional state of the hippocampus modulates how memories are stored — in some sense, the encoding state becomes part of the memory — retrieval may involve reinstating attentional state information. In other words, retrieval cues could direct behavior consistent with a goal state by reconfiguring the hippocampus to prioritize goal-relevant information [50, 56].

Consistent with these hypotheses, hippocampal processing can guide attention and eye movements to goal-relevant objects or locations. For example, memory for a goal location in a scene can aid attentional allocation to that location when the scene is presented again, and such memory-guided attention is mediated by the hippocampus [57–58]. Memory for associative or spatial information can also influence the way people move their eyes (and thus, their attention) to related objects and scenes, and this is also mediated by the hippocampus [82–84]. Thus, the same computations that allow the hippocampus to contribute to long-term memory may also support its role in attention-guided behavior.

Perception

A rich body of evidence suggests that the hippocampus contributes to high-level perceptual judgments, particularly perceptual discriminations of scene stimuli that have a high degree of feature overlap [13, 16, 28, 85]. We now consider which computations might enable the hippocampus to support perceptual decisions, including: encoding of relational representations, pattern separation, and match/mismatch detection.

A consistent finding in this literature is that the hippocampus is critical for perceptual discriminations based on the detection of relational differences in scenes, but it is not required when such judgments can be made on the basis of individual features or items

(Figure 5; [2, 13, 28, 85, 86]). Such relational representations are a hallmark of hippocampal mnemonic processing [31] and thus may also be leveraged to support online perceptual discriminations [13].

These relational representations may also support perceptual decisions at different levels of granularity. Spatial and mnemonic representations in anterior hippocampus (homologous to ventral hippocampus in rodents) are coarser than representations in posterior hippocampus (dorsal in rodents), which are more fine-grained [87]. For example, in rodents, place fields in ventral hippocampus are larger than those in dorsal hippocampus, and likewise, the scale of mnemonic representation changes in a similar gradient from anterior to posterior hippocampus in humans [88–89]. This is consistent with the finding that posterior/dorsal hippocampus is more important for fine-grained perceptual discriminations [90]. However, other results suggest that at a population level, fine-scaled information can emerge from individual cells with coarse coding [91]. Thus, the manner in which anterior vs. posterior hippocampus contribute to perceptual decisions at different scales remains to be elucidated.

Fine-scale discrimination is related to the ability of the hippocampus to form distinct representations for overlapping events — a process known as *pattern separation* [77–78]. This ability enables us to minimize interference when storing highly similar memories, such as where we parked our bike today vs. yesterday. Pattern separation can also enable the detection of subtle differences between highly similar percepts, for example, when the windows in a building are manipulated to be closer together in one photograph than another [85, 92]. Detection of such minute differences would benefit from the ability of the hippocampus to assign distinct representations to the two photographs despite their extensive overlap — in essence, they become different “items” to the hippocampus.

Finally, another hippocampal computation that can support perceptual decisions is the ability of the hippocampus — and most prominently the CA1 subfield — to act as a *comparator* that aids in match/mismatch detection [21, 93]. As mentioned earlier, the CA3 subfield is proposed to generate expectations by using environmental cues to retrieve related memories via pattern completion. The CA1 subfield receives these memory-based predictions, and, simultaneously, information about the current state of the world from entorhinal cortex. It is thus ideally situated to compare reality to expectations, and produce a signal based on how well they match or mismatch. This comparator function of CA1 has been found during long-term memory retrieval [93]. It may also enable the hippocampus to perform fine-grained perceptual discriminations [13, 94]. Indeed the continuous match/mismatch signals produced by such a comparator converge with the continuous hippocampal signals observed during perceptual decision-making [85, 94]. Thus, this provides another example of how a general hippocampal computation can contribute to multiple cognitive domains.

Before concluding, it is important to stress that we are not the first to discuss core hippocampal computations of pattern separation, pattern completion, relational binding, and match / mismatch detection [13, 15–16, 77–78]. The novel contribution of this paper is to make the argument that these core computations, in conjunction with diverse inputs to the hippocampus and flexible weighting of those inputs, enable the hippocampus to contribute

broadly to many cognitive functions. We also emphasize that communication with other brain regions and neuromodulatory influences play an important role (Section 4).

Although the focus of the current paper is the hippocampus, it is important to note that the mechanisms discussed above in the case studies of attention and perception are not unique to the hippocampus: Pattern separation, pattern completion, match/mismatch detection, and configural/relational representations also exist outside the hippocampus [95–97]. Moreover, mixed selectivity, diversity of input, and flexible weighting of input is not unique to the hippocampus — such properties are found in prefrontal cortex as well [10–11, 98–99]. We believe that these properties of the prefrontal cortex and hippocampus make them ideal for flexibly guiding online behavior using past experience or current attentional goals — and, importantly, these regions often work together to support these kinds of complex behaviors [19, 51]. Thus, the hippocampus is not the *only* region in the brain that is malleable, or that contributes to many cognitive functions. Rather, there is a hierarchy of malleability, and regions like the hippocampus and prefrontal cortex are at the top. Thus, our framework — that core computations operating on flexible and diverse inputs guide myriad cognitive functions — is not limited to the hippocampus, although we demonstrate the value of such a framework by using the hippocampus as a model system.

One question that arises, then, is: how and why is the hippocampus special? One possibility is that the hippocampus may be particularly good at rapid, one-shot learning [100]. That is, it may be able to form relational representations faster than other brain regions, and quickly modulate those representations when goals change [51]. Other brain regions may therefore depend on hippocampal representations when new learning must occur quickly. Under this view, the hippocampus is not qualitatively different from other high-order brain regions in terms of its malleability, but the difference arises in the rate of learning.

5. Conclusion

We propose that a core set of hippocampal computations may be deployed in different combinations, concurrently with flexible weighting of diverse inputs, to support various functions. Thus, the hippocampus is a malleable, configurable system, capable of prioritizing different types of information in support of behavioral goals.

Acknowledgments

Funding

This project was made possible through the support of the John Templeton Foundation and NIH (R01 EY021755). The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of these agencies.

References

1. Felleman DJ, Van Essen DC. Distributed hierarchical processing in the primate cerebral cortex. *Cer Cor.* 1991; 1:1–47.
2. Bussey TJ, Saksida LM. Object memory and perception in the medial temporal lobe: An alternative approach. *Curr Op Neurobiol.* 2005; 15:730–737. [PubMed: 16271459]

3. Lavenex P, Amaral DG. Hippocampal-neocortical interaction: A hierarchy of associativity. *Hippocampus*. 2000; 10:420–430. [PubMed: 10985281]
4. Hasson U, Chen J, Honey CJ. Hierarchical process memory: Memory as an integral component of information processing. *Trends Cog Sci*. 2015; 19:304–313.
5. Gavornik JP, Bear MF. Learned spatiotemporal sequence recognition and prediction in primary visual cortex. *Nat Neurosci*. 2014; 17:732–737. [PubMed: 24657967]
6. Kok P, Jehee JFM, de Lange FP. Less is more: Expectation sharpens representations in the primary visual cortex. *Neuron*. 2012; 75:265–270. [PubMed: 22841311]
7. Poort J, Khan AG, Pachitariu M, Nemri A, Orsolic I, et al. Learning enhances sensory and multiple non-sensory representations in primary visual cortex. *Neuron*. 2015; 86:1–13. [PubMed: 25856474]
8. Shuler MG, Bear MF. Reward timing in the primary visual cortex. *Science*. 2006; 311:1606–1609. [PubMed: 16543459]
9. Stanisor L, van der Togt C, Pennartz CMA, Roelfsema PR. A unified selection signal for attention and reward in primary visual cortex. *Proc Natl Acad Sci*. 2013; 110:9136–9141. [PubMed: 23676276]
10. Eichenbaum H. Barlow versus Hebb: When is it time to abandon the notion of feature detectors and adopt the cell assembly as the unit of cognition?. *Neurosci Lett*. 2017.
11. McKenzie S, Keene CS, Farovik A, Bladon J, Place R, Komorowski R, Eichenbaum H. Representation of memories in the cortical-hippocampal system: Results from the application of population similarity analyses. *Neurobiol Learn Mem*. 2016; 134:178–191. [PubMed: 26748022]
12. Aly M, Turk-Browne NB. How hippocampal memory shapes, and is shaped by, attention. In: Hannula DE, Duff MC, editors *The Hippocampus from Cells to Systems*. Springer International Publishing AG; Switzerland: 2017. 369–403.
13. Olsen RK, Moses SN, Riggs L, Ryan JD. The hippocampus supports multiple cognitive processes through relational binding and comparison. *Front Hum Neurosci*. 2012; 6:146–13. 1–13. [PubMed: 22661938]
14. Schiller D, Eichenbaum H, Buffalo EA, Davachi L, Foster DJ, Leutgeb S, Ranganath C. Memory and space: Towards an understanding of the cognitive map. *J Neurosci*. 2015; 35:13904–13911. [PubMed: 26468191]
15. Shohamy D, Turk-Browne NB. Mechanisms for widespread hippocampal involvement in cognition. *J Exp Psychol: Gen*. 2013; 142:1159–1170. [PubMed: 24246058]
16. Yonelinas AP. The hippocampus supports high-resolution binding in the service of perception, working memory, and long-term memory. *Beh Brain Res*. 2013; 254:34–44.
17. Rajasethupathy P, Sankaran S, Marshel JH, Kim CK, Ferenczi E, et al. Projections from neocortex mediate top-down control of memory retrieval. *Nature*. 2015; 526:653–659. [PubMed: 26436451]
18. Ito HT, Zhang SJ, Witter MP, Moser EI, Moser MB. A prefrontal-thalamo-hippocampal circuit for goal-directed spatial navigation. *Nature*. 2015; 522:50–55. [PubMed: 26017312]
19. Shapiro ML, Riceberg JS, Seip-Cammack K, Guise KG. Functional interactions of prefrontal cortex and hippocampus in learning and memory. In: Derdikman D, Knierim JJ, editors *Space, Time, and Memory in the Hippocampal Formation*. Springer; New York: 2014. 517–560.
20. Kempadoo KA, Mosharov EV, Choi SJ, Sulzer D, Kandel ER. Dopamine release from the locus coeruleus to the dorsal hippocampal promotes spatial learning and memory. *Proc Natl Acad Sci*. 2016; 113:14835–14840. [PubMed: 27930324]
21. Lisman JE, Grace AA. The hippocampal-VTA loop: Controlling the entry of information into long-term memory. *Neuron*. 2005; 46:703–713. [PubMed: 15924857]
22. Muzzio IA, Kentros C, Kandel E. What is remembered? Role of attention on the encoding and retrieval of hippocampal representations. *J Physiol*. 2009; 12:2837–2854.
23. Rowland DC, Kentros CG. Potential anatomical basis for attentional modulation of hippocampal neurons. *Ann NY Acad Sci*. 2008; 1129:213–224. [PubMed: 18591482]
24. Takeuchi T, Duzskiewicz AJ, Sonneborn A, Spooner PA, Yamasaki M, et al. Locus coeruleus and dopaminergic consolidation of everyday memory. *Nature*. 2016; 537:357–362. [PubMed: 27602521]

25. Backus AR, Bosch SE, Ekman M, Vicente-Grabovetsky A, Doeller CF. Mnemonic convergence in the human hippocampus. *Nat Comm.* 2016; 7:11991–9. 1–9.
26. Miši B, Goñi J, Betzel RF, Sporns O, McIntosh AR. A network convergence zone in the hippocampus. *PLoS Comp Biol.* 2014; 10:e1003982–10. 1–10.
27. Schedlbauer AM, Copara MS, Watrous AJ, Ekstrom AD. Multiple interaction brain areas underlie successful spatiotemporal memory retrieval in humans. *Scientific Reports.* 2014; 4:6431–9. 1–9. [PubMed: 25234342]
28. Lee ACH, Yeung LK, Barense MD. The hippocampus and visual perception. *Front Hum Neurosci.* 2012; 6:91–17. 1–17. [PubMed: 22529794]
29. Maguire EA, Mullally SL. The hippocampus: A manifesto for change. *J Exp Psychol: Gen.* 2013; 142:1180–1189. [PubMed: 23855494]
30. Nadel L, Peterson MA. The hippocampus: Part of an interactive posterior representational system spanning perceptual and memorial systems. *J Exp Psychol: Gen.* 2013; 142:1242–1254. [PubMed: 23895347]
31. Eichenbaum H, Cohen NJ. Can we reconcile the declarative memory and spatial navigation views on hippocampal function? *Neuron.* 2014; 83:764–770. [PubMed: 25144874]
32. Kumaran D, Hassabis D, McClelland JL. What learning systems do intelligent agents need? Complementary learning systems theory updated. *Trends Cog Sci.* 2016; 20:512–534.
33. Moses SN, Ryan JD. A comparison and evaluation of the predictions of relational and conjunctive accounts of hippocampal function. *Hippocampus.* 2006; 16:43–65. [PubMed: 16270317]
34. Horner AJ, Bisby JA, Bush D, Lin WJ, Burgess N. Evidence for holistic episodic recollection via hippocampal pattern completion. *Nat Comm.* 2015; 6:7462–11. 1–11.
35. Mack ML, Preston AR. Decisions about the past are guided by reinstatement of specific memories in the hippocampus and perirhinal cortex. *NeuroImage.* 2016; 127:144–157. [PubMed: 26702775]
36. Staresina BP, Michelmann S, Bonnefond M, Jensen O, Axmacher N, Fell J. Hippocampal pattern completion is linked to gamma power increases and alpha power decreases during recollection. *eLife.* 2016; 5:e17397–18. 1–18. [PubMed: 27508355]
37. Tompary A, Duncan K, Davachi L. High-resolution investigation of memory-specific reinstatement in the hippocampus and perirhinal cortex. *Hippocampus.* 2016; 26:995–1007. [PubMed: 26972485]
38. Kastner S, Ungerleider LG. Mechanisms of visual attention in the human cortex. *Annu Rev Neurosci.* 2000; 23:315–341. [PubMed: 10845067]
39. Dudukovic NM, Preston AR, Archie JJ, Glover GH, Wagner AD. High-resolution fMRI reveals match enhancement and attentional modulation in the human medial temporal lobe. *J Cogn Neurosci.* 2010; 23:670–682. [PubMed: 20433244]
40. Yamaguchi S, Hale LA, D’Esposito M, Knight RT. Rapid prefrontal-hippocampal habituation to novel events. *J Neurosci.* 2004; 24:5356–5363. [PubMed: 15190108]
41. Carr VA, Engel SA, Knowlton BJ. Top-down modulation of hippocampal encoding activity as measured by high-resolution functional MRI. *Neuropsychologia.* 51:1829–1837.
42. Davachi L, Wagner AD. Hippocampal contributions to episodic encoding: Insights from relational and item-based learning. *J Neurophysiol.* 2002; 88:982–990. [PubMed: 12163547]
43. Uncapher MR, Rugg MD. Selecting for memory? The influence of selective attention on the mnemonic binding of contextual information. *J Neurosci.* 2009; 29:8270–8279. [PubMed: 19553466]
44. Kelemen E, Fenton AA. Coordinating different representations in the hippocampus. *Neurobiol Learn Mem.* 2016; 129:50–59. [PubMed: 26748023]
45. Muzzio IA, Levita L, Kulkarni J, Monaco J, Kentros C, Stead M, Abbott LF, Kandel ER. Attention enhances the retrieval and stability of visuospatial and olfactory representations in the dorsal hippocampus. *PLoS Biol.* 2009; 7:e1000140–20. 1–20. [PubMed: 19564903]
46. Kentros CG, Agnihotri NT, Streater S, Hawkins RD, Kandel ER. Increased attention to spatial context increases both place field stability and spatial memory. *Neuron.* 2004; 42:283–295. [PubMed: 15091343]

47. Jackson J, Redish AD. Network dynamics of hippocampal cell-assemblies resemble multiple spatial maps within single tasks. *Hippocampus*. 2007; 17:1209–1229. [PubMed: 17764083]
48. Kelemen E, Fenton AA. Dynamic grouping of hippocampal neural activity during cognitive control of two spatial frames. *PLoS Biol*. 2010; 8:e1000403–14. 1–14. [PubMed: 20585373]
49. Fenton AA, Lytton WW, Barry JM, Lenck-Santini PP, Zinyuk LE, et al. Attention-like modulation of hippocampus place cell discharge. *J Neurosci*. 2010; 30:4613–4625. [PubMed: 20357112]
50. Aly M, Turk-Browne NB. Attention stabilizes representations in the human hippocampus. *Cer Cor*. 2016; 26:783–796.
51. Mack ML, Love BC, Preston AR. Dynamic updating of hippocampal object representations reflects new conceptual knowledge. *Proc Natl Acad Sci*. 2016; 113:13203–13208. [PubMed: 27803320]
52. Monaco JD, Rao G, Roth ED, Knierim JJ. Attentive scanning behavior drives one-trial potentiation of hippocampal place fields. *Nat Neurosci*. 2014; 17:725–731. [PubMed: 24686786]
53. Suthana N, Ekstrom A, Moshirvaziri S, Knowlton B, Bookheimer S. Dissociations within human hippocampal subregions during encoding and retrieval of spatial information. *Hippocampus*. 2011; 21:694–701. [PubMed: 20882543]
54. Wolosin SM, Zeithamova D, Preston AR. Distributed hippocampal patterns that discriminate reward context are associated with enhanced associative binding. *J Exp Psychol: Gen*. 2013; 142:1264–1276. [PubMed: 23834024]
55. Zeineh MM, Engel SA, Thompson PM, Bookheimer SY. Dynamics of the hippocampus during encoding and retrieval of face-name pairs. *Science*. 2003; 299:577–580. [PubMed: 12543980]
56. Aly M, Turk-Browne NB. Attention promotes episodic encoding by stabilizing hippocampal representations. *Proc Natl Acad Sci*. 2016; 113:E420–E429. [PubMed: 26755611]
57. Summerfield JJ, Lepsien J, Gitelman DR, Mesulam MM, Nobre AC. Orienting attention based on long-term memory experience. *Neuron*. 2006; 49:905–916. [PubMed: 16543137]
58. Stokes MG, Atherton K, Patai EZ, Nobre AC. Long-term memory prepares neural activity for perception. *Proc Natl Acad Sci*. 2012; 109:E360–E367. [PubMed: 22109554]
59. Voss JL, Cohen NJ. Hippocampal-cortical contributions to strategic exploration during perceptual discrimination. *Hippocampus*. 2017; 27:642–652. [PubMed: 28241401]
60. Buckner RL. The role of the hippocampus in prediction and imagination. *Ann Rev Psychol*. 2010; 61:27–48. [PubMed: 19958178]
61. Brown TI, Carr VA, LaRocque KF, Favila SE, Gordon AM, et al. Prospective representation of navigational goals in the human hippocampus. *Science*. 2016; 352:1323–1326. [PubMed: 27284194]
62. Balaguer J, Spiers H, Hassabis D, Summerfield C. Neural mechanisms of hierarchical planning in a virtual subway network. *Neuron*. 2016; 90:893–903. [PubMed: 27196978]
63. Johnson A, Redish AD. Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J Neurosci*. 2007; 27:12176–12189. [PubMed: 17989284]
64. Pfeiffer BE, Foster DJ. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*. 2013; 497:74–81. [PubMed: 23594744]
65. Ólafsdóttir HF, Barry C, Saleem AB, Hassabis D, Spiers HJ. Hippocampal place cells construct reward related sequences through unexplored space. *eLife*. 2015; 4:e06063–17. 1–17. [PubMed: 26112828]
66. Dudchenko PA, Wood ER. Splitter cells: Hippocampal place cells whose firing is modulated by where the animal is going or where it has been. In: Derdikman D, Knierim JJ, editors *Space, Time, and Memory in the Hippocampal Formation*. Springer; New York: 2014. 253–272.
67. Bornstein AM, Daw ND. Cortical and hippocampal correlates of deliberation during model-based decisions for rewards in humans. *PLoS Comp Biol*. 2013; 9:e1003387–19. 1–19.
68. Hindy NC, Ng FY, Turk-Browne NB. Linking pattern completion in the hippocampus to predictive coding in visual cortex. *Nat Neurosci*. 2016; 19:665–667. [PubMed: 27065363]
69. Schapiro AC, Kustner LV, Turk-Browne NB. Shaping of object representations in the human medial temporal lobe based on temporal regularities. *Curr Biol*. 2012; 22:1622–1627. [PubMed: 22885059]

70. Turk-Browne NB, Scholl BJ, Johnson MK, Chun MM. Implicit perceptual anticipation triggered by statistical learning. *J Neurosci*. 2010; 30:11177–11187. [PubMed: 20720125]
71. Hollup SA, Molden S, Donnett JG, Moser MB, Moser EI. Accumulation of hippocampal place fields at the goal location in an annular watermaze task. *J Neurosci*. 2001; 21:1635–1644. [PubMed: 11222654]
72. Hok V, Lenck-Santini PP, Roux S, Save E, Muller RU, Poucet B. Goal-related activity in hippocampal place cells. *J Neurosci*. 2007; 27:472–482. [PubMed: 17234580]
73. Kennedy PJ, Shapiro ML. Motivational states activate distinct hippocampal representations to guide goal-directed behaviors. *Proc Natl Acad Sci*. 2009; 106:10805–10810. [PubMed: 19528659]
74. Barron HC, Dolan RJ, Behrens TEJ. Online evaluation of novel choices by simultaneous representation of multiple memories. *Nat Neurosci*. 2013; 16:1492–1498. [PubMed: 24013592]
75. Newman EL, Gupta K, Climer JR, Monaghan CK, Hasselmo ME. Cholinergic modulation of cognitive processing: Insights drawn from computational models. *Front Beh Neurosci*. 2012; 6:24–19. 1–19.
76. Brazhnik ES, Muller RU, Fox SE. Muscarinic blockade slows and degrades the location-specific firing of hippocampal pyramidal cells. *J Neurosci*. 2003; 23:611–621. [PubMed: 12533621]
77. Knierim JJ, Neunuebel JP. Tracking the flow of hippocampal computation: Pattern separation, pattern completion, and attractor dynamics. *Neurobiol Learn Mem*. 2016; 129:38–49. [PubMed: 26514299]
78. Yassa MA, Stark CEL. Pattern separation in the hippocampus. *Trends Neurosci*. 2011; 34:515–525. [PubMed: 21788086]
79. Turk-Browne NB. Functional interactions as big data in the human brain. *Science*. 2013; 342:580–584. [PubMed: 24179218]
80. Córdova NI, Tompary A, Turk-Browne NB. Attentional modulation of background connectivity between ventral visual cortex and the medial temporal lobe. *Neurobiol Learn Mem*.
81. Ranganath C, Ritchey M. Two cortical systems for memory-guided behavior. *Nat Rev Neurosci*. 2012; 13:713–726. [PubMed: 22992647]
82. Hannula DE, Ranganath C. The eyes have it: Hippocampal activity predicts expression of memory in eye movements. *Neuron*. 2009; 63:592–599. [PubMed: 19755103]
83. Hannula DE, Ryan JD, Tranel D, Cohen NJ. Rapid onset relational memory effects are evident in eye movement behavior, but not in hippocampal amnesia. *J Cogn Neurosci*. 2007; 19:1690–1705. [PubMed: 17854282]
84. Ryals AJ, Wang JX, Polnaszek KL, Voss JL. Hippocampal contribution to implicit configuration memory expressed via eye movements during scene exploration. *Hippocampus*. 2015; 25:1028–1041. [PubMed: 25620526]
85. Aly M, Ranganath C, Yonelinas AP. Detecting changes in scenes: The hippocampus is critical for strength-based perception. *Neuron*. 2013; 78:1127–1137. [PubMed: 23791201]
86. Baxter MG. Involvement of medial temporal lobe structures in memory and perception. *Neuron*. 2009; 61:667–677. [PubMed: 19285463]
87. Poppenk J, Evensmoen HR, Moscovitch M, Nadel L. Long-axis specialization of the human hippocampus. *Trends Cog Sci*. 2013; 17:230–240.
88. Collin SHP, Milivojevic B, Doeller CF. Memory hierarchies map onto the hippocampal long axis in hippocampus. *Nat Neurosci*. 2015; 18:1562–1564. [PubMed: 26479587]
89. Strange BA, Witter MP, Lein ES, Moser EI. Functional organization of the hippocampal longitudinal axis. *Nat Rev Neurosci*. 2014; 15:655–669. [PubMed: 25234264]
90. McTighe SM, Mar AC, Romberg C, Bussey TJ, Sakisda LM. A new touchscreen test of pattern separation: Effect of hippocampal lesions. *NeuroReport*. 2009; 20:881–885. [PubMed: 19421077]
91. Keinath AT, Wang ME, Wann EG, Yuan RK, Dudman JT, Muzzio IA. Precise spatial coding is preserved along the longitudinal hippocampal axis. *Hippocampus*. 2014; 24:1533–1548. [PubMed: 25045084]
92. Aly M, Yonelinas AP. Bridging consciousness and cognition in memory and perception: Evidence for both state and strength processes. *PLoS ONE*. 2012; 7:e30231: 1–16. [PubMed: 22272314]

93. Duncan K, Ketz N, Inati SJ, Davachi L. Evidence for area CA1 as a match/mismatch detector: A high-resolution fMRI study of the human hippocampus. *Hippocampus*. 2012; 22:389–398. [PubMed: 21484934]
94. Elfman KW, Aly M, Yonelinas AP. Neurocomputational account of memory and perception: Thresholded and graded signals in the hippocampus. *Hippocampus*. 2014; 24:1672–1686. [PubMed: 25112784]
95. Carrillo-Rein L, Yang W, Bando Y, Peterka DS, Yuste R. Imprinting and recalling cortical ensembles. *Science*. 2016; 353:691–694. [PubMed: 27516599]
96. Kent BA, Hvoslef-Eide M, Saksida LM, Bussey TJ. The representational-hierarchical view of pattern separation: Not just hippocampus, not just space, not just memory? *Neurobiol. Learn Mem.* 2016; 129:99–106.
97. Nortmann N, Rekauzke S, Onat S, König P, Jancke D. Primary visual cortex represents the difference between past and present. *Cer Cor.* 2015; 25:1427–1440.
98. Miller EK, Cohen JD. An integrative theory of prefrontal cortex function. *Ann Rev Neurosci.* 2001; 24:167–202. [PubMed: 11283309]
99. Rigotti M, Barak O, Warden MR, Wang X-J, Daw ND, Miller EK, Fusi S. The importance of mixed selectivity in complex cognitive tasks. *Nature*. 2013; 497:585–590. [PubMed: 23685452]
100. Kumaran D, Hassabis D, McClelland JL. What learning systems do intelligent agents need? Complementary learning systems theory updated. *Trends Cog Sci.* 2016; 20:512–534.

Highlights

- The hippocampus contributes to many domains of cognition beyond long-term memory
- This results from a diversity of anatomical inputs and the flexibility of their weighting
- A core set of computations performed on these weighted inputs support its broad role

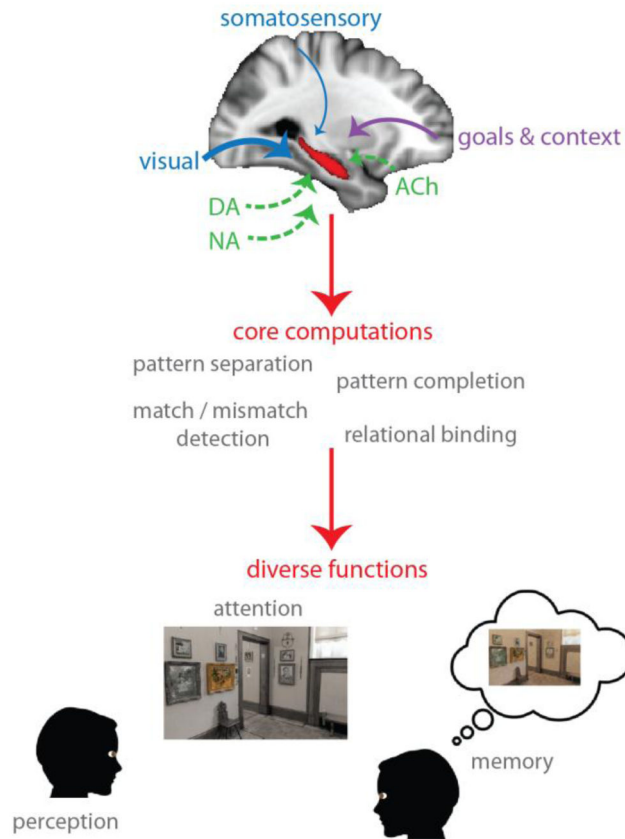


Figure 1.

Top: Multiple sensory modalities, including vision, audition, somatosensation, and olfaction, converge on the hippocampus. These inputs can be flexibly weighted based on behavioral goals and task context, which themselves are represented elsewhere, such as in frontoparietal cortex. In this example, visual signals are up-weighted (thicker arrow) while somatosensory signals are down-weighted (thinner arrow). Neuromodulatory systems, including dopaminergic (DA), cholinergic (ACh), and noradrenergic (NA) systems, can bias this flow of information and local processing. *Middle:* The hippocampus performs a core set of domain-general computations. *Bottom:* Flexibly weighted inputs, combined with some or all of these computations, enable the hippocampus to contribute to various cognitive functions.



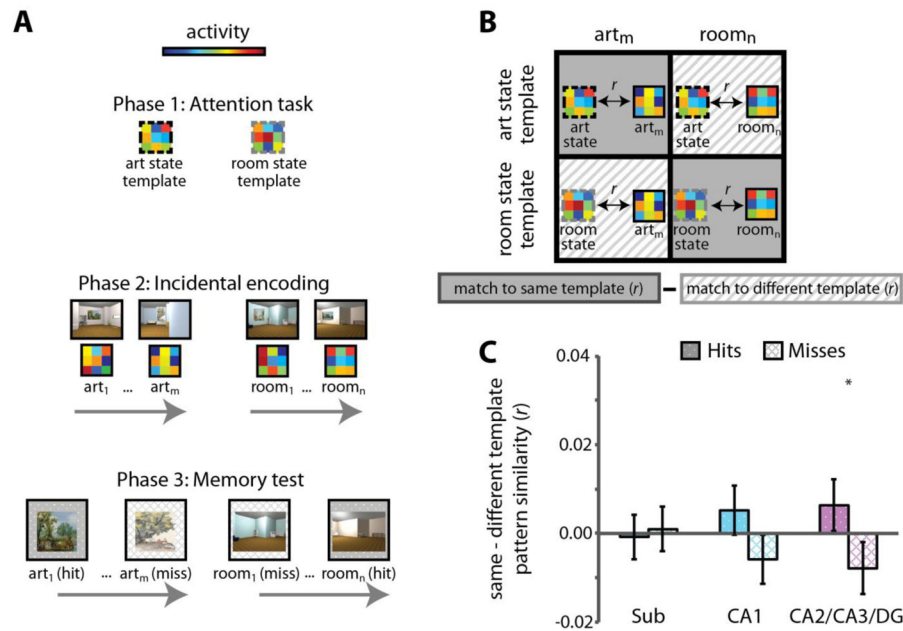


Figure 2.

Left: One of the most common neural measures of attention is the modulation of overall levels of activity in selective brain regions, whether firing rates in single neurons or BOLD signal in fMRI. *Right:* An alternative measure of attention is pattern stability — similar patterns of activity in a brain region across two or more instances of the same attentional state, and more distinct patterns of activity for different attentional states. The colored squares indicate the level of activity in nine voxels (or nine neurons).

**Figure 3.**

Attention stabilizes patterns of activity in the hippocampus. **(A)** While undergoing fMRI, participants viewed 3D-rendered rooms, each with a unique layout (furniture configuration, wall angles) and painting. For each “base image”, an “art match” was a room containing a different painting that was painted by the same artist as the painting in the base image, and a “room match” was a room that had the same spatial layout as the room in the base image viewed from a different perspective. **(B)** Participants were cued on each trial to attend to the style of the paintings (ART) or the layout of the rooms (ROOM). They then viewed a base image followed by a search set of four images. On art trials, they had to search for an art match. On room trials, they had to search for a room match. Finally, they were probed as to whether they had seen an art or a room match (the probe usually but not always matched the cue), and had to respond yes or no. **(C)** Activity patterns in each hippocampal subfield region of interest were more highly correlated for trials of the same attentional state (i.e., art/art and room/room) than trials of different attentional states (i.e., art/room), thus showing state-dependent patterns of activity. **(D)** In CA2/CA3/DG, individual differences in the stability of activity patterns for the room attentional state predicted behavior on the room task. $^{***} p < .001$. Figure adapted from [50].

**Figure 4.**

The fidelity of hippocampal attentional states during encoding of novel information predicts goal-directed memory. **(A)** While undergoing high-resolution fMRI, individuals performed the “art gallery” task used in [50] and depicted in Figure 3. This part of the experiment (Phase 1) was used to extract activity patterns in the hippocampus related to attending to paintings and attending to room layouts — i.e., attentional state “templates” for each task. Participants then encoded a novel set of 3D-rendered rooms in Phase 2, attending to art in one block and room layouts in the other. Finally, in Phase 3, they were taken out of the scanner and tested on their memory for the goal-relevant aspect of images from Phase 2: paintings from the art encoding task and room layouts from the room encoding task. **(B)** To investigate how hippocampal attentional states related to memory formation, the activity pattern for each Phase 2 encoding trial was correlated with the task-relevant and task-irrelevant attentional state templates, defined from Phase 1. For example, the activity pattern for an art encoding trial would be correlated with the art state template (task-relevant) and the room state template (task-irrelevant). The difference between these pattern similarity correlations provided a measure of the extent to which the hippocampus was in the correct attentional state during encoding. **(C)** In CA2/CA3/DG (and in the hippocampus treated as a single region of interest), activity patterns more closely resembled the task-relevant vs. task-irrelevant attentional state during successful vs. unsuccessful encoding. * $p < .05$. Figure adapted from [56].

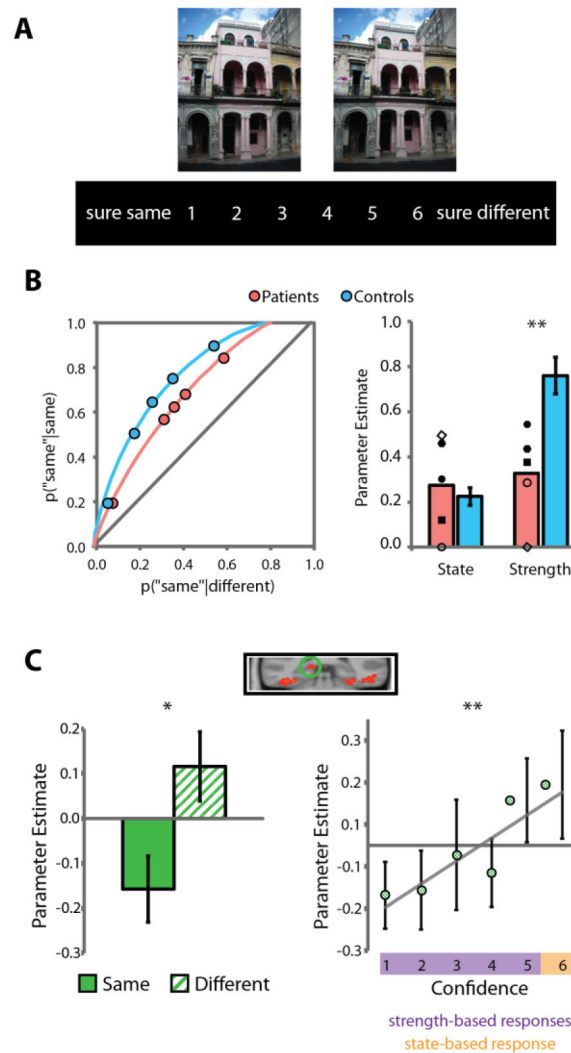


Figure 5. (A) Participants viewed pairs of scenes (presented simultaneously for the patient study shown in [B] and sequentially for the fMRI study shown in [C]) and reported their confidence that the two scenes were the same or different using a 1–6 scale. (B) Analysis of receiver-operating characteristics (ROCs) of patients with medial temporal lobe damage and healthy controls indicated that patients were selectively impaired on assessing the overall relational match of scenes (“strength-based perception”, measured by the curvature of the ROC) and were spared on perceptual judgments related to identifying specific feature-level differences (“state-based perception”, measured by the upper x-intercept of ROC). Patients with selective hippocampal damage are depicted with filled shapes in the bar plot; patients with more extensive MTL lesions are depicted with open shapes. The same results hold if only selective hippocampal lesion patients are considered. (C) fMRI in healthy adults showed greater hippocampal activity on trials with different vs. same scene pairs, and hippocampal activity continuously scaled with the strength of strength-based perceptual judgments. * $p < .05$ ** $p < .01$. Figure adapted from [85].